IPIC 2026

Conference Proceedings





12th International Physical Internet Conference

8-10 June 2026 | Bordeaux, France

SUPPLY CHAIN AND LOGISTICS INTELLIGENCE IN THE ERA OF THE PHYSICAL INTERNET: BRIDGING HIGH-TECH AND LOW-TECH SOLUTIONS

Conference Hosts

KEDGE BUSINESS SCHOOL

alice | Alliance for Logistics Innovation through Collaboration in Europe

IMT Mines Albi-Carmaux
École Mines-Télécom

IPIC 2026 Proceedings

Conference Paper Table of Contents

Submission 1. **Engineering a Physical Internet Concept for Food and Beverage Distribution in Sparse Geographies Using Network Optimization and Game-Theory** 9
Uday Venkatadri

Submission 3. **Dynamic Slot Reallocation for Physical Internet Hubs** 20
Anshul Vijay, Farzaneh Zirak, Russell Thompson and Neema Nassir

Submission 4. **Optimization and Resilience Assessment of Multi-Agent Delivery Using Evolutionary Neural Networks in the Context of Physical Internet** 30
Hisatoshi Naganawa and Enna Hirata

Submission 5. **Genetic Algorithm-Based Configuration of Logistics Hubs for Parcel Delivery: A Case Study of the Tokyo Metropolitan Area** 40
Muhammad Ilham Fahreza, Sinar Noe Corzo-Garcia, Masoud Kahalimoghadam, Hiroki Sakai, Greg Foliente, Russell G. Thompson and Enna Hirata

Submission 7. **The Strategic Value of AI-Driven Logistics Data Products** 50
Shulin He, Mengdi Zhang and George Q. Huang

Submission 8. **Promoting ESG in Energy Supply Chain with Digital Twin and Private Equity: An Evolutionary Game Analysis** 53
Shengan Yu, Mengdi Zhang, Zhiheng Zhao and George Q. Huang

Submission 9. **Optimization of Dynamically Hyperconnected Network Services for Heterogeneous Clients** 57
He Junkai, Labarthe Olivier, Wang Zhixin and Ding Chen

Submission 10. **PIMM: A Maturity Model for Advancing the Physical Internet from Concept to Practice** 67
Tadashi Mizutani, Yuki Kurakake, Kei Yazaki, Maki Aoki, Tomohiro Okuzumi and Satoshi Kagino

Submission 11. **A Generalized Requirement Engineering Application for PI-Hubs: Integrating Triple Bottom Line and Physical Internet Fundamental** 76
Monica-Juliana Perez, Tarik Chargui and Damien Trentesaux

Submission 13. **Towards Competitive Inland Waterway Logistics: Virtual Models for Modal Shift Assessment** 86
José Luis López and David Ciprés

Submission 14. **Urban Ride-Sharing as a Systemic Mobility Solution: Towards a Physical Internet of City Transport** 96
Maria Matusiewicz

Submission 15. **Implementing the Physical Internet in Multimodal Hubs: The AUTOSUP Digital Twin-Based Decision-Support Approach** 112
Aristea Maria Zafeiropoulo, Konstantinos Christidis, Filippos Adamidis, Athanasios Karydis, Babis Magoutas, Antonis Mygiakis and Ilias Gkotsis

Submission 16. **Enabling the Physical Internet Through AI-Driven Federated Interoperability: An Ontology Matching Perspective** 122
Leonardo Daou, Zhuoxun Zheng, Syed Juned Ali, Dominik Bork, Eva Petitdemange, Gregory Zacharewicz, Séverine Durieux and Nicolas Daclin

Submission 17. **Design Objectives and Requirements for a PI Dynamic Interoperability Tool** 132
Gero Niemann and Rod Franklin

Submission 19. **Heterogeneous Graph-Based Risk Assessment for Industrial Chain Cooperation** 142
Zihang Yu, Hanlin Tang, Dong Liu and Shuangxi Huang

Submission 20. **Bridging the Gap Between Physical Internet Vision and Urban Logistics Operations Through DISCO Living Labs** 155
Vasileios Giannoudis, Dimos Theodoros Touloumidis, Georgia Ayfantopoulou, Paola Cossu and Paola Astegiano

Submission 21. **Autonomous Edge-Integrated Digital Twins for Real-Time Physical Internet Optimization: A Multi-Agent Reinforcement Learning and Blockchain Framework** 163
Umesh Gangadhar, J Singh and P Kumar

Submission 22. **Contextual Optimization of Collaboration Openness in Physical Internet Systems** 165
Nafe Moradkhani, Olivier Labarthe, Matthieu Lauras, Yann Bouchery and Walid Klibi

Submission 23. **Digitalizing a Territorial Fresh-Produce Wholesale Market as a Physical-Internet Node: Container-Centric Traceability, Evidence Packages, and Controlled Transparency at Mercato Ortofrutticolo del Roero** 179
Claudio Salvadori, Franco Maciariello, Eddi Lorenzo, Massimo Torchio and Silvio Abrate

Submission 24. **Order Dispatching for On-Demand Delivery with Elevator Waiting Time** 189
Haitao Liu, Mouna Bamoumen, Juan-Carlos Piña Pardo, Zhaoxia Guo and Jan Fransoo

Submission 25. **The TREATMENT Project: A Physical-Internet-enabled Reference Architecture for End-to-End Visibility in Direct-to-Patient Clinical Trial Logistics** 198
Franco Maciariello, Eddi Lorenzo, Cristiana Caponio, Giuseppe Tortora and Claudio Salvadori

Submission 26. **Open Distribution Networks under Compound Stress: A Multi-Agent Stress-Testing Framework** 208
Mahmut Metin Inan, Ashwin Pothen, Matthieu Lauras and Benoit Montreuil

Submission 27. **Interactive Service and Blockchain Traceability in a Cyber Physical Internet Service Network** 218
Jialing Lin, Mengdi Zhang, Zhiheng Zhao and George Q. Huang

Submission 28. **Enabling Durable Horizontal Multilateral Collaboration by Integrating Collaborative Business Modelling with Multi-Party Optimization** 227
Ruben Fransen, Tamara Oukes, Jaco van Meijeren and Marijn van Adrichem

Submission 29. **Physics of Decision: A Data-Driven Discovery of Physical Internet Dynamics** 237
Benoit Morvan, Nafe Moradkhani, Victor Romero, Guillaume Martin, Benoit Montreuil and Frederick Benaben

Submission 30. **Optimizing Container Capacity for Light Cargo Parcel Delivery** 248
Luis Marques, Yann Bouchery, Olivier Labarthe and Nafe Moradkhani

Submission 31. **From Fragmented Digitalization to Physical Internet Coordination: Achieving Concrete Data Interoperability** 259
Liz Araceli Cristaldo and Marion Cottet

Submission 32. **Multimodal Itinerary Optimization for PI-Container Networks: Leveraging Inland Waterways for Sustainable Freight Consolidation** 269
Gabriel Zambrano Rey, Monica-Juliana Perez, Tarik Chargui, Yves Sallez and Damien Trentesaux

Submission 33. **From Digital to Physical Internet: CBoXX-Enabled LTL-on-Rail Demonstrators with Automated PI Nodes and Slot-Based Booking** 280
Hans-Juergen Weidemann

Submission 34. **Investable Interoperability for the PI: Participant-Centricity and a Viable Market for Common Services** 290
Inge Lucassen, Huib van Nieuwenhuijze and Harrie Bastiaansen

Submission 36. **Improving Transport Utilization via Proactive Consolidation and Coordinated Purchasing** 300
Najat Bara, Mohamed Nour Kitri, Onur Ozturk, Frédéric Gautier and Junkai He

Submission 37. **Adaptive Capacity Orchestration in Physical Internet Parcel Networks via Resource Pooling** 308
Yaxin Pan, Walid Klibi, Matthieu Lauras, Eva Petitdemange and Johan Leveque

Submission 39. **Towards Federated Routing Optimization for Road Transport** 318
Renaud De Landtsheer and Quentin Meurisse

Submission 40. **Policy-Based and Process-Aware Interoperability in the Physical Internet** 328
Philippe Michiels, Julián Rojas and Birger Schrevens

Submission 41. **Roadmap of Logistics Data Platforms for the Physical Internet: A DTLF-Aligned Assessment of Open Standard Data Exchange, Federated Data Spaces, and Governance for Adoption** 338
Paulo Cantiliano-Lizana, Gloria Giammei and Matteo Magnani

Submission 42. **H-CPI-ARP: A Hierarchical Proxy Mechanism Bridging High-Tech Hubs and Low-Tech Logistics Units** 349
Wenqing Lei, Yong Hong Kuo and George Q. Huang

Submission 43. **Operationalising Physical Internet Principles in Neighbourhood Urban Logistics: The IKIGAI Hybrid Fixed–Mobile Micro-Hub Living Lab** 359
Stefania Pasavento, Dimos Touloumidis and Sofoklis Dais

Submission 44. **Rethinking Blockchain Consensus Mechanisms: A Supply Chain Perspective Through the Physical Internet** 369
Imene Yasmine Benyelles, Sonia Mahjoub and Hicham Abbad

Submission 45. **A Physical Internet Approach to Disaster Management: Leveraging Open Distribution Networks for Humanitarian Supply Chains** 379
Mahmut Metin Inan, Matthieu Lauras and Benoit Montreuil

Submission 46. **A Roadmap-Aligned Decision-Support Methodology to Identify Scaling Mechanisms for Physical Internet-Compliant Logistics Innovations: The IKIGAI Approach** 389
Filippo Mauro, Fabio Cartolano, Paola Astegiano, Paola Cossu, Dimos Touloumidis, Sofoklis Dais, Georgia Ayfantopoulou, Alkis Pitelis and Ioanna Fergadiotou

Submission 47. **Shipment Planning Protocols in Cyber-Physical Internet for B2C Cross-Border E-Commerce Logistics** 399
Xinyue Xie, Xuegang Li, Bo Liu, Zhiheng Zhao and George Q. Huang

Submission 48. **ISO 14083 Guidance for Comparable CO_2e Reporting in Urban Road Freight Operations** 400
Dimos Touloumidis, Sofoklis Dais and Georgia Ayfantopoulou

Submission 50. **Look Around Cyber-Physical Internet Port Operations** 410
Yi You and George Q. Huang

Submission 51. **End-to-End Supply Chain Management in Plant Construction: A Simulation-Based, Data Space-Ready Approach** 411
Lucas Schreiber, Till Lemmer, Susanne Klöcker, Julia Wingenfeld and Michael Dominik Görtz

Submission 52. **Optimizing Physical Internet Container Circulation under Demand and Return Uncertainty: A Case Study of Global Rubber Shipment** 421
Yunyu Zhang, Song Huang, Fangchun Peng, Dong Yang, Saijun Shao, Ming Li and Zhiyuan Ouyang

Submission 53. **Modularization and Containerized Consolidation for Multimodal Three Tier City Logistics** 422
Najib Errami, Seyyed-Ehsan Hashemi-Petroodi, Jakob Puchinger and Walid Klibi

Submission 54. **Constraint-Based Spatial-Temporal Alignment for Traceable Modular Integrated Construction Logistics in a Cyber-Physical Internet Framework** 432
Pengjun Yue, Zhiheng Zhao and George Q. Huang

Submission 55. **Open Parcel Lockers as Physical Internet Nodes: A Framework for Accelerated and Sustainable Last-Mile Delivery** 442
Ola Qasseer, Péter Bajor and Norina Szander

Submission 56. **Proposed Framework for Modular Container Usage Estimation Based on Regional Contextualization** 453
Priyali Bandla, Yinzhu Quan, Praveen Muthukrishnan and Benoit Montreuil

Submission 57. **Automatic Risk-Conscious Hyperconnected Logistics Location Selection via LLM-Based Multi-Agent Systems** 464
Yinzhu Quan, Frederick Benaben and Benoit Montreuil

Submission 58. **A Generic Meta-Model-Driven Simulation Engine for Configurable Supply Chain Structures** 474
Camélia Belkepeau, Camille Durthaller Renard, Benoit Morvan, Victor Romero, Guillaume Martin, Benoit Montreuil and Frédéric Benaben

Submission 59. **Implementing the Physical Internet: A Digital Twin Architecture for Collaborative and Cost-Aware Freight Networks** 484
Masoud Kahalimoghadam, Russell Thompson, Lele Zhang and Michael Kirley

Submission 60. **Interoperable π-Containers: Physical-Internet Network Design and Multi-Firm Adoption under Uncertainty** 494
Jinghan Huang, Mingxi Chen, Yumeng Ye, Wenbin Wang, Timothy Tian, Libin Zhang and Zhaohui Lyu

Submission 61. **Multi-Agent Reinforcement Learning Framework with Resource-Affinity Modeling for Production-Logistics Collaborative Scheduling in MiC Factories** 496
Yujie Han, Zhiheng Zhao and George Q. Huang

Submission 62. **LLM-Driven Multi-Agent Group Chat Framework for Decentralized Garment Supply Chain Coordination** 506
Kexin Sun, Zhiheng Zhao and George Q. Huang

Submission 63. **Inventory Management of Traditional and Physical Internet-Enabled Logistics Network: A Simulation Study** 516
Sajan T John and K M Niyas

Submission 64. **A CPI-Enabled Resilient Flexible Job Shop Scheduling Framework for Shipyard Manufacturing in Industry 5.0** 518
Yanying Wang, Zhiheng Zhao and Guoquan Huang

Submission 66. **LLM-Augmented Strategic Framing for PI-Oriented Supply Chain Networks** 528
Ghada Ben Meriem, Shenle Pan, Walid Klibi and Mohamed Heythem Selmi

Submission 67. **Optimization in Multimodal Hubs under the Physical Internet Paradigm Through Environmental KPIs** 533
Jihane Lamri, Monica-Juliana Perez and Tarik Chargui

Submission 69. **Incorporation of a Social Fatigue KPI in Physical Internet Hub Operations: A Comparative Study Between FIFO and Social Policies** 537
Pedro Nunes Maia, Monica-Juliana Perez and Tarik Chargui

Submission 71. **Optimizing Multi-Modal Physical Internet Hubs: A Mathematical Model for Collision-Free Container Routing** 540
Tarik Chargui, Gabriel Zambrano Rey and Monica-Juliana Perez

Submission 72. **Hybrid Hyper-Heuristic Reinforcement Learning Agent for Disruption-Resilient Synchromodal Transport** 544
Thomas Kroiss, Felix Kamhuber, Sandra Stein and Fazel Ansari

Submission 73. **Cyber-Physical Internet-Enabled Logistics and Assembly Synchronization for Modular Integrated Construction in the Greater Bay Area** 546
Wanyi Song, Zhiheng Zhao and George Q. Huang

Submission 74. **An Urban Circular Delivery Service for the Implementation of Proximity-Kilometer Logistics** 547
Fouzia Ounnar and Patrick Pujo

Submission 75. **Multi-Stakeholder Governance Model for the Physical Internet** 557
Attallah Qweider, Shafagh Alaei and Koen Mommens

Submission 76. **Toward a Precursor System for the Physical Internet: Minimal Objectives and Requirements for Sustained Interconnection** 561
Malte Spanuth, Rod Franklin, Sahrish Jaleel Shaikh, Tiankuo Zhang and Paria Nourmuhammadi

Submission 77. **Federated Autonomy and Availability Intelligence in Multi-Organizational Physical Internet Ecosystems for Critical Supply Resilience** 571
Ashwin Pothen, Benoit Montreuil and Sahrish Jaleel Shaikh

Submission 78. **Toward a Physical Internet for Food Aid: Measuring Node Capacity and Need to Improve Cost, Reliability, and Fairness** 574
Abdullah Konak and Sadan Kulturel-Konak

Submission 79. **Orchestrating Sovereign Logistics Networks Through Performance Commitments in the Physical Internet** 584
Sahrish Jaleel Shaikh, Benoit Montreuil, Alan Brere, Tiankuo Zhang, Malte Spanuth and Paria Nourmuhammadi

Submission 80. **From Digital Supply Chains to Network Intelligence in the Physical Internet: An Empirical Consolidation of Capability Pathways** 597
Nadia Hmitti and Ahmed Maghni

Submission 82. **PI-Container Flow Optimization in a Multi-Echelon Physical Internet Network** 608
Hadjer Louzim, Abdelghani Bekrar, Fouad Maliki, Tarik Chargui, Atika Rivenq and Marwane Ayaida

Submission 83. **A Conceptual Physical Internet Framework for EV Battery Reverse Logistics** 618
Martyna Zielińska and Marta Cudziło

Submission 84. **Enabling Scalable and Resilient Supply Chains Through Hyperconnected Mobile Supply Chains** 628
Julien Maurice, Benoit Montreuil, Walid Klibi, Olivier Labarthe, Eunhye Song and Todd Ullom

Submission 85. **Autonomous Transport and the Physical Internet: Towards a Shared Research Agenda** 638
Elisha Van Kempen, Hannah Onwerwagt, Matthias Santing and Jaco van Meijeren

Submission 86. **Dynamic Pricing and Bayesian Demand Learning for Loading Dock Management** 645
Hamrah Kor, Lele Zhang, Mark Fackrell and Russell Thompson

Submission 87. **Station Simulations for Hyperloop Systems** 655
Lukas Eschment, Irina Jackiva, Jurijs Tolujevs, Thomas Schvänig and Walter Neu

Submission 88. **Designing Scalable Multi-Tier Networks for Physical Internet–Enabled Hyperconnected Urban Logistics** 665
Praveen Muthukrishnan, Walid Klibi and Benoit Montreuil

Submission 89. **Digital Shadow–Enabled Validation of Deployable π-Nodes for Distributed Modular Construction Networks** 676
Miguel Campos, Montreuil Benoit, Mcginnis Leon, Mingze Li, Zhihan Liu, Julien Maurice, Jorge Garcia, Mahsa Mahdavisharif, Jihye Jung, Kamyar Fatemifar, Negin Nadj, Reem Ashima and Todd Ullom

Submission 91. **Operationalizing Book & Claim Through a Digital Chain of Custody for Freight Decarbonization in the Physical Internet** 689
Harris Niavis, Ettore Gaulandi, Andrea Condotta and Ioanna Fergadiotou

Submission 93. **Strategic Infrastructure Design for a Hyperconnected Urban Delivery System Integrating Public Transport** 699
Cécile Dupouy, François Clautieux, Walid Klibi and Olivier Labarthe

Submission 95. **Design Principles for Human-Centered and Resilient Autonomous Logistics Systems** 709
Frank Phillipson and Ruben Fransen

Submission 96. **Converged Digital–Physical Identity Resolution Architecture for Enabling Traceability in the Physical Internet** 719
Wenqing Lei, Ming Li, George Huang and Yong-Hong Kuo

Submission 98. **Assessing the Role of Firm Size Distribution and Hub Localization in Collaborative Supply Chain Design under the PI Paradigm** 725
Luz Helena Mancera, William Guerrero and Jairo R. Montoya-Torres

Submission 99. **Resource Pooling for Multiple Stakeholders in Hyperconnected Warehouse Networks under Uncertainty** 734
Yaxin Pang, Shenle Pan, Olivier Labarthe and Walid Klibi

Submission 103. **Cyber-Physical Internet-Enabled Spatial-Temporal Traceability for Production Logistics Synchronization: An Agentic Reasoning Approach** 743
Jinpeng Li, Zhiheng Zhao and George Q. Huang

Submission 104. **Blockchain-Enabled IoT for Smart Food Traceability: Cost and Performance Insights** 753
Hari Sannamuri

Submission 106. **A Cyber-Physical-Internet-Based Self-Organizing Path Planning Method for Production Logistics** 759
Haoran Liu, Bin Huang, Zhiheng Zhao and George Q. Huang

Submission 107. **Digital Twins: Accelerating the Transition of Port Community Systems Toward the Physical Internet – The Atlantic Corridor Case Study** 768
João Nabais, Raquel Pereira and Carlos Batista

Submission 108. **Physical Internet Across Sovereign Borders: A Formalized Design Framework and the Price of π-Sovereignty** 780
Oscar Hasburn-Babich, Benoit Montreuil and Craig Tovey

Submission 109. **IoT Enabled Smart Dry Containers as Foundational Infrastructure for the Physical Internet: A Standards-Aligned Architecture for Fire Risk and Door Event Detection in Multimodal Transport** 792
Marianna Levtov

Submission 111. **Cyber Physical Internet-Enabled Spatial-Temporal Traceability for Green Methanol Production and Logistics** 804
Qi Liu, Leidi Shen, Zhichen Wu, Zhiheng Zhao and George Q. Huang

Submission 113. **Mapping the European Logistics Data Ecosystem for the Physical Internet: A Structural Analysis of Gaps and Interoperability** 814
Yu Liu and Eric Ballot

Submission 114. **Toward an Aviation-Enabled Physical Internet: Designing High-Speed Layers for African Multimodal Logistics Networks** 824
Abiy Asrat Jiru

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE
DALHOUSIE UNIVERSITY

Engineering a Physical Internet Concept for Food and Beverage Distribution in Sparse Geographies using Network Optimization and Game-Theory

Uday Venkatadri
Dalhousie University, Halifax, Canada
Corresponding author: Uday.Venkatadri@dal.ca

Abstract: *Sparse geographies present challenges for transportation and distribution. Examples of such geographies include Atlantic Canada and Australia. The grow local movement has been a response to globalization over the last couple of decades. However, with increased tariff threats from the United States, this movement is expected to gather more momentum. In the food and beverage industry, this implies growth opportunities for local producers. The cost of distribution for local producers has been estimated at 20-35% of the sales price of products. In fact, distribution costs represent a Catch-22 dilemma in sparse geographies. When price is high due to high distribution costs, demand is low. The low demand results in high production and distribution costs due to the inability of producers to exploit economies of scale. However, an area of opportunity for local producers in the current tariff environment is to form a shared private logistics network to reduce landed costs to end customers and exploit their market potential. In this paper, we look at how collaborative warehousing and transportation structures using Physical Internet concepts can be designed in Atlantic Canada to increase customer responsiveness and reduce transportation and warehousing costs by exploiting the structure of the road networks. A methodology based on network optimization and game theory is proposed.*

Keywords: *Physical Internet, Sparse geographies, Economies of Scale, Shared logistics resources, Distribution and inventory costs, Benefits of cooperation and coordination, Network optimization, Game-theory.*

Physical Internet (PI) Roadmap Fitness*: ☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories), ☐ Transportation Equipment, ☒ PI Networks, ☒ System of Logistics Networks*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

1 Introduction

Nova Scotia (the most populus province in Atlantic Canada) has a loop and split cross-hair structure that connects most rural producing communities with the larger urban centre located in Halifax which has most of the customer demand. Figure 1 (taken from Google maps) shows a map of Nova Scotia as situated on the East Coast of Canada. There are four provinces in Atlantic Canada (New Brunswick, Prince Edward Island, PEI, Newfoundland and Labrador, situated to the Northeast in the map and Nova Scotia). The island of St. Pierre and Miquelon is French territory and comes under the European Union. In this paper, the focus of interest is Nova Scotia, although there are implications for distribution in the food and beverage industry through interactions with producers and markets in the other jurisdictions that have low densities.

Within Nova Scotia, the major highway map (Figure 2) reveals a very interesting structure from the physical network (PI) perspective. The upper branch of the route in green is the Nova Scotia segment of the Trans-Canada highway the world's longest national road, spanning approximately 7,476 km across all ten Canadian provinces from Victoria, B.C., to St. John's,

Newfoundland. The connection to Newfoundland is via boat (equipped to carry tractor-trailers) from Sydney. The connections to New Brunswick and PEI is via Amherst in the West of the Province. Halifax, New Glasgow and Sydney are the largest population centres (and therefore markets) in the Province, with Halifax's population representing approximately 50% of the total population. Truro is at the intersection of the green route(s) and the blue artery to Halifax and the blue loop from Halifax to Kentville to Yarmouth back to Halifax. The routes in blue are major highways. The routes in black are arterial highways that could be used for collection (or distribution) in smaller trucks.



Figure 1: Location of Nova Scotia in the Mid-Atlantic (Source: Google Maps)

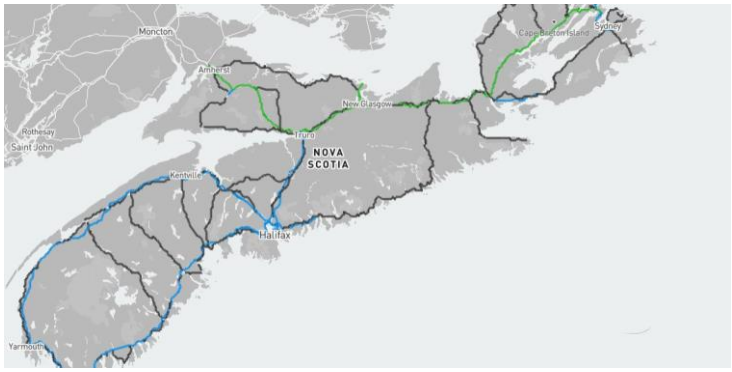


Figure 2: Publicly Available Highway Map Provided by the Government of Nova Scotia (https://data.novascotia.ca/Roads-Driving-and-Transport/Nova-Scotia-Road-Network-NSRN-Map/pg9a-kkaq)

The region around Kentville is known for fruit (particularly apple) and vegetable production. There are fisheries along the South Shore (Yarmouth-Halifax), dairy in the region of Truo, blueberry farms between Truro and Amherst and various small scale food and beverage producers on Cape Breton island, from where the two green routes split into a major route to Sydney (Hwy 105) and a shorter segment (Hwy 104) to St. Peter's. There are wineries along the segment from Halifax to Kentville (about 2/3^{rds} of the way) and distilleries in Cape Breton as well as in other parts of the province.

By the same token, the customers for these products are often spread across the sparse geography of 1.1 million residents, half of whom live within a denser core sorrounding Halifax. The customer locations include daily or weekly farmer's markets, high-end restaurants, small speciality stores and much larger grocery stores which often dedicate sections to small food and beverage producers for various reasons such as higher margins, corporate social responsibility

to promote local products, and the willingness of customers to pay a premium for local products in the face of globalization and tariff regimes (Praveen et al., 2025; Charlebois, 2026).

There are a variety of small producers who would benefit from a shared logistics network, since each producer may not have the necessary volume/economy of scale to reach markets effectively. They can rely on private distributors, of course, but there are transportation premiums that they would have to incur and it would be difficult for these companies to get block rates or have a say in how an efficient collection/distribution network could be engineered. However, for local producers, choosing a distributor often adds 20-35% to the sales price (Farm Credit Canada, 2026). Fortunately, there is an industry association across the four provinces called Food and Beverage Atlantic (https://atlanticfood.ca/), which has member companies producing fresh (including greenhouse grown) or processed fruits/vegetables, cheese, alcholic and non-alchoholic beverages, etc. Their website lists members in the food and beverage category. Such an industry association can potentially promote a coalition of logistics partners working together under PI principles such as hyperconnected logistics, open and shared infrastructure, standardization (PI-containers, pi-boxes, or PI-wagons, packet-based logistics, decentralized routing and interconnectivity (Pan et al., 2017).

From a (PI) perspective, the research questions addressed in this paper are: 1) How can the food and distribution network be engineered based on PI principles? 2) What are the cost and environmental benefits (the latter being out of scope in this paper) of collaboration? 3) How do marginal costs of distribution reduce as more players enter the network and how can a coalition of existing partners in a shared network be set up and hopefully expanded?

To this end, we present a combined network flow optimization and game-theoretic model which is formulated to provide insights within a sequential decision-making framework. The network flow optimization model is used to configure a PI-network for existing coalition partners. The model could eventually include mode and modular PI-Container choices as a function of transportation, inventory and warehousing costs. The game-theoretic model is used to estimate that a new partner will join the network. Finally, a sequential decision-making framework is proposed to predict how the network will evolve over time as information is passed to potential partners to facilitate network growth. The paper will focus on the modelling elements rather than present mathematical or computational results which will necessarily need to be done for a more complete proof-of-concept. More modelling elements are currently being formulated.

2 Literature Review

We briefly review the literature pertinent to the logistics ideas in this proposal. PI is a very broad framework for shared logistics networks design and includes several layers (from policy to operational) and concepts (such as open connectivity). These ideas can be traced from the original papers on PI (Montreul, 2011; Montreuil et al, 2012; Meller et al, 2012) to later work such as Pan et al. (2017). Kay (2024) present the concept of a public logistics network (PLN) and write that:

“In a PLN, the different functions of the network would be separated so that a single firm is not required for coordination. This would enable scale economies to be realized in performing each logistics function since each element of the network has access to potentially all of the network’s demand. The increase in scale would make it economical to have many more transshipment points. Each transshipment point, or DC, could be an independently operated

facility that serves as both a freight terminal and a public warehouse, and could be established in small cities and towns that would never have such facilities if they were served as part of a proprietary, private logistics network"

Such networks do not exist in Atlantic Canada. Therefore, in the short term, food and beverage distribution would require a shared (perhaps industry group) initiative to jump start a coalition. This of course requires collaboration, trust, fairness, etc. which are alluded to in the road map of Ballot et al (2020). True to the spirit of the roadmap timeline, this work should also be seen as a bottom-up approach, i.e., "Get organized, start small, and grow."

From the methodological point of view, Niu et al (2022) present a game-theoretic analysis of a logistics service supply chain differentiated by a short-distance and a long-distance logistics service provider. They conclude that the long-distance case always favours PI while preference for PI in the short-distance case depends on the potential for efficiency improvement and service competition intensity. In our own reflection, game theoretic considerations are relevant to this paper because a shared logistics distribution network first requires a coalition of partners. As the network grows and gains in economies of scale, other partners would find the coalition interesting to join. Naturally, a question to answer is what their benefit should be given that they did not assume the initial risk. Dong and Franklin (2021) present a conceptual framework for PI with a stylized network flow model. Network optimization is a foundational methodology to engineer the logistics network presented in this paper. To this end, Naganawa et al (2024) show how PI enables multiple companies to efficiently share warehouses and trucks and reduce CO_2 emissions. They use a genetic algorithm and the Lin–Kernighan heuristic with real-life simulations to evaluate the performance of an optimization model and show how the resulting solution reduces CO_2 emissions by 54%.

A few other papers are worth mentioning. Using the illustrative case of the Mexican automotive supply chain, Chadha et al. (2022) show how peddling, a consolidation strategy in traditional logistics, can complement the PI supply chain. By the same token, same-day driver return was always a consideration in PI, right from the beginning (Montreuil, 2011: "Modern Cowboy" Trucking). In their Eastern Canada road network case study (with some of the network overlapping with the Nova Scotia road network), Fazili et al (2017) explicitly model driver assignment constraints to limit driver assignment durations before returning home. Venkatadri et al (2016) show how, among other things, that two-way routing is more efficient than P2P routing from the point of view of the cost-time tradeoff in logistics dispatch. This ties back to the theme of sustainability where empty returns can Finally, collection of recyclable materials (such as aluminium and plastics) through the road network in Nova Scotia has been studied in Chari et al (2016). Since this study was designed for collection of recyclable materials in sparse geographies, it models collection frequencies. Sanders and Montreuil (2021) investigate the use of mobile hubs for consolidation and distribution of goods in the restaurant supply chain where food travels from rural farms to suburban and urban restaurants. The authors state in their abstract that:

"This paper computes suitability and location analysis through a novel multi-criterion scoring methodology utilizing kernel density and network analysis. The effectiveness of these mobile hubs is assessed through strategic routing, where the routes are optimized for time and distance."

This paper is closest in intent to our scan of the literature. However, to the best of our knowledge, there have been few case studies in the literature dealing with food and beverage distribution, particularly in sparse geographies. The major contribution of this paper is that food and beverage distribution in sparse geographies such as Nova Scotia offer unique challenges (lack of density and thereby economies of scale) and opportunities (a very specialized road network that actually favours PI adoption, as we will see in the next sections) worth exploring. Another gap in the literature is that there has been no attempt to understand how a shared logistics network can grow once setup. This is a methodological contribution of the paper.

3 Engineering the Physical Internet for Low Volume Food and Distribution in Nova Scotia

Since Nova Scotia does not have a shared logistics network, a private network according to PI principles can be designed for the case study in question. In the methodological framework presented below, it is assumed that a coalition of food and beverage producers is interested in setting up the network. This assumes of course that the network would be beneficial to the partners in the coalition (a hypothesis that needs to be validated with real data). We then conceive a game-theoretic two-party sequential game where a new member wishes to enter the coalition over time. Their decision to enter would be based on the benefits incurred because of access to a larger market based on a reduction in the landed cost of the product at the markets (as expected to occur).

3.1 Optimizing PI hub to PI flows and PI node capacities in the network

A standard network optimization model is first presented to optimize hub to hub flows and PI-node capacities. The decision maker in this model is the logistics planner of the coalition. In this model, the coalition makes hub sizing decisions and flow decisions (movement of PI-transporters between hubs). For the most part, a weekly decision model works well for this case study. Although seasonality is important in the industry, we omit seasonal indices for simplicity of exposition with the understanding that the model below needs to be indexed for seasons.

The sets in the model (summarized in Table 1) are the supplier/product combinations *i* (it is assumed for now that the suppliers are responsible for delivery to the nearest PI-hub), PI-nodes *j* (intersection points of the road network in Figure 2) and markets *m*. For example, Halifax would not have a direct connection to New Glasgow or Sydney. The markets are large and small stores, farmer's markets in the regions and possibly eCommerce drop-offs. Every supplier-product combination is delivered to the closest PI-hub of origin.

Table 1: Network Optimization Model Sets

Set	Index	Description
P	*i*	Supplier-product combinations
H	*j(or k)*	Set of PI-hubs (also representing markets)
H*j*	*k*	Set of PI-hubs connected directly to hub *j*

S	*s*	Set of product storage categories (refrigerated, room temperature controlled, etc.)

The parameters of the model (summarized in Table 2) are the fixed amortization costs of operating the PI-hubs within the weekly planning timescale as a function G_s (..) of flow of products in product storage category *s*, the fixed and variable costs of PI-container operations (*g* and v_{ij} respectively), and the last mile delivery cost of products to markets. The fixed cost function G_s(..) is modelled as an economy of scale cost as is customary in the supply chain literature. However, the fixed cost of the PI-containers is assumed to depend on the number of containers. In this case study, we assume that one container with multiple weekly trips between adjacent PI-hubs is sufficient, as our preliminary understanding of the case study data indicates. It is worth mentioning that the PI architecture here is very robust. Since dedicated containers move back and forth between hubs and both producers and customers are located throughout Nova Scotia, shared logistical access to markets is remarkably high (as seen in Figure 2) while the risk of deadhead travel or unused PI-container capacity is potentially low.

In this concept, the coalition would normally route smaller trucks from the PI-hubs to reach these end points using peddling (Chadha et al, 2022). Again, for simplicity, we ignore these costs and how they are shared and simply use an exogenous variable cd_{ij} to denote the last mile delivery cost in market *j* for supplier-product *i* estimated from an external capacitated vehicle routing model (omitted for simplicity's sake and due to page limitations). For convenience, the markets are assumed to be located at the PI-hubs. The demand of supplier-product combination *i* is assumed to be d_{ij}. The PI containers are assumed to have a capacity of C_s PI-cube units for each storage class *s*. O_i is the origin hub of supplier-product combination *i*.

Table 2: Network Optimization Model Parameters

Parameter	Description
cd_{ij}	Last mile delivery cost in market *j* for supplier-product *i* estimated from an external capacitated vehicle routing model for one PI-cube unit
d_{ij}	Demand for supplier product *i* at PI-hub *j* in PI-cube units
O_i	Origin hub of supplier-product combination *i*
G_s(..)	Functional parameter representing amortized weekly fixed costs of operating PI-hubs as a function of flow through PI-cube units of storage class *s*. This is an economy of scale function.
g	Amortized weekly fixed cost of a PI-Container (we assume one PI-container dedicated to PI-hub to PI-hub pair for a total of *n* PI-Containers)
v_{ij}	Variable round-trip cost of a PI-Container return between PI hubs *i* and *j*
C_s	Total capacity of storage class s in a PI-Container
a_{is}	A zero/one parameter indicating whether supplier-product i belongs to storage class

The decision variables are the frequency of (f_{jk}) weekly roundtrip PI-container shipments from eligible PI-node to PI-node and the flow of supplier-combination i between PI-hubs j and k (x_{ijk}) in weekly PI-cube units flows. For PI-transport, we assume standard multi-zone containers (appropriately sized for different product classes) and PI-cubes (for simplicity's sake, we assume only one size in this paper) into which the products can be packed.

Table 3: Network Optimization Decision Variables

Variable	Description
f_{jk}	Frequency of weekly PI-Container shipments between PI-hubs j and k with a direct connection
x_{ijk}	Flow of supplier-product i from PI-hubs j and k with a direct connection in PI-cube units

The network optimization model to configure the PI-network for the coalition may now be written as:

Minimize (PI-NFOM):
$\sum_j \sum_k v_{jk} f_{jk} + \sum_j \sum_s G_s(\sum_i \sum_k a_i s (x_{ijk} + x_{ikj})) + \sum_i \sum_j cd_{ij} d_{ij} + gn$

Subject to:
$\sum_k x_{ijk} = \sum_k d_{ik} \; \forall \; i \; (j = O_i)$ (1)

$\sum_k x_{ikj} - \sum_k x_{ijk} = d_{ij} \; \forall \; i \; (j \neq O_i)$ (2)

$\sum_k s_i x_{ijk} \leq C_s f_{ij} \; \forall$ i $\forall \; j$ *(3)*

All f's and x's $>=0$

The objective in the PI Network Flow Optimization model (PI-NFOM) is to minimize the total network fixed and variable costs over a weekly (amortized) period. The first term minimizes PI-container transfer costs between pairs of adjacent PI-hubs. The second term is shown as a cost function but can be linearized using approximations such as piece-wise linear programming. The third term represents the last mile delivery cost which would ideally be linked to another capacitated vehicle routing problem formulation for last mile delivery to the customers closest to a given PI-hub. It is a constant and may be omitted from the objective. The fourth term is simply the total fixed PI-container cost which may also be omitted in this version of PI-NFOM. Equation (1) specifies the supply quantity to the origin node for any supplier-product i. Equation (2) enforces flow balance for all non-origin nodes. Constraint (3) enforces that the flows out of a node for any supplier-product combination respects the capacity of the weekly storage class capacity of PI-Container flows.

Once the problem is solved, protocols can be developed for cost-sharing between coalition partners. One such allocation could be fixed cost sharing based on PI-cube volume and variable cost sharing based on PI-cube volume-distance.

3.2 Game-Theoretic Modelling of New Partner Entry

Under the hypothesis that a PI-Coalition will be economically feasible and would offer economies of scale in logistics handling (which we believe is the case), it is in the interest of the coalition to recruit new members. By the same token, a new member would have an incentive to join the coalition if its current logistics costs can be reduced.

We assume that the volume weighted logistics cost is one driver of this decision (within scope). The other driver would be market expansion (out of scope in this paper). Figure 3 illustrates the concept. A potential new entrant could calculate its total volume weighted logistics cost, regardless of whether it has its own distribution network or uses contracted logistics as well as define a low threshold value (*LowT*). The idea is that the probability of joining the coalition would be one if the low threshold value is proposed to it by the coalition. On the other hand, if the weighted cost of distribution is greater than equal to its current operation point (*COP*), the probability of joining the coalition is zero.

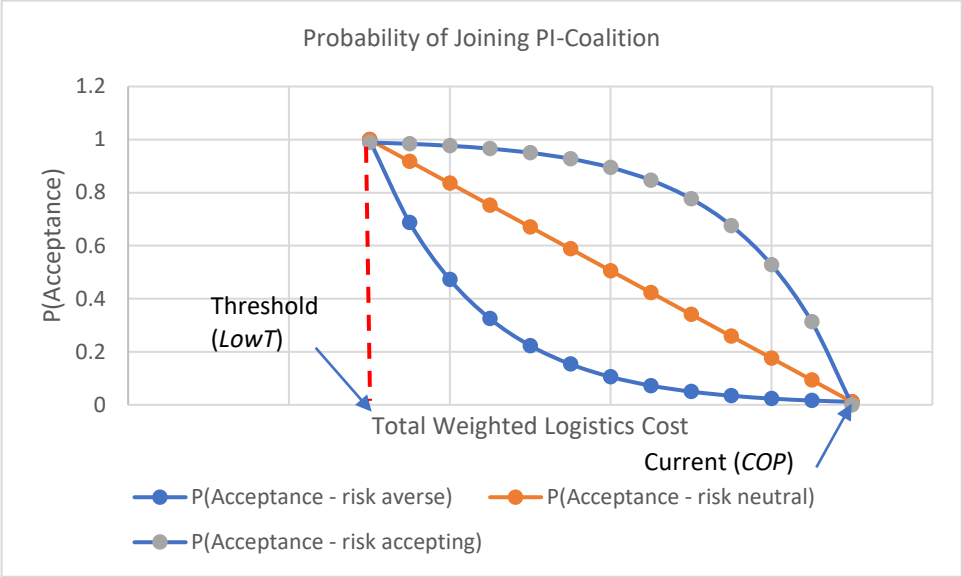


Figure 3: Joining the PI-Coalition Based on Total Weighted Logistics Cost

The exponential curves with parameter λ could model a risk averse or risk accepting potential new entrant.

$P(A|\text{Risk averse}) = LowP\ e^{-\lambda x}$

$P(A|\text{Risk accepting}) = COP^{\,e^{-\lambda(1-x)}}$

Where x is the total weighted cost of distribution after joining the PI-Coalition and $\lambda = -ln(\varepsilon/(COP+LowT)/COP)$

3.3 Proposed Solution Methodology

To understand how the PI-Network modelled would grow, we propose a methodology in which a new partner's entry would be simulated and their requirements entered into the PI-NFOM (as seen in Figure 4). This could be from a database of industry association partners. Once the model is solved, a price could be quoted to the potential partner (space limitations prohibit development of quoting mechanisms). The aim would be to quote a price based on existing network costs (if no infrastructure expansion is required to the PI network) or a fair

allocation of expansion costs if required. Depending on the quoted price, the potential entrant would evaluate their response based on Section 3.2. If accepted, the network would be expanded for future entrants. Through this sequential process, the growth of the network could be predicted over time and the efficacy of various cost-sharing mechanisms evaluated.

4 Conclusion and Future Work

The modelling framework presented in this paper can be used to inform producers as well as logistics parties (transportation and shared logistic hub operators) on the benefits of a shared collaboration model for distribution. The aim of the paper was to demonstrate that 1) The road network and low customer density of Nova Scotia offer opportunities for distribution organized according to PI principles and 2) Provide decision makers in the industry a methodological framework using network optimization and game-theory concepts to evaluate how a PI-based network could be setup and how it might evolve.

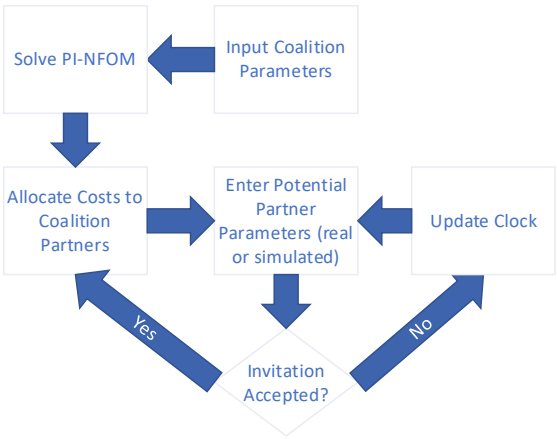


Figure 4: Sequential Optimization Procedure to Model PI-Coalition Growth

A numerical implementation of this methodology would provide more insights and lead to industry specific or more general knowledge contributions to PI and logistics research. Future work could include exploring industry context specific multiple PI-container configurations, multiple pickup and delivery PI-movers, modular design of PI-hubs based on fractal layout principles and modelling of GHG emissions. On the game-theoretical side, the coalition partners at the beginning assume risks. So, the question then becomes how to allocate costs to partners who come into the coalition later, depending on their time of entry. If the cost is higher, how would it decay? Similarly, if there is a capacity expansion needed due to a partner's entry, how would it impact their allocated cost of entry (since the existing coalition would also benefit)?

Acknowledgements

The author of this paper would like to thank the Natural Sciences and Engineering Research Council of Canada for continuing to support his basic research through their Discovery Grants program. He would also like to thank Mr. Dwayne Boudreau, then a brand and growth specialist at Food & Beverage Atlantic, Dr. Gumataw Abebe and Mr. Brennan Gillis at the Faculty of Agriculture at Dalhousie University for numerous thoughtful conversations.

References

- Ballot, E., Barbarino, S., van Bree, B., Liesa, F., Franklin, J.R, 't Hooft, D., Nettsträter, A., Paganelli, P., Tavasszy, L.A. (2020): ALICE-ETP, Roadmap to the Physical Internet,

https://www.etp-logistics.eu/wp-content/uploads/2022/11/Roadmap-to-Physical-Intenet-Executive-Version_Final-web.pdf, retrieved on 2026-03-31 at 5:30pm Atlantic Standard time

- Chadha, S.S., Ülkü, M.A., Venkatadri U. (2022): Freight delivery in a Physical Internet Supply Chain: an applied optimisation model with peddling and shipment consolidation, International Journal of Production Research, Volume 60, Issue 16, Pages 4995-5011
- Chari, N., Venkatadri, U., & Diallo, C. (2016): Design of a reverse logistics network for recyclable collection in Nova Scotia using compaction trailers, *INFOR: Information Systems and Operational Research*, *54*(1), 1–18
- Charlebois, Sylvain (2026): Trade, Investment and Nutrition: Lessons from the Canadian Case for Food Systems Governance. Public Health Nutrition 29, no. 1 (2026): e47, https://doi.org/10.1017/S1368980026102134.
- Dong, C. and Franklin, J.R. (2021): From the Digital Internet to the Physical Internet: A Conceptual Framework with a Stylized Network Model, J Bus Logist, 42: 108-119, https://doi.org/10.1111/jbl.12253
- Farm Credit Canada (2026): Consider the 'other' costs when calculating selling price, https://www.fcc-fac.ca/en/knowledge/consider-other-costs-when-calculating-selling-price, retrieved on 2026-03-31 at 5:30pm Atlantic Standard time
- Mehran Fazili, M., Uday Venkatadri, U., Pemberton Cyrus, P., Tajbakhsh, M. (2017): Physical Internet, conventional and hybrid logistic systems: a routing optimisation-based comparison using the Eastern Canada road network case study, Int. J. Prod. Res. 55(9), pp. 2703-2730
- Kay, M.G. (2024): Protocol Design for a Public Logistics Network (2004), In Progress in Material Handling Research: 2004, R. Meller et al., Eds., MHI, Charlotte, NC, 2004, pp. 181–188 (presented at the Int. Material Handling Res. Colloq., June 13–17, 2004, Graz, Austria).
- Meller R.D., K.P. Ellis, B. Loftis (2012): From Horizontal Collaboration to the Physical Internet: Quantifying the Effects on Sustainability and Profits When Shifting to Interconnected Logistics Systems, Final Research Report of the CELDi Physical Internet Project, Phase I, http://faculty.ineg.uark.edu/rmeller/web/CELDi-PI/Final%20Report%20for%20Phase%20I.pdf, 2014/02/20.
- Montreuil, B., Meller, R.D., Ballot (2012): Physical Internet Foundations, IFAC Proceedings Volumes, Vol. 45, Issue 6, Pages 26-30, https://doi.org/10.3182/20120523-3-RO-2023.00444.
- Montreuil B. (2011): Towards a Physical Internet: Meeting the Global Logistics Sustainability Grand Challenge, Logistics Research, v3, no2-3, 71-87, https://doi.org/10.1007/s12159-011-0045-x
- Naganawa, H.; Hirata, E., Firdausiyah, N., Thompson, R.G. (2024): Logistics Hub and Route Optimization in the Physical Internet Paradigm, Logistics, Vol. 8, No. 37. https://doi.org/10.3390/logistics8020037
- Niu, B., Dai, Z., Liu, Y., Jin, Y. (2022): The role of Physical Internet in building trackable and sustainable logistics service supply chains: A game analysis, International Journal of Production Economics, Volume 247
- Pan, S., Ballot, E., Huang, G. Q., & Montreuil, B. (2017): Physical Internet and interconnected logistics services: research and applications. International Journal of Production Research, 55(9), 2603–2609, https://doi.org/10.1080/00207543.2017.1302620
- Praveen S., Ramaprasad, B.S., Prabhu, K.P.N,, Rao, K.P.S, Phadke, R.A., Kizhekepat, S. (2025): Rooted in loyalty: How consumer ethnocentrism shapes consumer satisfaction with homegrown brands. Acta Psychologica, Volume 260

- Sanders, I.T., Montreuil, B. (2021): Reducing Transport Miles Through the Use of Mobile Hubs: A Case Study in Local Food Supply Chains, Proceedings of the 54th Hawaii International Conference on System Sciences
- Venkatadri, U., Krishna, K.S., Ülkü, M.A. (2016): On Physical Internet logistics: modeling the impact of consolidation on transportation and inventory costs, IEEE Transactions on Automation Science and Engineering, Volume 13, Issue 4, Pages 1517-1527

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE
THE UNIVERSITY OF MELBOURNE

Dynamic Slot Reallocation for Physical Internet Hubs

Anshul Vijay [1], Farzaneh Zirak[1], Russell G. Thompson[1], Neema Nassir[1]
1. The University of Melbourne, Melbourne, Australia

Corresponding author: anshul.vijay@unimelb.edu.au

Abstract: *This study assesses how GS1 Scan4Transport data sharing technology for fleet location sharing from independent carriers can assist in solving the Dynamic Slot Reallocation Problem for PI-Hub operators. To develop the network and real-time vehicle location data, survey of the Sydney Courier Hub was conducted. Utilised as an open transhipment facility at the border of Sydney CBD, Australia, numerous courier organisations incorporate it within their delivery operations. Qualitative responses from users of the facility informed the test network and quantitative analysis of video footage of hub activity informed the real-time location data of incoming vehicles.*

Keywords: *City Logistics, Physical Internet, Hyperconnected City Logistics, Artificial Intelligence, Optimisation, Machine Learning, Deep Reinforcement Learning, Multi-Agent System, Surveying*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☒ *PI Networks,* ☐ *System of Logistics Networks,* ☒ *Vertical Supply Consolidation,* ☒ *Horizontal Supply Chain Alignment,* ☒ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

1 Introduction

City Logistics's primary concern is to minimise negative externalities arising due to urban freight activities (Taniguchi & Thompson, 2001). These externalities include: social, environmental, economic, safety and energy costs. Based on the systems approach, the CL framework encourages competition via free market dynamics for driving innovation and network efficiencies.

Physical Internet seeks to enhance sustainability of supply chain networks through the incorporation of: open & shared networks; standardised & modular load carriers; track & trace protocols, & certificates. PI envisions seamless interconnectivity between network participants by integrating goods, fleets and transhipment facilities (Ballot et al., 2021).

The integration of PI that endorses open, shared & standardized freight networks with CL, that seeks to minimize urban freight costs, culminates in to Hyperconnected City Logistics (HCL) (Crainic & Montreuil, 2016). Functioning as nodes where goods are reconsolidated & transshipped to other vehicles within the HCL network, efficient PI-Hub operations are critical to the overall success of PI (Montreuil et al., 2010). Nevertheless, there have been limited studies that have quantitatively assessed its potential.

Average speed for vehicles along a road, especially within the urban area, varies through the day due to congestion (McKinnon, 2015). Uncertainty in link travel times due to variable congestion levels

may make the initial master docking schedules for vehicles inaccurate. Hence, these initial schedules need to be further adjusted based with the new information of vehicle locations. There has been sparse literature on investigating uncertainty in arrival times of vehicles at PI-Hubs to assist with the slot allocation problem (Hillerstrom et al., 2018).

This study proposes a Dynamic Slot Rescheduling Framework for PI-Transporters at PI-Hubs. A ML Classifier predicts ETA of incoming vehicles, & a proximal policy optimisation (PPO) powered scheduling agent decides when to reschedule. A Binary Integer Linear Program is used to create the schedule. To the researchers' best knowledge, no study to date has integrated RL for the dynamic slot reallocation problem (DSRP) within cross-dock and PI-Hub literature. To develop the network and real-time vehicle location data, survey of the Sydney Courier Hub was conducted. Utilised as an open transhipment facility at the border of Sydney CBD, Australia, numerous courier organisations incorporate it within their delivery operations. Qualitative responses from users of the facility informed the test network and quantitative analysis of video footage of hub activity informed the real-time location data of incoming vehicles.

This study assesses how GS1 Scan4Transport data sharing technology for fleet location sharing from independent carriers can assist in solving the DSRP for PI-Hub operators. Furthermore, this study investigates how infrastructure-based travel time estimates can enhance slot reallocations for PI-Hubs by enabling increased accuracy in vehicle arrival times.

2. Literature Review

This section provides an overview of literature related to rescheduling problems within the cross-docks and data sharing.

2.1 Rescheduling Problems

Despite the inherent uncertainties, majority of the literature within cross-dock scheduling operations have not considered the stochastic environment, where rescheduling is required (Nasiri et al., 2022; Prakoso et al., 2022). Those that do consider the uncertainty arising from vehicle arrival times, have traditionally focused on developing a robust master schedule that seeks to incorporate the worst-case scenarios (Acar, Yalcin, & Yankov, 2012; Chen et al., 2013; Konur & Golias, 2013; Ladier and Alpan, 2016). Although not prevalent in cross-dock literature, rescheduling problems have been considered in other domains including: the manufacturing environment (Li & Li, 2020; Rahmani & Ramezanian, 2016); healthcare setting for nurse scheduling (Wolbeck et al., 2020); and public transportation such as trains (Binder et al., 2017) and buses (Gkiotsalitis, 2021).

2.2 Data Sharing

To enable collaboration amongst the industry players, data standards are an important component. Le Roch et al. (2015) applied the GS1 based, EPC global standards to a PI-enabled FMCG supply chain, enabling the management of reusable pallets between the manufacturers and retailers' distribution centers. Meanwhile, K. Chen et al. (2017) introduced a cloud based, PI- enabled Building Information Modelling System, that sought to improve resource allocation efficiency and human errors within the pre-fabricated construction industry.

Hakimi et al. (2009) conceptualised an information architecture for a collaborative supply chain network, along with business intelligence tools for its member constituents. This was extended by Hakimi et al. (2010) who introduced a multi-agent simulation platform to enable more informed decision making by managers operating within a dynamic, collaborative environment.

GS1 has also developed the Scan4Transport (GS1S4T) standards that enables data sharing among network participants via a 2D barcode. Data such as origin, destination, delivery handling conditions are encoded within the barcode, enabling prompt access to relevant information for all parties within the supply chain (*Encoding Transport Process Information - GS1 Implementation Guideline*, 2021). In 2021, three types of pilots were conducted using the GS1S4T (Scan4Transport Pilot Report, 2021):

1. A desktop pilot which consisted of assessing impact on the participating business if GS1S4T was implemented within their operations.
2. A lab pilot, where a small scale implementation was tested within the participant's business.
3. A physical pilot, where the GS1S4T was physically implemented across a company's operations.

However, impact of utilising GS1S4T within a collaborative environment as demanded by PI has had limited consideration. Therefore, opportunity exists to investigate how PI-network participants may be impacted through the use of GS1S4T and enhance network performance and sustainability.

3. GS1 Scan4Transport

Developed by GS1, Scan4Transport (GS1S4T) data standards enable data sharing among SCN network participants via a 2D barcode, including carriers, shippers and customers (Figure 1). It also enables location updates of goods with the supply chain and within transshipment facilities. Data includes key elements among others: origin, destination, delivery handling conditions are encoded within the barcode, enabling prompt access to relevant information for all parties within the supply chain (*Encoding Transport Process Information - GS1 Implementation Guideline*, 2021). Customers can update destination and time window information that is accessible by all relevant SCN participants, along with confirmation of having received goods.

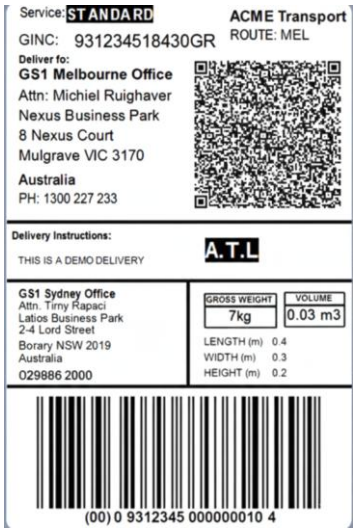


Figure 1 – GS1S4T Label

4. Problem Statement

PI-Hubs are responsible for the efficient transference of goods contained within PI-Containers from one carrier to another. The transition process involves the scheduling of PI-Transporters at PI-Docks, unloading of PI-Containers from PI-Transporters, reconsolidation, & loading of PI-Containers onto their respective outbound PI-Transporters.

Uncertainty in link travel times due to variable congestion levels may make the initial master docking schedules for vehicles inaccurate. Hence, these schedules can be further aligned with the knowledge of the vehicles' real time positions.

As considered by Prakoso et al. (2022), this study develops a predictive-proactive methodology to tackling DSRP. While Prakoso et al. (2022) utilised a rolling-horizon stochastic optimisation informed by predicted presence probabilities using ML classifiers, this study combines ML based ETA predictions and DRL policy for proactive decision making. The DRL agent proactively decides when to reschedule PI-Transporters to PI-Docks given the latest state of the system. It seeks to minimise the total time in PI-Hub experienced by the PI-Transporters, while considering future consequences of its current choices.

5. Data Collection and Synthesis Methodology

Data was gathered from the Courier Hub, located in Goulburn St. car park, Sydney. Utilised by several small and medium-sized carriers as an open consolidation hub, where goods are consolidated, transferred and/or stored temporarily in cages or parcel lockers to serve the CBD.

In-situ surveys were conducted with facility users to inform the test network (vehicle origins), while video analysis of vehicle and parcel movements within the hub informed the following characteristics: Vehicle Type (van, car, walker, cargo-bike); actual arrival time; Day (Tuesday, Wednesday, Thursday); Date; Month (September – March); Year (2022, 2023).

The master schedule of the vehicles, nor their real-time locations were known, thus they were determined using the observed vehicular activity. Table 1 outlines the procedure used.

Table 1 – Overview of Data

Procedural Step	Description
Step 1	Vehicle Assignments
Step 2	Vehicle Origins
Step 3	Schedule Development
Step 4	Real Time Locations

Step 1: Vehicle Assignments

In Step 1, inbound & outbound vehicles were determined. Vans & cars were allocated as inbound vehicles as this reflects how they are utilised in practice, requiring to travel greater distances with greater loads. Conversely, walkers and cargo-bikes were defined as outbound vehicles serving the nearby CBD.

Step 2: Vehicle Origins

In Step 2, random allocation of origins were assigned to each inbound and outbound vehicle. The sets used for each were determined based on the survey of the users of the courier hub. Table 2 & Table 3 indicate typical origin and destinations of vehicles utilising the Courier Hub. These locations informed the network that was formulated for the PI-Hub. The origin sets used for inbound and outbound vehicles respectively were {10, 20, 30} & {1, 2, 3} respectively.

Table 2 – Origins for Inbound Vehicles

Typical Locations	Approx. Distance from Courier Hub (km)	Distance from PI-Hub (km)
Ashfield	9	10
Chatswood	10	
Burwood	12	
Olympic Park	15	
Bankstown	20	20
Pemulwuy	30	30
Liverpool	33	

Origins for inbound vehicles established via survey of Courier Hub users.

Table 3 - Origins for Outbound Vehicles

Typical Locations	Approx. Distance from Courier Hub (km)	Distance from PI-Hub (km)
Pitt Street	1	1
Castlereagh Street	1	
Martin Place	1	
O'Connell Street	1.2	
Wynyard	1.5	2
Barangaroo	2.5	
Pyrmont	3	3

Origins for outbound vehicles established via survey of Courier Hub users.

Step 3: Schedule Development

In Step 3, scheduled arrival and departure times were defined. Key parameters used are defined as per Table 4 below.

It is important to note that it was necessary to synthetically generate the additional necessary parameters since they were not directly available for research purposes. In practice, they could be ascertained via various technologies including but not limited to: GS1S4T pings; GPS for carrier-based fleet location data (shared via EPCIS: Electronic Product Code Information Standard); ICT for infrastructure based ETA updates.

Table 4 – Key Parameters

Parameter	Description	Observed or Synthesised
ADT	Actual Departure Time	Synthesised
AAT	Actual Arrival Time	Observed
ATT	Actual Travel Time	Synthesised
STT	Scheduled Travel Time	Synthesised
SAT	Scheduled Arrival Time	Synthesised
$RSAT$	Rounded Scheduled Arrival Time	Synthesised
SDT	Scheduled Departure Time	Synthesised
CF_{Hourly}	Congestion Factor (Hourly)	Synthesised
CF_{Daily}	Congestion Factor (Daily)	Synthesised
$CF_{Monthly}$	Congestion Factor (Monthly)	Synthesised
e_1, e_2	Random value between [-5,5], Normal Distribution	-
v_{av}	Average Speed	Synthesised
D	Distance between origin and destination PI-Hub	Synthesised
d	Distance between current location and destination PI-Hub	Synthesised

Scheduled arrival (at PI-Hub) and scheduled departure (from origin) times were calculated first. Actual Departure Time (ADT) calculated via equation 1, where Actual Arrival Time (AAT) was ascertained through direct observation via video analysis. e_1 is a normally distributed value with a mean of 0 and S.D of 5 mins.

$$ADT = AAT - ATT \quad (1)$$

$$where\ ATT = STT * CF_{Hourly} + e_1 \quad (2)$$

$$\&\ STT = \frac{D}{v_{av}} \quad (3)$$

Similarly, Scheduled Arrival Time (SAT) was calculated using Equation 4, where AAT was observed via video analysis:

$$SAT = AAT + Earliness/Delay \quad (4)$$

$$where\ Earliness/Delay = ATT - STT \quad (5)$$

Scheduled Departure Time (SDT) was calculated by equation 6. e_2 is a random value normally distributed, with a mean of 0 and S.D of 5 mins.

$$SDT = SAT - STT + e_2 \quad (6)$$

It is important to note that to ensure discrepancy between ADT & SDT (as arises in practice), $e_1 != e_2$. e_1 determines ATT, which in turn determines ADT. e_2 is directly used to distort STT which is used to determine SDT.

Step 4: Real Time Locations

In Step 4, real-time, per minute locations were calculated. The distance travelled per minute followed a normal distribution, with mean as calculated by equation 7, and standard deviation of 10%. The last minute was allocated the remaining distance to be travelled ensuring that the overall travel time equated to the actual travel time as determined above.

$$Mean\ Distance/Minute = \frac{D}{ATT}$$

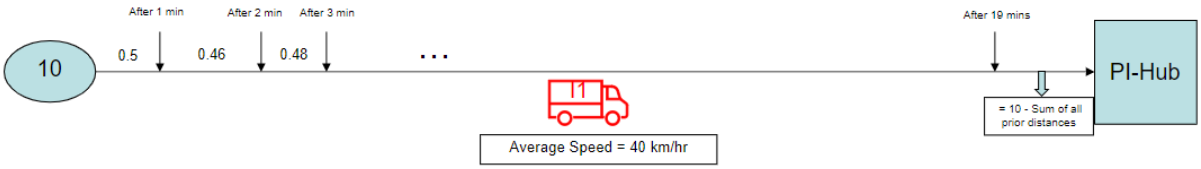


Figure 2 – Real Time Locations

$$ATT = STT * CF_{Hourly} + e_1$$

$$ATT = \frac{10\ [km]}{40[\frac{km}{hr}]} * 1.1 + 4\ [min]$$

$$ATT = 20.5\ mins$$

$$Mean\ Distance/Minute = \frac{10\ [km]}{20.5\ [mins]}$$

$$Mean\ Distance/Minute = 0.49\ [\frac{km}{min}]$$

$$Standard\ Deviation = 0.49 * 0.1 = 0.049$$

Figure 2 & the subsequent equations provide an example of these steps. Inbound vehicle I1 travelling at an average speed of 40 km/hr takes 20.5 mins to travel 10 km from origin to PI-Hub. It's mean distance/minute is calculated to be 0.49 km/min with a standard deviation of 0.049. 0.5 km are allocated to the first minute, 0.46 km to the second minute, and 0.48 km to the third minute. Once the last minute of the journey is reached, the remainder sum is allocated, calculated by subtracting the cumulative distance travelled up to the 19th minute mark from the overall distance between the origin and the PI-Hub, 10 km.

6. Results

Courier Hub data from video analysis spans 13 days across 7 months. The scenarios derived from the first 11 days were used for training the ML & DRL based models. Scenarios of the last two days were used for evaluation purposes.

Two use cases of GS1S4T are considered:

- Improving Visibility: real-time tracking of goods and vehicles as they move through the supply chain including transit facilities;
- Leveraging GEO Location Data: National Location Registry (NLN) is a central database containing precise location of pickup and delivery sites, along with additional attributes such as: safety constraints such as requirements of wearing hard hats; weight/height restrictions; dedicated entry/exit points for specific vehicle types. GS1S4T enables ready access to this database, enabling for more efficient pickup and delivery process.

6.1 Use Case: Improving Visibility

Use Case of improvising visibility investigates the impact of leveraging real-time location updates of PI-Transporters through the road network. Figures 3 and 4 demonstrate the superior performance of models where there are real-time updates, with the average waiting time and total time in PI-Hub consistently less than against the base cases of Naïve and FIFO. Thus, PI-Hubs leveraging real-time vehicle location data for slot reallocations can enhance operations as compared to the Naïve and FIFO (original schedule) base models.

Further, it is evident from these test results that GS1S4T's single location update is suitable for performing as well as the GPS (every 5 mins) and Infrastructure (every 1 min) based rescheduling. Couriers and PI-Hub operators opting for GS1S4T can expect to achieve similar levels of performance as that achieved by GPS and Infrastructure based rescheduling operations. Increased variability in travel times may impact relative performance of the three granularities and needs to be further investigated.

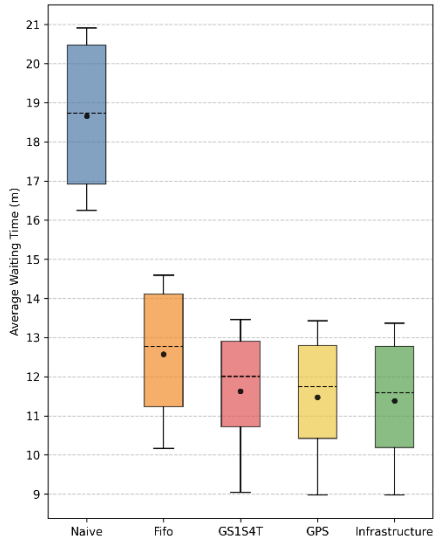


Figure 3 – Use Case: Improving Visibility. Average Waiting Times. Model Comparisons with configuration of 6 hr time periods, 2 PI-Docks & max pairings of 8 PI-Transporters.

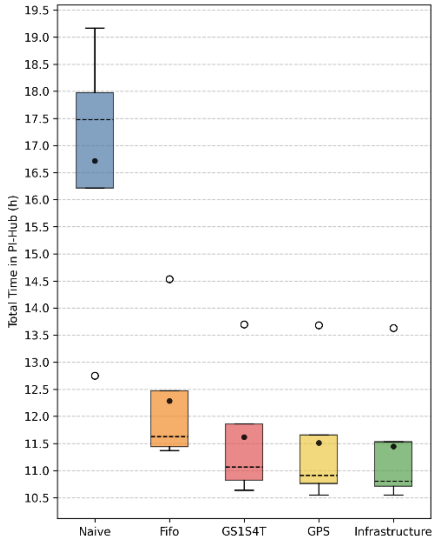


Figure 4 – Use Case: Improving Visibility. Total Time in PI-Hub. Model Comparisons with configuration of 6 hr time periods, 2 PI-Docks & max pairings of 8 PI-Transporters.

6.2 Use Case: Leveraging GEO-Location Data

Use case of leveraging GEO-Location data investigates the role of having access to NLN for conducting deliveries. Without GEO-location data, there are unexpected delays experienced by some vehicles.

Figure 5 & 6 show average waiting times and total time in PI-Hub respectively. It is evident that Naïve performs the worst with a median of approximately 18.75 minutes of average waiting time for each vehicle. Meanwhile, the other base model, FIFO, performs significantly better, with a median of approximately 12.8 minutes. GS1S4T not only has a lower median of 12 minutes of waiting time, but its spread is also the narrowest, indicating that its approach is much more consistent and therefore reliable.

A similar relative performance is found between the three models for the total time in PI-Hub. Naïve has a median of 17.5 hrs, while FIFO performs significantly better with a median of approximately 11.6 hrs. GS1S4T performs best with median of just over 11.5 hrs. Further, with the spread for GS1S4T being the narrowest, it again exemplifies its superior stability in performance.

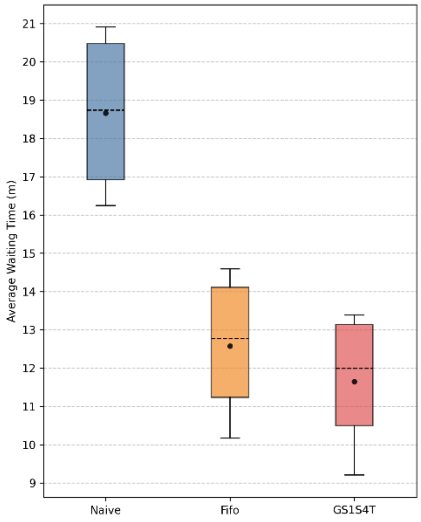


Figure 5 – Use Case: Leveraging GEO-Location data. Average Waiting Times. Model Comparisons with configuration of 6 hr time periods, 2 PI-Docks & max pairings of 8 PI-Transporters.

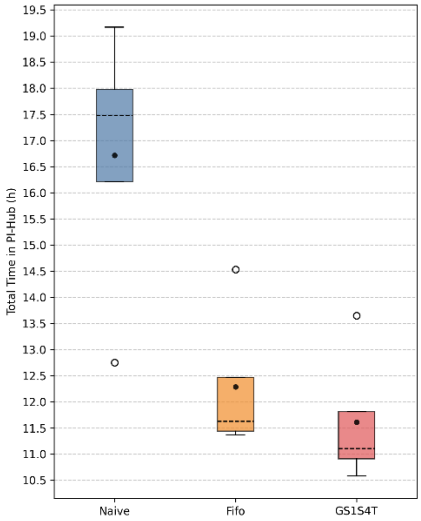


Figure 6 – Use Case: Leveraging GEO-Location data. Total Time in PI-Hub. Model Comparisons with configuration of 6 hr time periods, 2 PI-Docks & max pairings of 8 PI-Transporters.

7. Conclusion

This study proposes a Dynamic Slot Reallocation Framework for PI-Transporters at PI-Hubs to solve the DSRP. Two GS1S4T use cases were tested: Improving Visibility; and Leveraging GEO-Location Data. It was found GS1S4T data sharing technology for fleet location sharing from independent carriers can assist in solving the dynamic slot reallocation problem for PI-Hub operators. Across both use cases, it was found that GS1S4T outperformed the base models Naïve and FIFO. All real-time location sharing models outperformed the Naïve and FIFO base models. Further, GS1S4T performed at similar levels to the GPS and Infrastructure based granularities. This suggests that PI-Hub operators may extract as much a benefit from GS1S4T than with investment and integration of either GPS or Infrastructure based technologies.

References

Acar, K., Yalcin, A., & Yankov, D. (2012). Robust door assignment in less-than-truckload terminals. *Computers & Industrial Engineering*, 63(4), 729–738.

Ballot, E., Montreuil, B., & Zacharia, Z. G. (2021). Physical Internet: First results and next challenges. *Journal of Business Logistics*, *42*(1), 101–107.

Chen, G., Govindan, K., and Yang, Z. (2013). Managing Truck Arrivals with Time Windows to Alleviate Gate Congestion at Container Terminals. *International Journal of Production Economics*. 141, 179–188. Meta-heuristics for manufacturing scheduling and logistics problems. doi:10.1016/j.ijpe.2012.03.033

Chen, K., Xu, G., Xue, F., Zhong, R. Y., Liu, D., & Lu, W. (2017). A Physical Internet-enabled Building Information Modelling System for prefabricated construction. *International Journal of Computer Integrated Manufacturing*, *31*(4–5), 349–361. https://doi.org/10.1080/0951192X.2017.1379095

Crainic, T. G., & Montreuil, B. (2016). Physical Internet Enabled Hyperconnected City Logistics. *Transportation Research Procedia*, *12*, 383–398.

Gkiotsalitis, K. (2021). Bus rescheduling in rolling horizons for regularity-based services. *Journal of Intelligent Transportation Systems*, *25*(4), 356–375. https://doi.org/10.1080/15472450.2019.1681992

Hakimi, D., Montreuil, B. and Labarthe, O. (2009). Supply web: concept and technology. *7th Annual International Symposium on Supply Chain Management.* Toronto, Canada.

Hakimi, D., Montreuil, B., & Labarthe, O. (2010). Supply Web Agent-Based Simulation Platform. *International Conference on Information and Systems, Logistics and Supply Chain.* Casablanca, Morocco.

Hillerstrom, F., Kruithof, M., & Oggero, S. (2018). Arriving on Time Using Uncertainty Aware Deep Learning. *IPIC 2018. International Physical Internet Conference*. Groningen, Netherlands.

Konur, D., & Golias, M. M. (2013). Analysis of different approaches to cross-dock truck scheduling with truck arrival time uncertainty. *Computers & Industrial Engineering*, 65 (4), 663–672.

Ladier, A.-L., & Alpan, G. (2016). Robust cross-dock scheduling with time windows. *Computers & Industrial Engineering*, 99, 16–28.

Le Roch, Y., Ballot, E., & Perraudin, X. (2015). A New Framework for the Management of Returnable "Containers" Within Open Supply Networks. In T. Borangiu, A. Thomas, & D. Trentesaux (Eds.), *Service Orientation in Holonic and Multi-agent Manufacturing* (Vol. 594, pp. 293–305). Springer International Publishing. https://doi.org/10.1007/978-3-319-15159-5_27

Li, C.-L., & Li, F. (2020). Rescheduling production and outbound deliveries when transportation service is disrupted. *European Journal of Operational Research*, *286*(1), 138–148. https://doi.org/10.1016/j.ejor.2020.03.033

McKinnon, A. C. (Ed.). (2015). *Green logistics: Improving the environmental sustainability of logistics* (Third edition). Kogan Page.

Nasiri, M. M., Ahmadi, N., Konur, D., & Rahbari, A. (2022). A predictive-reactive cross-dock rescheduling system under truck arrival uncertainty. *Expert Systems with Applications*, *188*, 115986. https://doi.org/10.1016/j.eswa.2021.115986

Prakoso, E. F., Maknoon, Y., Pel, A., Tavasszy, L. A., & Vanga, R. (2022). A Predictive–Proactive Approach for Slot Management of a Loading Facility With Truck ETA Information. *Frontiers in Future Transportation*, *3*, 815267. https://doi.org/10.3389/ffutr.2022.815267

Rahmani, D., & Ramezanian, R. (2016). A stable reactive approach in dynamic flexible flow shop scheduling with unexpected disruptions: A case study. *Computers & Industrial Engineering*, *98*, 360–372. https://doi.org/10.1016/j.cie.2016.06.018

Scan4Transport Pilot Report. (2021). GS1.

Taniguchi, E., & Thompson, R. G. (2001). *City Logistics: Network Modelling and Intelligent Transport Systems* (1st edition). Emerald Group Publishing Limited.

Wolbeck, L., Kliewer, N., & Marques, I. (2020). Fair shift change penalization scheme for nurse rescheduling problems. *European Journal of Operational Research*, *284*(3), 1121–1135. https://doi.org/10.1016/j.ejor.2020.01.042

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE

Optimization and Resilience Assessment of Multi-Agent Delivery Using Evolutionary Neural Networks in the context of Physical Internet

Hisatoshi Naganawa[1] and Enna Hirata[2]
1. International Maritime Center, Kobe University, Kobe, Japan
2. Graduate School of Maritime Sciences, Kobe University, Kobe, Japan
Corresponding author: enna.hirata@platinum.kobe-u.ac.jp

Abstract: *The Physical Internet (PI) paradigm aims to realize efficient and sustainable global logistics networks by orchestrating Industry 4.0 technologies. This study addresses the multi-agent delivery problem in a PI environment, proposing a decision-making method using an evolutionary neural network. Each delivery vehicle operates under multiple constraints, including battery level, loading capacity, cargo type, and driving time limits, while performing route planning that accounts for loading/unloading operations and cross-docking at distribution centers and hubs. The proposed method evolves a recurrent neural network for vehicle control using neuroevolution of augmenting topologies (NEAT), comprehensively evaluating multiple metrics including delivery efficiency, number of undelivered modules, number of service level agreement (SLA) violations, and CO_2 emissions. Furthermore, it quantitatively analyzes the fault tolerance of the PI system through sensitivity analysis and resilience evaluation in the event of vehicle failures or route disruptions. Simulation results demonstrate that the P-based delivery strategy is effective in terms of both efficiency and resilience.*

Keywords: *Physical Internet; Multi-Agent Delivery; NeuroEvolution of Augmenting Topologies; Logistics Optimization; Resilience Analysis; Sustainable Transportation*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☒ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☒ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☐ *In-Person presentation*

1 Introduction

The Physical Internet (PI) proposes a standardized, open logistics paradigm to solve the inefficiencies of traditional hub-and-spoke models (Montreuil et al., 2010, Pan et al., 2021, Naganawa et al., 2024). However, practical PI operations face high complexity due to mixed-fleet dynamics, EV charging constraints, and time-sensitive requirements (Ambra et al., 2019, Betti Sorbelli, 2024). Traditional centralized optimization often lacks the agility to handle real-world disruptions like vehicle breakdowns or road closures. While distributed AI models show promise(Liu et al., 2023), many rely on static neural architectures that struggle with the fluid nature of PI. This study proposes a distributed multi-agent delivery method utilizing NEAT (NeuroEvolution of Augmenting Topologies). By evolving both the structure and weights of recurrent neural networks, each vehicle agent learns to make autonomous decisions under constraints like battery life, load capacity, and cargo compatibility. A core focus is resilience.

We introduce a framework to quantify performance under stress, such as vehicle failures and network disruptions. Our main contributions include:

1. **PI Formulation** Modeling dynamic, multi-constraint, heterogeneous delivery.
2. **Decision Model** Distributed control via evolutionary recurrent neural networks.
3. **Multi-Purpose Index** Balancing costs, service quality, and environmental impact.

2 Construction of a Delivery Decision Model Using NEAT

This chapter details the NEAT algorithm, which forms the core of the proposed method, and its application in this research.

2.1 Basic Principles of NEAT and Reasons for Selection

NEAT is a method that simultaneously evolves the "weights" and "topology" of neural networks using genetic algorithms (Stanley and Miikkulainen, 2002). Unlike conventional fixed-layer networks, it starts from the minimal configuration of only inputs and outputs, automatically adding hidden layer nodes and connections as needed. This enables the construction of optimal models that are neither excessive nor insufficient for multivariate and nonlinear problems like those in PI environments. Furthermore, when crossing networks with different structures, it incorporates a mechanism to trace the lineage of genes, preventing fatal defects caused by structural mismatches.

2.2 The Effectiveness of NEAT in PI Environments

Complex delivery planning in PI environments requires simultaneous consideration of dynamic and uncertain factors such as vehicle breakdowns and road closures, in addition to constraints like remaining battery charge, load capacity, and delivery time windows (Mahmoodi et al. 2025, Amrovani et al. 2025). Consequently, methods based on predefined rules or fixed network structures often prove insufficient for adequately responding to sudden environmental changes. The NEAT algorithm employed in this research demonstrates high effectiveness in such PI environments. A major advantage of NEAT is its ability to evolve not only network weights but the topology itself. This enables it to acquire flexible decision criteria tailored to the situation, even when sudden route changes or unexpected problems occur. Consequently, it facilitates adaptive decision-making that was difficult with fixed-structure models. Furthermore, NEAT has high affinity with multi-objective optimization. It can integrate multiple evaluation criteria as fitness functions, including not only efficiency metrics like delivery distance and time, but also environmental impact indicators such as the number of undelivered items, suppression of SLA violations, and CO_2 emissions. Consequently, it enables learning realistic and balanced delivery plans that avoid bias toward a single metric.

3 Formulation and System Concept

This chapter formulates the dynamic delivery planning problem under PI environment as the Capacitated Multi-Vehicle Routing Problem with Time Windows (CMVRPTW). This problem is a dynamic stochastic optimization problem that considers a mixed fleet of heterogeneous vehicles, diverse cargo characteristics, and operational constraints of electric vehicles.

3.1 Design Philosophy of Fully Decentralized Decision-Making

The core concept underpinning this research is "fully decentralized decision-making" independent of a central authority. Within the PI network, each vehicle functions as an independent intelligent agent. Each agent detects dynamic changes in the network in real time—such as new delivery requests, charging station congestion, sudden vehicle failures, or road

closures—and autonomously determines the optimal next node to proceed to via its own neural network. This approach prevents load concentration on specific servers while achieving high robustness, ensuring that partial vehicle failures do not halt the entire system.

3.2 Networks and Components

The delivery network is represented by a directed graph $\mathcal{G} = (\mathcal{N}, \mathcal{A})$. Here, $\mathcal{N} = \{0, 1, \dots, N\}$ is the node set, and $\mathcal{A} = \{e_1, \dots, e_M\}$ is the edge set.

3.2.1 Node set ($\boldsymbol{\mathcal{N}}$)

Node sets are classified into the following four types.

$$\mathcal{N} = \mathcal{N}_{dep} \cup \mathcal{N}_{hub} \cup \mathcal{N}_{cst} \cup \mathcal{N}_{charger} \quad (1)$$

Each node $i \in \mathcal{N}$ has coordinates (x_i, y_i), service time δ_i, and node capacity C_i. In particular, the hub node $\mathcal{N}_{hub}$ possesses cross-docking functionality and serves as a temporary cargo buffer.

3.2.2 Vehicle Assembly ($\boldsymbol{\mathcal{V}}$)

Define a heterogeneous vehicle set $\mathcal{V} = \{1, \dots K\}$. Each vehicle $k \in \mathcal{V}$ has the following attributes.

Q_k	Maximum load capacity
B_k	Battery capacity
ρ	Energy consumption rate
T_k	Types of Cargo That Can Be Loaded (Refrigerated, Hazardous Materials, General Cargo)

3.2.3 Module Assembly ($\boldsymbol{\mathcal{M}}$)

Let $\mathcal{M} = \{1, \dots, M\}$ denote the set of modules to be delivered. Each module $m \in \mathcal{M}$ has an origin o_m, a destination d_m, a size s_m, a delivery time t_m, and a type τ_m, and moves dynamically within the network.

3.3 Dynamic States and Decision Making Using Neural Networks

This study adopts a sequential decision-making model involving distributed vehicles rather than static route planning. Let $S_{k,t}$ denote the state vector of vehicle k at time t. When vehicle k is at node i, the policy π_θ for determining the next node $j \in \mathcal{N} \mid \{i\}$ to visit is approximated by a neural network with parameter θ.

$$j^* = \arg\max_{j \in \mathcal{N}} f_\theta(x_{i,j,k,t}) \quad (2)$$

Here, the input feature vector $x_{i,j,k,t}$ is composed of the following elements, which integrate local and global information.

$$x_{i,j,k,t} = \left[\tilde{d}_{ij}, \frac{t}{T_{max}}, \eta_k, \delta_j, \phi_j, \frac{b_{k,t}}{B^{max}}, \psi_{k,t}, \mathbb{I}_{charge}(j), \mathbb{I}_{comp}(k,j)\right]^\top \quad (3)$$

3.4 Constraints

The system must comply with the following constraints.

(1) Battery and Range Constraints: Let $\Delta E_{ij} = d_{ij}\rho$ denote the energy required for vehicle k to move from node i to node j. It must always hold that $b_{k,t} \geq \Delta E_{ij}$. If reaching the destination is impossible, the vehicle retreats to the nearest chargeable node $h \in \mathcal{N}_{charger}$. If retreat is also impossible, the vehicle becomes dysfunctional.

(2) Loading Compatibility Constraints: For vehicle k to load module m, it must satisfy the following exclusive control in addition to the capacity constraint.

$$\tau_m \in T_k \wedge \forall l \in \mathcal{L}_k, \neg\, \text{Conflict}(\tau_m, \tau_l) \quad (4)$$

Here, $\mathcal{L}_k$ denotes the set of modules being loaded, and $\text{Conflict}(\cdot)$ is a function defining prohibited combinations such as "frozen goods" and "hazardous materials."

(3) Working Hours Constraints: A driver's continuous driving time $\tau_{k,t}^{drive}$e must not exceed the legal limit T_{limit}. When the limit is reached, a break time $T_{break}$is required.

3.5 Objective Function

The objective of this problem is to maximize the sustainability and service level of the PI. This is equivalent to minimizing the following cost function J.

$$Minimize\ J = \sum_{k \in \mathcal{V}} C_{op}(k) + \sum_{m \in \mathcal{M}} C_{serve}(m) + P_{env} \tag{5}$$

Details for each item are as follows.

(1) Operating cost C_{op}:

$$C_{op}(k) = \omega_1 D_k + \omega_2 T_k^{over} \tag{6}$$

Here, D_k denotes the total distance traveled, and T_k^{over} indicates overtime hours exceeding the specified time.

(2) Service cost $C_{service}$:

$$C_{service}(m) = \omega_3 \mathbb{I}_{miss}(m) + \omega_4 \max(0, t_m^{del} - t_m^{due}) + \omega_5 \Delta t_m^{dwell} \tag{7}$$

Item 1 represents penalties for undelivered shipments, Item 2 represents penalties for delivery delays (SLA violations), and Item 3 represents penalties for excessive dwell time at the hub.

(3) Environmental cost P_{env}:

$$P_{env} = \omega_6 \sum_{k \in \mathcal{V}} E_k^{CO2} \tag{8}$$

E_k is the CO2 emissions calculated in proportion to the load weight and distance.

4 Simulation Setup

To verify the effectiveness of the proposed method, we constructed a simulation environment based on the actual logistics infrastructure in Hyogo Prefecture, Japan.

4.1 Geographic Data and Node Network

This study conducted simulations covering the entire Hyogo Prefecture, encompassing diverse logistics characteristics from urban to suburban areas.

Table 1. Node type

Type	Description
PI Depot	The starting point of the logistics network. Designated as a major center for Sagawa Express.
PI Hubs	A central hub within the prefecture. Designated as a Sagawa Express branch office. Promotes large-scale consolidation and cross-docking.
PI Nodes	As the end-points of the Physical Internet, we extracted location data for approximately 2,500 public parcel lockers across the prefecture.

4.2 Agent and Environmental Parameters

The specific parameter settings used in the simulation are shown in the table below.

Table 2. System Scale and Temporal Constraints

Category	Parameter (Unit)	Value / Details
Network Scale	Customers / Vehicles	100,000 / 200 units
Time Windows	Business Hours / Max Time	08:00–19:00 / 24:00 (min)
Work Limits	Max Drive / Break / Total	200 / 30 / 1000 (min)
Nodes (Svc Time)	Depot / Hub / Customer	10 / 5 / 2 (min)

Table 3. Fleet and Cargo Specifications

Type	Capacity / Battery	Rate / Weight	Size Range
Small Van / Normal	30 / 120 kWh	0.8 / 65%	(4, 8)
Med Truck / Chilled	30 / 200 kWh	1.2 / 20%	(6, 10)
Refrigerated / Frozen	30 / 180 kWh	1.5 / 10%	(8, 12)
Express / Dangerous	30 / 100 kWh	0.9 / 5%	(1, 3)

Table 4. Reward and Penalty Weights

Type	Factor	Weight	Objective
Rewards	Load Fill / Loop	8.0 / 4.0	Efficiency & Throughput
Penalties	Stranded (Battery)	5000.0	System Reliability
	SLA / Hub Dwell	2.0 / 1.0	Service Quality
	CO2 / Energy	0.5 / 5.0	Sustainability

5 Result

This study conducted simulations targeting Hyogo Prefecture to evaluate the effectiveness of the proposed method. The following outlines key findings regarding geographical conditions and agent behavior.

5.1 Distribution and Network Formation

The locations where delivery demand originates in the simulation were set based on the population density distribution of Hyogo Prefecture. In this study, the points of origin for delivery demand in the simulation were set based on Hyogo Prefecture's population density distribution. As shown in the population density map in Figure 1, demand is highly concentrated in coastal urban areas such as Himeji, Akashi, Kobe, and Nishinomiya. Under these geographical constraints, examining the visualization results of the multi-agent delivery routes shown in Figure 2 demonstrated that the evolved agents autonomously formed an efficient network targeting high-demand areas. Specifically analyzing route dynamics reveals a dynamic "cross-docking" operation: each vehicle departing from a depot (red square) passes through multiple PI hubs (blue triangle) before delivering to the final target customer (green dot). This suggests that each agent can appropriately select and utilize hubs as strategic relay points by considering its own loading capacity and cargo compatibility in real time.

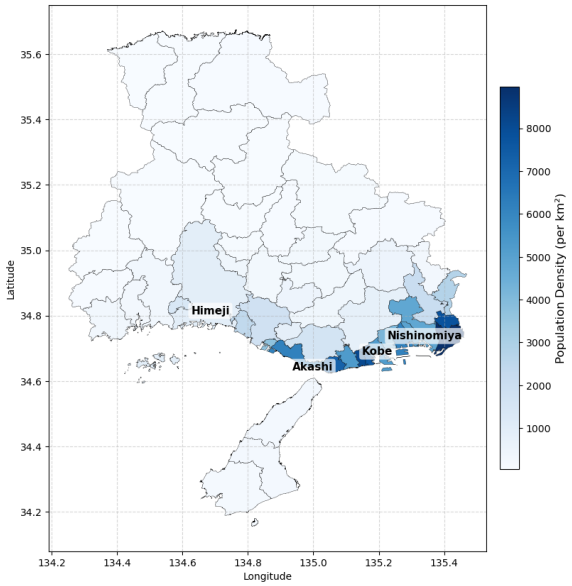


Figure 1. Population Density

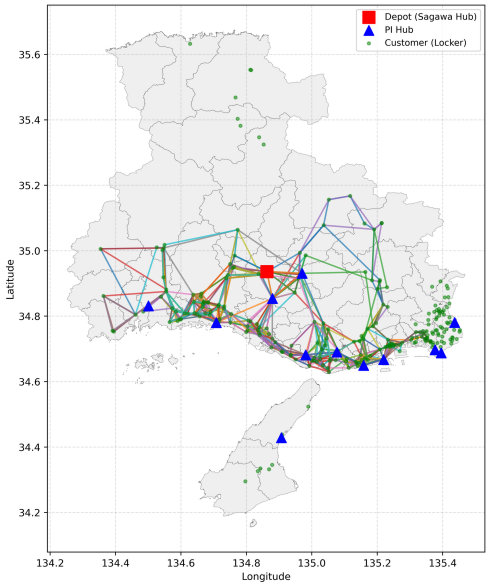


Figure 2. Nodes and Network

5.2 Fitness Changes in Evolutionary Processes

The fitness evolution in Figure 3 clearly illustrates how NEAT progressively approaches the optimal solution in the complex mathematical optimization problem discussed here. First, the overall trend of the graph shows that the global maximum fitness, indicated by the red dashed line, increases with each generation. The particularly rapid improvement during the first 10 generations indicates that the system swiftly acquired a structure capable of satisfying basic delivery constraints despite starting from a random initial state. Subsequently, the discontinuous, step-like improvements observed around generations 40 and 80 represent the discovery of innovative strategies that break out of existing local optima. This is achieved through NEAT's characteristic addition of new topologies and species differentiation. Meanwhile, the sharp fluctuations in the solid green line (the maximum fitness per generation) are evidence of vigorous exploration within the evolutionary process. While adding new neurons or connections carries the “evolutionary cost” of temporarily lowering fitness, NEAT's protection mechanism safeguards these new structures, enabling them to contribute to optimization in subsequent generations. Thus, the fact that the global maximum is steadily updated despite significant fluctuations confirms convergence toward an extremely robust decision-making model through trial and error of diverse strategies.

5.3 Correlation between Loading Efficiency and Environmental Impact

Figure 3 shows the changes in the load capacity of each vehicle over time. As evident from the graph, the values for each agent rose rapidly immediately after operation began and subsequently fluctuated within a certain range. This suggests that agents not only travel the shortest distance but also continue operation while maximizing load capacity and minimizing empty runs whenever possible. The fact that the distributed control model evolved by NEAT autonomously optimizes both cost reduction and environmental impact reduction is considered an extremely important achievement in building a sustainable logistics network.

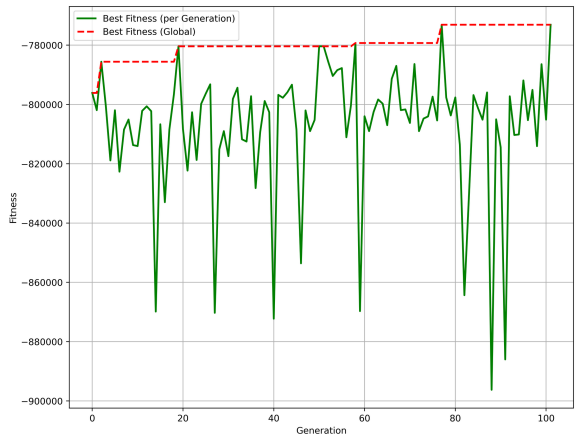


Figure 3. Fitness over Generations

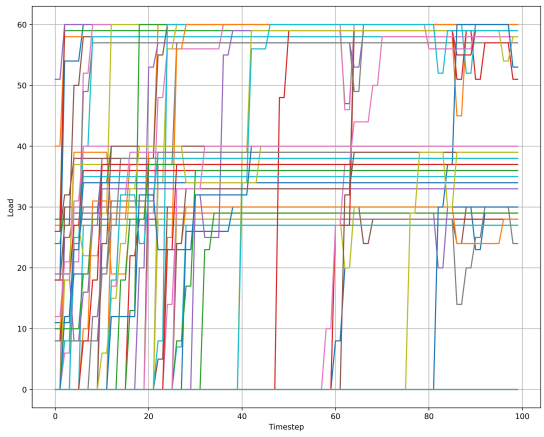


Figure 4. Vehicle Load

5.4 Vehicle Fleet Operational Dynamics and Resource Allocation

The heatmap group in Figure 5 visualizes the resource consumption and operational status of all vehicle agents over time, providing key insights into the operational efficiency of the proposed method.

5.4.1 Optimization of Load Factor and Differences Between Vehicles

The “Vehicle Load” heatmap on the left clearly shows the load factor trends for each vehicle ID. While a specific group of vehicles (the dark blue area in the center) consistently maintains a high load rate (around 60%), other vehicles dynamically adjust their load volume based on delivery conditions. This demonstrates that agents evolved by NEAT do not operate uniformly across the entire fleet. Instead, they autonomously adjust the system's overall loading efficiency by dividing roles according to each vehicle's current location, performance, and surrounding demand density.

5.4.2 Driving Time Management and Compliance with Labor Restrictions

The central “Vehicle Overtime” heatmap shows a gradual increase in overtime hours across all vehicles over time, with no sudden spikes. This is likely the result of each agent selecting routes that consider driving time limits while ensuring no single vehicle bears an excessive load. This suggests the potential for stable service delivery while maintaining quality working conditions, even within large-scale delivery networks.

5.4.3 Stability of Energy Consumption and Range Management

The heatmap on the right labeled “Vehicle Battery Level” shows that within a fleet of vehicles with varying battery specifications, each group maintains stable energy consumption patterns. This planned energy expenditure aligned with delivery tasks confirms that agents are learning efficient power management strategies to minimize “downtime for charging” while achieving continuous operation.

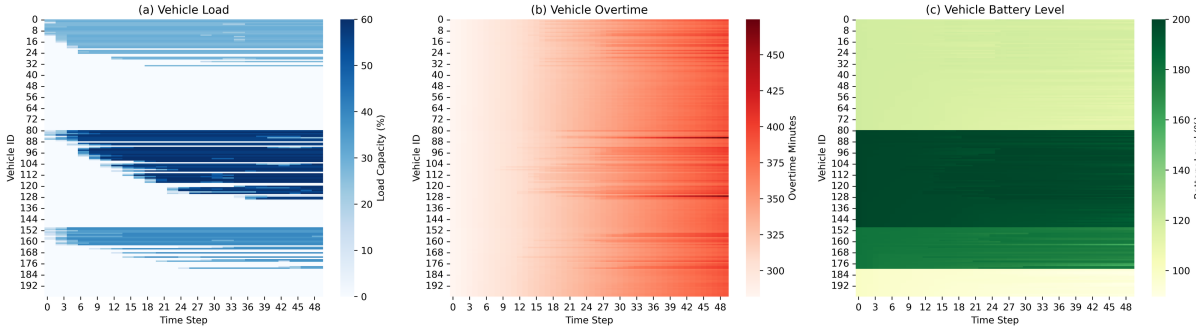


Figure 5. Fleet Performance

6 Conclusion

This study proposed an autonomous decentralized decision-making model using NEAT to optimize multi-point delivery and enhance resilience within a PI environment.

6.1 Research Findings

The simulations and analyses conducted in this study yielded significant insights regarding delivery optimization within a PI environment. First, the introduction of a recurrent neural network evolved using NEAT demonstrated that each delivery vehicle can autonomously select routes responsive to dynamic environmental changes, even under multiple constraints such as remaining battery charge, load capacity, cargo compatibility, and legal working hours. This ensures the flexibility for real-time decision-making that was difficult with conventional fixed algorithms. Furthermore, the proposed method successfully balanced multiple objectives at a high level, going beyond mere reduction of delivery distance or time. It effectively minimized SLA violations and undelivered modules while also reducing CO_2 emissions. This confirmed that the PI philosophy, aiming to achieve both economic efficiency and reduced environmental impact, can be embodied as a concrete delivery strategy.

6.2 Societal and Academic Significance

This research demonstrates that AI using evolutionary computation is highly effective in addressing uncertainties difficult to handle with fixed rule-based systems within PI, a sustainable logistics foundation for the industry 4.0 era. Specifically, the “fully distributed model,” where vehicles act as independent agents, significantly contributes to distributing computational load in large-scale, complex logistics networks and maintaining logistics networks during emergencies such as disasters.

6.3 Future Prospects

Future challenges include expanding to more realistic scenarios that incorporate intermodal transport, including coordination with freight rail and shipping. Additionally, game theory methods will be introduced to optimize coordination of interests among shippers and model the consensus-building process. Furthermore, plans are in place to refine the model through verification using real-world data, establish evaluation metrics for resilience against unforeseen events such as disasters and accidents, and advance preparations for societal implementation.

Appendix

Algorithm 1: NEAT-based Logistics Policy Optimization

procedure NEAT_TRAINING(*Inputs, Outputs*)
  Initialize NEAT population P with minimal topologies
  while termination condition is not met **do**
    for each individual π_θ **in** P **do**
      Initialize environment $E(\theta)$: $G=(N,A)$, V, M
      for $t = 0$ **to** T **do**
        for each available vehicle k **in** V **do**
          $i \leftarrow GetCurrentNode(k)$, $CandidateNodes \leftarrow N \setminus \{i\}$
          for each j **in** $CandidateNodes$ **do**
            $x_{i,j,k,t} \leftarrow [d_{ij}, \eta_k, \delta_j, \phi_j, b_{k,t}, \psi_{k,t}, I_{charge}, I_{comp}]^T$
            $score_j \leftarrow f_\theta(x_{i,j,k,t})$
          end for
          $FeasibleNodes \leftarrow FilterConstraints(CandidateNodes, b, \tau, Capacity)$
          if $FeasibleNodes$ is empty **then**
            $j^* \leftarrow nearest_charger_or_dysfunctional(k)$
          else
            $j^* \leftarrow argmax_{j \in FeasibleNodes}\ (score_j)$
          end if
          UpdateVehicleState(k, j^*)
        end for
        UpdateEnvironmentState(t)
      end for
      $Fitness(\pi_\theta) \leftarrow 1.0 / (1.0 + CalculateTotalCost(V, M))$
    end for
  $P \leftarrow NEAT_Evolution(P)$
  end while
  return π^*_θ
end procedure

References

- Ambra, T., Caris, A., Macharis, C., 2019. Towards freight transport system unification: reviewing and combining the advancements in the physical internet and synchromodal transport research. International Journal of Production Research 57, 1606–1623. https://doi.org/10.1080/00207543.2018.1494392
- Amrovani, M.A., Askarian-Abyaneh, H., Gharibi, M.A., Mozaffari, M., 2025. Urban grid resilience assessment framework: Leveraging electric vehicles, time-based analysis, and mobile distributed generators for repair crew strategic deployment. Sustainable Energy, Grids and Networks 41, 101588. https://doi.org/10.1016/j.segan.2024.101588
- Betti Sorbelli, F., 2024. UAV-Based Delivery Systems: A Systematic Review, Current Trends, and Research Challenges. ACM J. Auton. Transport. Syst. 1, 12:1-12:40. https://doi.org/10.1145/3649224
- Liu, Y., Tao, X., Li, X., Colombo, A.W., Hu, S., 2023. Artificial Intelligence in Smart Logistics Cyber-Physical Systems: State-of-The-Arts and Potential Applications. IEEE Transactions on Industrial Cyber-Physical Systems 1, 1–20. https://doi.org/10.1109/TICPS.2023.3283230
- Mahmoodi, A., Sajadi, S.M., Sadeq, A.M., Narenji, M., Eshaghi, M., Jasemi, M., 2025. Enhancing unmanned aerial vehicles logistics for dynamic delivery: a hybrid non-dominated sorting genetic algorithm II with Bayesian belief networks. Ann Oper Res. https://doi.org/10.1007/s10479-025-06504-z

- Montreuil, B., Meller, R., Ballot, E., 2010. Towards a Physical Internet: the Impact on Logistics Facilities and Material Handling Systems Design and Innovation. Progress in Material Handling Research: 2010.
- Naganawa, H., Hirata, E., Firdausiyah, N., Thompson, R.G., 2024. Logistics Hub and Route Optimization in the Physical Internet Paradigm. Logistics 8, 37. https://doi.org/10.3390/logistics8020037
- Pan, S., Zhang, L., Thompson, R.G., Ghaderi, H., 2021. A parcel network flow approach for joint delivery networks using parcel lockers. International Journal of Production Research 59, 2090–2115. https://doi.org/10.1080/00207543.2020.1856440
- Stanley, K.O., Miikkulainen, R., 2002. Evolving Neural Networks through Augmenting Topologies. Evolutionary Computation 10, 99–127. https://doi.org/10.1162/106365602320169811

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Genetic Algorithm-Based Configuration of Logistics Hubs for Parcel Delivery: A Case Study of the Tokyo Metropolitan Area

Muhammad Ilham Fahreza[1], Sinar Noe Corzo-Garcia[1], Enna Hirata[1], Hiroki Sakai[1], Russell G. Thompson[2], Greg Foliente[2], Masoud Kahalimoghadam[2]

1. Graduate School of Maritime Sciences, Kobe University, Kobe, Japan
2. Department of Infrastructure Engineering, The University of Melbourne, Melbourne, Australia

Corresponding author: 239w205w@gsuite.kobe-u.ac.jp

Abstract: *The logistics sector in Japan is facing multifaceted problems in the present day, exacerbated by the increase in e-commerce activities. The increasing demand for parcel delivery has put a strain on the already overwhelmed network. Tokyo, as one of the metropolitan areas in Japan, suffers from this problem as well, as the number of small parcel deliveries has steadily increased. The current population decline in Japan impacts labour availability. On the other hand, the enactment of work reform legislation limits the working hours for the already strained labour force. Combined with the high demand for parcel delivery, this poses a threat to the current logistics sector. Optimisation of logistics facilities as hubs can fix the current parcel delivery problem. The objective of this study is to optimise the number and location of LSP hubs through a combination of the genetic algorithm (GA) and Taguchi method to produce the hub configuration, while taking other factors into consideration. Through Taguchi-based GA, hubs are reduced from 1543 to 419, corresponding to a 72.84% reduction, whilst reducing average delivery time, distance, and emissions, without reducing delivery capability.*

Keywords: *physical internet, logistics service provider, optimisation, genetic algorithm, Taguchi method*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☒ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☒ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

1 Introduction

The parcel delivery sector in Japan is currently facing pressure from multiple aspects. These include expanding e-commerce activities, high urban demand, shortage of labour, and new working time limitations in work-style reform legislation. The total parcel deliveries in Japan are 5,031,472,000 parcels, an increase of 0.25% from the previous year (MLIT, 2025). The market share for e-commerce has shown a steady increase throughout the years, with a 3.7% increase in FY2024 (METI, 2025). This increases the demand for parcel deliveries. According to Sakai et al. (2024), more than 90% of parcel deliveries are handled by the three logistics service providers (LSPs): Sagawa Express, Yamato Transport, and Japan Post. High demand for deliveries can generally be found in densely populated areas, which leads to a high number of logistics facilities in those areas (Sakai *et al.*, 2024). This includes Tokyo as the capital of Japan. The shrinking labour force is also affecting the sector, with an annual decline rate of

around 0.54% (Fernández-Villaverde, Ventura and Yao, 2025). The Japanese government passed several law reforms, dubbed the 'Work Style Reform-related Laws', in June 2018, which place a cap on maximum working hours (MHLW, 2023; JTA, 2024). This can potentially increase the personnel cost needed for each parcel delivery.

These factors, combined with the current independent, conventional-based logistics network, need to be optimised. This can be achieved through a configuration of hubs in a shared-asset network under a PI-aligned concept to improve the utilisation of existing assets and potentially optimise other performance indicators that include average delivery time, distance, and emissions. In this study, the optimised shared-hub system is interpreted as an early-stage PI network configuration. The model does not assume a fully open PI system yet; rather, it evaluates a realistic intermediate form of PI adoption in which three dominant parcel carriers coordinate hub usage under common capacity and routing rules.

2 Research gaps and objectives

Related studies

Previous studies about hub and network optimisation have been conducted under different conditions. These include multi-hub consolidation, service level, stochastic demand, and spatial and externality factors. Kweon *et al.* (2024) analysed the operation of each hub in the parcel delivery network optimisation within a multi-hub system. Bayram, Aydoğan and Kargar (2024) examined multi-period hub network design with multi-period hub network with congestion, service level, and stochastic demand. For Japan-related research, Sakai, Kawamura and Hyodo (2020) found that each hub placement in the Tokyo capital region is affected by unique spatial and externality factors. Another study by Nakayama and Yan (2019) the accessibility issue in convenience-store-based pickup delivery. However, studies related to the reconfiguration of real logistics facilities as shared PI-aligned logistics nodes are still limited.

Identified research gaps

Despite several recent studies contributing to hub optimisation research to some degree, some gaps still exist. First, limited studies have been done on optimisation and configuration of logistics facilities into shared hubs. Second, studies on hub reduction based on a demographic bias score, combined with multiple performance indicators such as travel time, distance, emissions, and capacity, remain limited. Lastly, PI-oriented studies in Japan, particularly in the Tokyo Metropolitan Area, incorporating real-world data from three major logistics companies can be developed. Conducting a study that focuses on one of the largest economic agglomeration areas in Japan would further the PI-related research in Japan and provide the preliminary path towards similar, if not more comprehensive, research for Tokyo or other areas in Japan in the future.

Research Objectives and Contributions

This study proposes optimising logistics hubs in the Tokyo Metropolitan Area by reducing the number and locations of hubs. Other factors, such as distance, time, and emissions, are also aimed at being reduced to provide a comprehensive optimisation study that takes multiple considerations into account. Finally, this study also provides an empirically grounded PI-aligned optimisation framework that serves as the basis for further study on PI-based logistics networks in the Tokyo metropolitan area. The main contribution is therefore not only the use of GA and Taguchi tuning, but the application of this optimisation framework to a real multi-

operator parcel network, where existing logistics facilities are evaluated as candidate shared PI nodes under demand, cost, capacity, and demographic suitability considerations.

3 Methodology overview and formulation

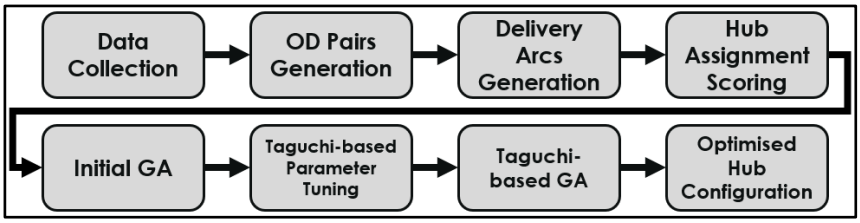


Figure 1: Methodology of the study

The methodology proposed in this study consists of several steps, shown in Figure 1. In this study, parcel movements in the Tokyo Metropolitan Area are considered. Each OD pair represents daily parcel demand between the origin and destination cluster. The notations are used to differentiate multiple parcel flows, origin and destination nodes, and assigned routes. The nodes, both origin and destination, are office and PUDO clusters, which represent businesses and households in a parcel delivery scenario. LSP offices, distribution centres, and PUDO lockers represent candidate hub nodes which can be opened or closed by GA. This explains the duality of PUDO lockers as both OD nodes in a cluster form to represent the residential areas and logistics hubs to represent the suitable candidates for parcel delivery. Overall, the arcs in the network consist of direct cluster–hub and hub–cluster arcs, as well as cluster–PUDO, PUDO–hub, hub–PUDO, and PUDO–cluster arcs. The model works under the assumption that all LSP companies jointly operate the offices under a PI-aligned shared asset concept. This enables a transitional step towards full implementation of PI in the future.

Genetic algorithm

The genetic algorithm (GA) is an optimisation method to improve potential solutions based on numerous iterations through selection, crossover, and mutation processes. (Azadeh *et al.*, 2017). In this study, GA is used to modify a chromosome, which consists of binary vectors, with 1 being an open hub and 0 being a closed one, with length equal to the number of candidate hubs. The optimal LSP hubs configuration that contains the lowest-cost open hubs will be assigned to OD parcel flows for each chromosome.

Taguchi method for parameter tuning

The Taguchi method is a statistical method intended for finding the ideal setting for parameters. Two main parameters are orthogonal arrays (OA), a combination of parameters arranged in a matrix shape, and the signal-to-noise (S/N) ratio, a ratio of wanted and unwanted variations of the response. (Azadeh et al. 2017; Vaikkathe, Benaini and Boukachour, 2025). This study applies the Taguchi method to evaluate GA parameters through an L16 orthogonal array. It evaluates four aspects: population size, maximum number of generations, crossover rate, and mutation rate. To identify the most robust parameter setting, the S/N ratio will be observed.

Mathematical model

3.1.1 Sets and parameters

Sets and parameters can be seen in Table 1.

Table 1: Sets and parameter variables

Symbol	Definition	Symbol	Definition	Symbol	Definition	Symbol	Definition	Symbol	Definition
C	Set of office clusters	$o(k)$	Origin cluster	C_{cl}^{eff}	Effective cost between cluster c and hub l	ϕ_l	Bias score	C_{CAP}	Capacity overflow penalty factor
P	Set of PUDO locker clusters	$d(k)$	Destinatior cluster	$C_{l_1 l_2}^{\text{hub}}$	Interhub cost	T_d^{base}	Baseline transport cost	M_{big}	Large infeasibility penalty
L	Set of LSP hubs candidates	d_k	Daily demand for the parcel k	Cap_l	Daily handling capacity of hub l	H^{base}	Baseline hub cost	z_l	Binary decision variable, 1 if l is open
Symbol	Definition	Symbol	Definition	Symbol	Definition	Symbol	Definition	Symbol	Definition
K	Set of OD parcel-demand	C_{ij}	Truck cost on arc (i, j)	f_l	Fixed opening cost of hub l	W_T, W_H and W_ϕ	Objective-function weights		

The demographic bias score ϕ_l represents the attractive score for candidate hub placement l that is based on population and worker density. A higher value places the hub closer to areas with a higher likelihood of parcel-generation or parcel-receiving activities. This enables consideration of hub selection based not only on cost reduction, but also on retaining them in relevant areas.

3.1.2 Mathematical model

Objective Function:

$$\min J = W_T \cdot \frac{T_d}{T_d^{base}} + W_H \cdot \frac{H}{H^{base}} - W_\phi \cdot \frac{\phi_{hub}}{\Phi_{max}} \quad (1)$$

The costs and biases are defined as:

$$T_d = \sum_{k \in K} d_k \left(C_{eff}[o(k), h(o(k))] + C_{hub}[h(o(k)), h(d(k))] + C_{eff}[d(k), h(d(k))] \right) \quad (2)$$

$$H = C_{CAP} \sum_{l \in L} \frac{\max\{0, Flow_l - Cap_l\}}{Cap_l} + \sum_{l \in L} f_l z_l \quad (3)$$

$$Flow_l = \sum_{k \in K} d_k \left[1_{h(o(k))=l} + 1_{h(d(k))=l} \right], \forall l \in L \quad (4)$$

$$\phi_{hub} = \frac{\sum_{l \in L} \phi_l z_l}{\sum_{l \in L} z_l} \quad (5)$$

3.1.3 Constraints

$$z_l \in \{0,1\}, \forall l \in L \quad (6)$$

$$\sum_{l \in L} z_l \geq 1 \quad (7)$$

$$h(c) = \arg \min_{l \in L: z_l = 1} C_{eff}[c, l], \forall c \in C \quad (8)$$

$$C_{hub}[h(o(k)), h(d(k))] < +\infty, \forall k \in K \quad (9)$$

$$Flow_p^O \leq Cap_{PUDO,p}^O, \forall p \in P \quad (10)$$

$$Flow_p^D \leq Cap_{PUDO,p}^D, \forall p \in P \quad (11)$$

If any feasibility condition in (7)-(12) is violated, then

$$J = M_{big} \quad (12)$$

The objective function in Equation (1) considers a weighted combination of three components for the minimisation process: normalised transport cost, hub-related cost, and demographic reward. Total transport cost T_d is defined in Equation (2), which includes the cluster-to-hub cost, inter-hub cost, and hub-to-cluster cost for all OD parcels. Hub-related cost H is defined in Equation (3), which includes the hub capacity overflow penalty and fixed cost of open hubs. Total throughput handled by the hub is calculated in Equation (4). Lastly, Equation (5) calculates the average demographic suitability for the chosen hubs.

The variable for the binary hub-opening decision is defined by Constraint (6). To ensure that at least one hub is open, Constraint (7) is added. Constraint (8) enforces the routing rule where each cluster is assigned to the cheapest reachable hub. To ensure that every cluster, both for the origin and destination sides, is connected through an inter-hub connection, Constraint (9) is added. To enforce the capacity limit for PUDO lockers located on both the origin and destination sides, Constraints (10) and (11) are added. Lastly, Constraint (13) adds a penalty to any infeasible solution.

Study Area and data

3.1.4 23 wards of Tokyo

The movement of parcel deliveries simulated in this study is based on the intra-ward movement within the Tokyo Metropolitan Area. This is due to the high frequency of movement of parcels within the area and the abundance of LSPs' hubs in the area, which can be further analysed later. Due to this, the area analysed in this study is restricted to the 23 wards of Tokyo that are considered to be the metropolitan area of Tokyo, as shown in Figure 2.

3.1.5 LSPs' hubs datasets

The dataset consists of LSP offices and distribution centres obtained from three company websites via Python-based data scraping: Sagawa Express, Yamato Transport, and Japan Post. They are treated as part of candidate hubs in the optimisation model. The total number of hubs is 1543: 205 hubs for Sagawa Express, 272 hubs for Yamato Transport, and 1066 hubs for Japan Post. The average value is obtained by taking a sample of the hubs owned by each LSP

company in each ward. In total, 69 hubs were sampled, and the average area of each hub is 635.504 m^3. These hubs serve as the real-world data that will be used in this study.



Figure 2: The wards of the Tokyo Metropolitan Area

3.1.6 Supporting data for the analysis

Supporting data are used to support the optimisation analysis conducted in this study. Estimated data include vehicle capacity, parcel size, and demographic-based bias score for hub selection. Vans and trucks are chosen for the parcel delivery task in this study. The smallest available size for a parcel is used, which is the Japanese size-60 parcel. It is the total sum of the length, width, and height (Kase, Murayama and Suzuki, 2023). Population and worker density are used to produce a suitability score to generate a hub configuration that represents expected potential parcel demand for office and residential sides in the Tokyo Metropolitan Area. The density maps can be seen in Figure 3.

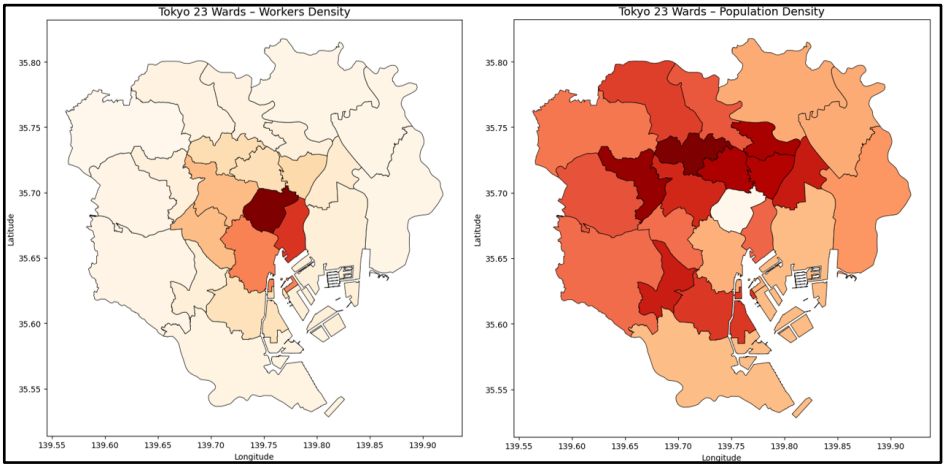


Figure 3: The Workers and Population Density

4 Results and discussions

Genetic algorithm and Taguchi method

GA is performed on the aforementioned data to find the selected hub configuration. Several initial parameters are used in the settings. They are indicated with P for population, Cr for

crossover rate, Mr for mutation rate, and G_{max} for the maximum generation of the stopping parameter. The initial parameter settings are 256 for P, 0.7 for Cr, 0.02 for Mr, and 200 for G_{max}. The Taguchi method is utilised to tune the parameters through 16 different combinations. The S/N ratio is then evaluated for the robustness of each combination. Three different solutions are selected based on the highest score, shortest runtime, and the compromise between the two. Taking computational time into consideration, the compromise option is chosen with the parameters of 256 for P, 0.8 for Cr, 0.01 for Mr, and 150 for G_{max}.

Baseline vs optimised network

The hub optimisation problem is solved by the GA across all three LSP companies. For comparison, the baseline value is provided through the number of hubs currently operating in 23 wards of the Tokyo Metropolitan Area. The results from the initial GA parameters are also obtained to provide the optimisation process before the Taguchi method is utilised. Lastly, the results were obtained through the Taguchi-based parameters. The number of hubs under all three conditions can be seen in Table 2.

Table 2: Number of hubs under baseline and GA-optimised networks

Type	Hubs (Baseline)	Hubs (GA before Taguchi)	Hubs (GA after Taguchi)
Sagawa Express	205	82	43
Yamato Transport	272	124	76
Japan Post	1066	442	300
All Companies	1543	648	419

Under the initial parameter settings, the number of hubs is 648, or a 58% reduction compared to the baseline value. Taguchi-based settings gave a further reduction of 419 hubs, or a 72.84% reduction compared to the baseline value. This shows that GA, under the tested ideal parameter settings, can produce high-quality heuristic solutions that produce a significant reduction of hubs. The spread of hubs from all three LSP companies can be seen in Figure 4.

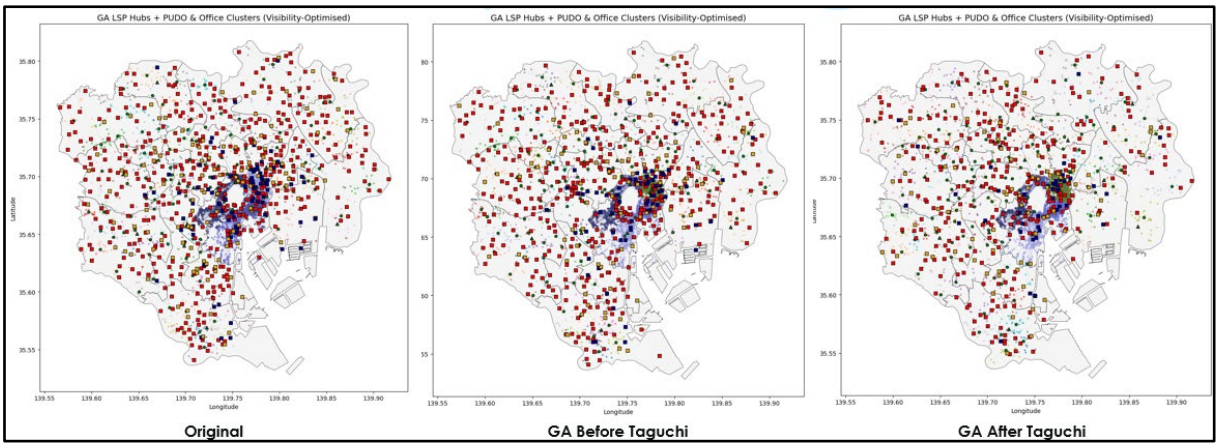


Figure 4: Spread of LSP hubs in the Tokyo Metropolitan Area

Performance improvement

In addition to hub reduction, Taguchi-optimised, GA-based hub configuration also improves operational indicators: average time, distance, and CO_2 emissions. In line with the concept of shared hubs, the values obtained for those aforementioned objectives are based on the joint

operation of assets. The baseline values are obtained from each company's independent operation, using all existing hubs to fulfil all parcel delivery demands in the Tokyo Metropolitan Area. The average time needed for a parcel to travel from origin to destination dropped from the baseline value of 9.4525 hours to 3.0789 hours, or 67.43% reduction, the average distance travelled for a parcel to travel from origin to destination dropped from the baseline value of 23.1729 to 7.2964 kilometres, or 68.51% reduction, and the average emissions generated by a parcel to travel from origin to destination dropped from the baseline value of 1.1164 kg CO_2 per trip to 0.3759 kg CO_2 per trip, or 66.33% reduction.

Sensitivity analysis

GA is a stochastic optimisation method. Therefore, its stability under repeated runs was assessed through the Taguchi method. The method employed an L16 design for the OA matrix to obtain the robust parameter settings for GA, with four main GA parameters as the examination objects: population size, maximum generations, crossover rate, and mutation rate. Each trial in the Taguchi-based tuning was repeated three times with different random seeds, and the results of the objective value and runtime were logged.

The evaluation for each combination's robustness was performed using the S/N ratio, in which the bigger the ratio, the better. Conducting this approach is important as it helps to analyse not only the average performance value but also the consistency of performance under multiple runs. The mean runtime is also recorded in parallel to provide an additional aspect to consider in the trade-off between solution robustness and the computational time required. This can be seen in Figure 5.

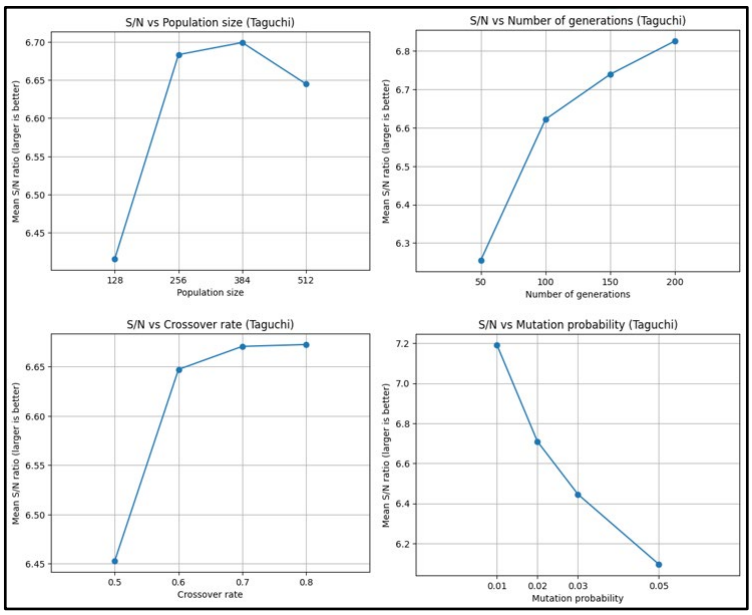


Figure 5: Results of the Taguchi method

5 Conclusions

This study generated a GA-based hub optimisation model for parcel delivery in the Tokyo Metropolitan Area. Real-life logistics facilities are used as shared π-node candidates under a PI-aligned shared hub concept. The findings show that the hubs can be reduced whilst improving average time, distance, and CO_2 emissions generated from parcel delivery activities. The findings also suggest the asset-overlapping phenomenon in the independently operated LSP networks, which can be optimised through shared hub configuration.

The contributions of the study are twofold. First, it demonstrates that a shared hub-based network can be generated based on sets of real-life logistics facilities in a PI-aligned concept. Second, the shared-hub-based network shows that it can improve multiple operational objectives, which are the reduction in hub count, average time, distance, and CO_2 emissions, whilst maintaining the same parcel delivery capability under the independent networks. Such findings can help to pave the way for future PI-related research in Japan, which in turn may tackle its multifaceted problem of labour scarcity, high demand, and labour-protective yet restrictive legislation.

Several limitations need to be acknowledged. First, the optimisation is done under the assumption that all three LSP companies agree to jointly use their assets. Second, detailed operational and governance procedures may need to be further studied in the future. Third, the reliance on assumption-based data due to unavailability needs to be addressed to improve the results further. Lastly, the GA-based results need to be interpreted as a heuristic solution, not a universal guaranteed one.

Future research can be explored by using this study as the basis: expansion of network optimisation for multiple cities, analysis of required governance and operational procedures for a shared parcel delivery network, and integration of full-fledged synchromodality routing in the network. These improvements would enable the concept analysed in this study to be developed further to help with PI-based research and, in the future, the full implementation of PI all around the world.

References

- Azadeh, A. *et al.* (2017) 'A genetic algorithm-Taguchi based approach to inventory routing problem of a single perishable product with transshipment', *Computers & Industrial Engineering*, 104, pp. 124–133. Available at: https://doi.org/10.1016/j.cie.2016.12.019.
- Bayram, V., Aydoğan, Ç. and Kargar, K. (2024) 'Multi-Period Hub Network Design from a Dual Perspective: An Integrated Approach Considering Congestion, Demand Uncertainty, and Service Quality Optimization', *SSRN Electronic Journal* [Preprint]. Available at: https://doi.org/10.2139/ssrn.4752030.
- Fernández-Villaverde, J., Ventura, G. and Yao, W. (2025) 'The wealth of working nations', *European Economic Review*, 173, p. 104962. Available at: https://doi.org/10.1016/j.euroecorev.2025.104962.
- JTA (2024) *Japan Trucking Industry: Current Situation and Problems*. Japan: Japan Trucking Association. Available at: https://jta.or.jp/wp-content/themes/jta_theme/pdf/yusosangyo2024.pdf.
- Kase, H., Murayama, Y. and Suzuki, K. (2023) 'A Study on Improvement of Measures for Redelivery of Courier Services Compared to Parcel Box Typology', Journal of Japan Society of Directories, 21(1), pp. 108–118. Available at: https://doi.org/10.50987/jsod.21.1_108.
- Kweon, O. *et al.* (2024) 'Parcel delivery network optimization problem considering multiple hubs and consolidation of small-sized parcels', *Computers & Industrial Engineering*, 191, p. 110113. Available at: https://doi.org/10.1016/j.cie.2024.110113.
- METI (2025) *E-Commerce Market Survey FY2024*. Ministry of Economy, Trade and Industry, p. 115. Available at: https://www.meti.go.jp/english/press/2025/0826_003.html (Accessed: 2 February 2026).
- MHLW (2023) *Handbook on Work Style Reform-related Laws*. Japan: Ministry of Health, Labour, and Welfare, p. 24. Available at: https://www.mhlw.go.jp/content/001140961.pdf.
- Nakayama, S. and Yan, W. (2019) 'The package redelivery problem, convenience store solution, and the delivery desert: Case study in Aoba Ward, Yokohama', *Journal of Urban Management*, 8(3), pp. 355–363. Available at: https://doi.org/10.1016/j.jum.2019.08.001.
- Sakai, T. *et al.* (2024) 'Locations of logistics facilities for e-commerce: a case of the Tokyo Metropolitan Area', *Research in Transportation Business & Management*, 56, p. 101174. Available at: https://doi.org/10.1016/j.rtbm.2024.101174.
- Sakai, T., Kawamura, K. and Hyodo, T. (2020) 'Logistics facilities for intra and inter-regional shipping: Spatial distributions, location choice factors, and externality', *Journal of Transport Geography*, 86, p. 102783. Available at: https://doi.org/10.1016/j.jtrangeo.2020.102783.
- Vaikkathe, A., Benaini, A. and Boukachour, J. (2025) 'Solving synchromodal container transportation problem using a genetic algorithm', *Multimodal Transportation*, 4(4), p. 100229. Available at: https://doi.org/10.1016/j.multra.2025.100229.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

The strategic value of AI-driven logistics data products

Shulin He[1], Mengdi Zhang[1], and George Q. Huang[1]

1. Department of Industrial and Systems Engineering, The Hong Kong Polytechnic University, Hong Kong, China

Corresponding author: Shulin He

Keywords: *Physical Internet, Cyber-Physical Internet, artificial intelligence, data products, channel structure, pricing, value-added service, game theory*

Conference Domain Fitness: *This study contributes to the IPIC 2026 domain of Hyperconnected supply chain & logistics systems by examining how Cyber-Physical Internet (CPI) services influence data quality and connectivity in the freight market. Our analysis of single versus dual-channel structures for AI-driven data transactions specifically targets Platform business models and marketplaces, offering new insights into the strategic dynamics between data providers and integrators. Additionally, by exploring the monetization of data products and value-added subscriptions, the research extends the concept of Servitization and Logistics-as-a-Service. Finally, the managerial implications derived from our game-theoretic model provide Transition strategies and implementation roadmaps for logistics stakeholders seeking to optimize economic efficiency in the AI era.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:☐ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),☐ Transportation Equipment, ☐ PI Networks, ☐ System of Logistics Networks,☐Vertical Supply Consolidation,☐Horizontal Supply Chain Alignment,☒Logistics/Commercial Data Platform,☐Access and Adoption,☐Governance.*

Targeted Delivery Mode-s*: ☐ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

Research Contribution Abstract (Extended Abstract)

Motivation and Background:

The rapid advancement of artificial intelligence (AI), particularly the deployment of large-scale models, is reshaping how information assets are created and exchanged. Within the freight transportation sector, AI-driven data products that are derived from logistics operations have emerged as a new class of economic goods. Unlike traditional big data supply chains that examine the operational impact of data, the logistics data supply chain (LDSC) treats processed data as a tradable commodity. This shift introduces a novel business environment where raw logistics information is collected, transformed into structured datasets, and eventually packaged as AI-driven data products for end users.

Two types of stakeholders play central roles in this ecosystem. A logistics data provider (LDP) owns original data from daily operations and can directly sell AI-generated insights to the market. A logistics data integrator (LDI), such as FreightWaves Sonar, does not own raw data but consolidates

multiple sources to deliver more comprehensive and accurate analytics. The LDI's offerings are functionally interchangeable with those of a single LDP but are superior in quality due to broader data coverage. This creates a co-opetitive relationship when the LDP supplies its datasets to the LDI for a licensing fee while also maintaining its own direct sales channel.

Key Challenges and Research Questions:

Two critical issues arise in this setting. First, logistics data exhibits strong spatial and temporal correlations, and its granularity and freshness directly affect product effectiveness. Even the most advanced AI cannot compensate for poor raw data quality. The Cyber-Physical Internet (CPI) has been proposed as a novel information technology for logistics, using technologies like Internet of Things, digital twins, and blockchain to enhance real-time data collection and processing. CPI offers value-added services that raise the quality level of AI-driven data products, but at a subscription cost. The impact of such services on the LDP's profit, especially under channel competition, remains unclear.

Second, existing research on virtual product pricing (e.g., data trading, digital goods) largely focuses on single-channel settings. Little is known about pricing AI-driven data products across multiple channels where two different product types coexist: datasets sold by the LDP to the LDI, and AI-driven data products sold by both parties to consumers. Moreover, the economic value of datasets depends on application scenarios and analysis depth, making fixed wholesale pricing unsuitable. A revenue-sharing contract that is responsive to market feedback better captures the long-term relationships in AI training.

Thus, our work addresses three questions:

(1) How should AI-driven data products be priced across different distribution channels given quality considerations?

(2) Should the LDP adopt a dual-channel structure by collaborating with an LDI?

(3) Does subscribing to CPI services improve the LDP's profit, and how does this interact with channel choice?

Methodological Approach:

We develop a game-theoretic model of the LDSC with one LDP and one LDI. Two channel structures are compared: single-channel (LDP sells only directly) and dual-channel (LDP sells directly and also supplies datasets to the LDI for a licensing fee, with both parties then selling substitutable AI products). Within each structure, the LDP may either subscribe to CPI services (paying a fixed fee to enhance quality) or not. Demand functions incorporate cross-channel effects: the competitor's price positively affects own demand, while the competitor's quality negatively affects own demand. We assume that the LDI's quality is exogenously higher than that of the LDP in the base model.

Principal Findings:

First, contrary to conventional wisdom, raising data product quality does not always increase the LDP's profit. In a single channel, a threshold exists beyond which further quality improvement incurs costs that outweigh revenue gains. More strikingly, in a dual channel with a high licensing rate and a market dominated by affluent consumers (who prioritize quality over price), improving the LDP's quality can reduce his own profit. This occurs because higher quality raises his price, driving consumers to the LDI, and also directly reduces the LDI's demand. When the licensing rate is sufficiently large, the LDP's profit aligns with the LDI's, and damaging her sales ultimately harms him.

Second, the dual-channel structure is advantageous only under specific conditions. If the licensing rate is high enough, or if the market consists largely of budget-conscious consumers (who focus on price differences), then adding the indirect channel boosts the LDP's profit. In the former case, a high licensing rate allows the LDP to benefit from the LDI's revenue even while sacrificing some direct sales. In the latter, strong cross-price sensitivity enables the LDP to raise his price and capture more demand. Otherwise, a single channel remains superior.

Third, subscription to CPI services interacts strongly with channel structure. In a single channel, the LDP subscribes only if the quality improvement generates net benefits after covering the fixed fee. In a dual channel, however, spillover effects emerge. When the licensing rate is high and affluent consumers dominate, the LDP's quality upgrade harms the LDI's competitiveness, reducing the licensing revenue he receives. Under such conditions, even a zero-cost CPI subscription may be undesirable. Conversely, when budget-conscious consumers prevail, the LDP's higher quality raises his price, shifting demand to the LDI, a positive spillover that makes subscription profitable.

Managerial Implications and Conclusion:

These results offer practical guidance for the LDSC management. First, quality decisions must be calibrated to market composition and channel strategy; higher quality is not always better. Second, opening a dual channel is not universally beneficial. It requires either a high licensing rate or price-sensitive consumers. Third, value-added services like CPI should be adopted cautiously in dual-channel settings due to negative spillover effects. The study also introduces a revenue-sharing contract for dataset transactions, distinguishing this work from traditional dual-channel literature. Overall, this work provides a strategic framework for stakeholders in the emerging market of AI-driven logistics data products, enabling more efficient channel and subscription decisions under diverse market conditions.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Promoting ESG in energy supply chain with digital twin and private equity: An evolutionary game analysis

Shengan Yu, Mengdi Zhang, Zhiheng Zhao, George Q. Huang
Department of Industrial and Systems Engineering, The Hong Kong Polytechnic University, Hong Kong, China

Corresponding author: melodyd.zhang@polyu.edu.hk

Keywords: *Environmental, Social, and Governance (ESG), energy supply chain, private equity, digital twin, evolutionary game, physical internet*

Conference Domain Fitness: *This study contributes to efficient logistics by advancing the digital technology empowerment of the Physical Internet (PI) for sustainable supply chain transformation. Based on the PI vision of open, interconnected, and interoperable logistics, we explore how digital twins (DT) promote intelligent and ESG decision-making in energy supply chains. This paper adopts a tripartite evolutionary game framework, integrating government policies, private equity investment, and corporate technology adoption to evaluate the feasibility and sustainability of PI digital transformation. This study aligns with IPIC themes on modeling and methodology, feasibility and sustainability assessment, and combines DT with institutional and financial mechanisms to support the development of resilient and sustainable Physical Internet logistics systems.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☐ *PI Networks,* ☐ *System of Logistics Networks,* ☒ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☒ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☒ *Governance.*

Targeted Delivery Mode-s*:* ☐ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

Research Contribution Abstract

Currently, an increasing number of energy enterprises are actively engaging in Environmental, Social and Governance (ESG) practices to facilitate high-quality development through their own initiatives (Lu and Li, 2024). For instance, State Grid ICT released its 2023 ESG report, emphasizing energy digitization and intelligent system development, including the "blockchain+5G" demand response system selected by the National Energy Administration in 2023 (State Grid ICT, 2024). At the industry level, the disclosure rate of ESG-related reports across China's A-share market increased steadily from 30.18% in 2021 to 41.36% in 2023 (National Business Daily, 2024), reflecting the growing importance of ESG practices in data-driven and digitally enabled energy supply chains.

The growing prominence of ESG provides investors with a multidimensional framework to evaluate firm performance beyond purely financial metrics, thereby supporting more informed investment decisions. In this context, private equity organizations (PEs) are expanding their role by facilitating flexible investment strategies and enhancing long-term firm value. Prior studies have shown that ESG considerations are increasingly embedded in supply chain sustainability decisions (Dai and Tang,

2022). At the same time, the development of digital supply chains highlights the importance of advanced technologies in improving operational efficiency and value creation (Büyüközkan and Göçer, 2018). Digitalization is reshaping energy supply chains in ways aligned with the vision of the Physical Internet (PI), which emphasizes interconnected and intelligent logistics systems. Among emerging technologies, digital twin (DT) plays a critical role by enabling real-time data integration, system visibility, and process optimization across energy operations (Zhao et al., 2022). However, despite its potential benefits, the high implementation cost of DT remains a significant barrier, limiting firms' incentives to invest in ESG technological innovation (Zhang et al., 2024). In this process, private equity investment can alleviate financial constraints and support ESG-driven innovation, while governments play a critical role in shaping regulatory and incentive mechanisms. Therefore, the strategic interactions among energy enterprises, PEs, and governments become essential for enabling sustainable and intelligent energy supply chains, motivating further analysis of their long-term decision-making behavior.

Existing studies focuses on single dimension of ESG and supply chain, such as performance evaluation or sustainability outcomes. Energy enterprises are often reluctant to adopt DT technology due to its substantial implementation cost, which limits the technological investment in ESG. In this context, PEs investment plays a crucial role in stimulating enterprise innovation, particularly in promoting substantive technological advancement (Liang et al., 2023). By providing long-term capital support, PEs can facilitate the adoption of DT and accelerate the digitalization of energy supply chains, which is essential for enabling intelligent and interconnected operations aligned with the PI paradigm. Investments are also associated with considerable risks, given the complexity of exit mechanisms and uncertainty in market conditions (Calafiore et al., 2020). However, the existing literature largely overlooks the interactive and dynamic decision-making processes among multiple stakeholders, particularly in the context of digitally enabled and increasingly interconnected supply chains aligned with the PI paradigm. In particular, the strategic interplay among governments, energy enterprises, and PEs in driving ESG transformation through DT adoption remains insufficiently explored.

This paper develops a tripartite evolutionary game model to examine the strategic interactions among energy enterprises, PEs, and governments. Energy enterprises face the trade-off between enhancing ESG performance through DT adoption and the associated implementation costs. PEs, as long-term investors, provide critical financial support that can facilitate firms' ESG transformation and technological upgrading. Governments, as regulatory and supporting entities, aim to promote technological innovation and sustainable development while managing policy costs and incentives. Within the context of digitally enabled and increasingly interconnected energy supply chains aligned with the PI, the model captures the dynamic interplay among these stakeholders and identifies the conditions for optimal strategic choices.

Furthermore, the analysis investigates how key parameters influence equilibrium outcomes and strategic stability, providing insights into ESG transformation and policy design. Accordingly, this study addresses the following research questions: (1) How can strategic interactions among energy enterprises, PEs, and governments be characterized within an evolutionary game framework? (2) What are the equilibrium conditions and evolutionary stable strategies of the three stakeholders? (3) How do key factors influence the strategic choices and equilibrium outcomes of each stakeholder?

To address these issues, this study develops a tripartite evolutionary game model involving PEs, energy enterprises, and governments to characterize their strategic interactions and equilibrium outcomes. The conditions for evolutionary stability are derived based on Lyapunov's first method, followed by a sensitivity analysis to examine the effects of key parameters on stakeholders' strategic choices. The main findings are summarized as follows. First, when environmental penalties are relatively low, PEs are more inclined to invest. However, their strategies remain conservative due to uncertainty and complex exit mechanisms. The government's fines can encourage energy enterprises

to carry out technological innovations in their supply chains and enhance their ESG performance. Second, under high DT implementation costs, governments tend to strengthen support for ESG-related initiatives. While PEs place greater emphasis on long-term returns and show stronger investment incentives when technological barriers are present. Third, social returns from government decisions reflect market dynamics, thereby influencing the strategic decisions of both energy enterprises and PEs in the context of supply chains aligned with the PI.

This paper makes three main contributions. First, this paper constructs a three-party evolutionary game framework to examine the decisions made by the stakeholders, thereby expanding the existing research. By incorporating PEs, this study provides a new perspective on how financial investors influence the ESG decisions of energy enterprises. Second, the study integrates DT technology into the analysis of ESG transformation from an operational management. We emphasize the crucial role of DT as a key driving factor, especially in the context where energy enterprises are confronted with high investment costs and technical barriers, how to improve ESG. This model further reveals the interaction between government policies and the decisions of enterprises and PEs, thereby promoting the improvement of ESG. Third, this paper examines the long-term strategic evolution under uncertainty by considering the cost and risk associated with DT adoption. This study analyzed how the adoption of DT technology affects the dynamic decision-making process of stakeholders, providing important insights for supply chain management in an interconnected PI paradigm. By analyzing the equilibrium conditions and key influencing factors, this study provides insights on management and policies for energy enterprises, PEs and governments, which is conducive to promoting the development of sustainable supply chains.

Reference

Büyüközkan, G. and Göçer, F. (2018), "Digital Supply Chain: Literature review and a proposed framework for future research", *Computers in Industry*, Elsevier B.V., Vol. 97, pp. 157–177, doi: 10.1016/j.compind.2018.02.010.

Calafiore, G.C., Morales, M.H., Tiozzo, V., Fracastoro, G. and Marquie, S. (2020), "Survival and Neural Models for Private Equity Exit Prediction", *IFAC-PapersOnLine*, Vol. 53, Elsevier B.V., pp. 16983–16988, doi: 10.1016/j.ifacol.2020.12.1248.

Dai, T. and Tang, C. (2022), "Frontiers in Service Science: Integrating ESG Measures and Supply Chain Management: Research Opportunities in the Postpandemic Era", *Service Science*, INFORMS Inst.for Operations Res.and the Management Sciences, Vol. 14 No. 1, pp. 1–12, doi: 10.1287/SERV.2021.0295.

Liang, W., Wang, Y. and Hou, L. (2023), "Does private equity investment influence enterprise innovation strategy?", *Finance Research Letters*, Elsevier Ltd, Vol. 58, doi: 10.1016/j.frl.2023.104593.

Lu, J. and Li, H. (2024), "The impact of ESG ratings on low carbon investment: Evidence from renewable energy companies", *Renewable Energy*, Elsevier Ltd, Vol. 223, doi: 10.1016/j.renene.2024.119984.

National Business Daily. (2024), "A-share ESG Disclosure Observations: Scope 1+Scope 2 Disclosure Rate Less Than 20% Experts Suggest Additional Disclosure of Energy Structure", available at: https://www.nbd.com.cn/rss/tonghuashun/articles/3597011.html (accessed 20 April 2026)

State Grid ICT. (2024), "2023 Environmental, Social and Governance (ESG) report", available at: http://www.sgitc.com/u/cms/sgitc/other/202406/2c34025774dc4511afc8c9dc856faa3e.pdf (accessed 20 April 2026)

Zhang, M., Yang, W., Zhao, Z., Wang, S. and Huang, G.Q. (2024), "Do fairness concerns matter for ESG decision-making? Strategic interactions in digital twin-enabled sustainable semiconductor supply chain", *International Journal of Production Economics*, Elsevier B.V., Vol. 276, doi: 10.1016/j.ijpe.2024.109370.

Zhao, Z., Zhang, M., Chen, J., Qu, T. and Huang, G.Q. (2022), "Digital twin-enabled dynamic spatial-temporal knowledge graph for production logistics resource allocation", *Computers and Industrial Engineering*, Elsevier Ltd, Vol. 171, doi: 10.1016/j.cie.2022.108454.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Optimization of Dynamically Hyperconnected Network Services for Heterogeneous Clients

Junkai HE[1], Olivier LABARTHE[1], Zhixin WANG[2], and Ding CHEN[3]

1. *Centre Excellence of Supply Chain Innovation & Transportation, KEDGE Business School, Talence, France*
2. *International School of Low Carbon Studies, Shandong University of Finance and Economics, Jinan, China*
3. *School of Transportation, Fujian University of Technology, Fuzhou, China*

Corresponding author: olivier.labarthe@kedgebs.com

Abstract: *This paper studies the problem of embedding heterogeneous enterprises into a Dynamically Hyperconnected Network (DHN) to access production and logistics resources under time-varying and geographically distributed demand. Demand changes across products, quantities, and locations, affecting network utilization and enterprise decisions. Enterprises must decide whether to join or leave the network and whether to share resources or operate independently. These choices influence facility activation, routing, and overall system performance. A key challenge is coordinating decentralized decisions while maintaining efficient and resilient operations. Uncoordinated decisions may lead to fragmented and inefficient networks, whereas coordinated embedding through resource sharing can improve overall performance. Understanding this trade-off is essential for designing Physical Internet-enabled networks. We propose a dynamic embedding optimization model that integrates enterprises, facilities, transportation routes, and distributed demand within a unified framework. The model minimizes total system cost through joint decisions on facility allocation and hub-based transportation, while ensuring feasibility and adaptability. It captures enterprise embedding, resource sharing, capacity allocation, routing, and network reconfiguration over time, enabling efficient responses to demand changes.*

Keywords: *Physical Internet; Dynamically Hyperconnected Network; Facility Allocation; Geo-Graphically Dispersed Customers; Demand Variation; Mixed-Linear Programming*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),☐ Transportation Equipment, ☒ PI Networks, ☐ System of Logistics Networks, ☐ Vertical Supply Consolidation, ☐ Horizontal Supply Chain Alignment, ☐ Logistics/Commercial Data Platform, ☐ Access and Adoption, ☐ Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

1 Introduction

The transition toward the Physical Internet (PI) has shifted the logistics paradigm from private, siloed supply chains toward open, interconnected ecosystems (Montreuil, 2011). Central to this evolution is the *Dynamically Hyperconnected Network* (DHN), as the core infrastructure. Managed by multiple stakeholders, the DHN consists of production facilities and transportation routes that evolve over time. Its primary role is to provide a shared pool of resources that can

be activated or deactivated based on shifting needs. Unlike traditional networks, the DHN is characterized by its dynamic nature. Facilities and corridors can open to meet demand or support new products and close when utilization is low or companies exit the market. This flexibility allows the network to adapt over time, remaining efficient and responsive.

Within this infrastructure, *heterogeneous enterprises* act as key actors, embedding into the DHN to acquire the resources necessary for production and logistics. These enterprises are diverse in their operational needs, and their interactions are characterized by a tension between collaboration and competition. Collaborative enterprises improve system efficiency by sharing production resources, enabling flexible allocation and reducing costs. In contrast, competitive enterprises may seek exclusive access to specific facilities or routes to protect proprietary processes or ensure service levels. Because these enterprises can choose to join or leave the network and may shift their preferences for sharing over time, their behavior introduces continuous changes that require coordination.

Network activity is driven by *geographically dispersed customers*, whose orders generate demand for products supplied by different enterprises. These customers exhibit varying demand patterns, requiring the enterprises and the DHN to remain adaptive. Demand may be concentrated in a specific area, placing pressure on local DHN nodes. In other cases, demand may be more widely distributed, requiring hub-based transportation. As customers change their requirements in terms of product types and quantities, they affect the entire network. These changes force enterprises to adjust their network design, which in turn requires the DHN to adapt its facilities and transportation routes.

This shift is supported by growing research. Early studies highlighted the need for standardized modular containers and cross-docking to enable resource sharing (Ballot et al., 2014; Montreuil et al., 2013). More recent work shows how hyperconnectivity addresses underutilized proprietary assets (Treiblmaier et al., 2020; Zhang et al., 2026). Classical facility location studies focus on fixed horizons (Jayaswal & Vidyarthi, 2013; Melo et al., 2009), while recent research emphasizes agile and reconfigurable supply chains that can withstand disruptions (Fahimnia et al., 2015; Ivanov et al., 2019) and remain stable despite the entry and exit of enterprises (Govindan et al., 2017). In addition, the success of these networks depends on managing the tension between horizontal collaboration and competitive isolation (Soosay & Hyland, 2015; Pan et al., 2017). However, existing models often treat facility activation and enterprise participation as independent decisions (Schulte et al., 2023; Wang et al., 2024; Liu & Chen, 2025; Kim et al., 2025). There is a need for quantitative frameworks that integrate network dynamics, real-time demand fluctuations, and enterprise behavior in a unified framework. Integrating collaboration and isolation with dynamic facility location and real-time reconfiguration provides a more realistic and actionable approach to managing hyperconnected logistics. Addressing this gap is key to moving the PI from a conceptual vision to practical, scalable implementation.

Several challenges arise. (1) A key challenge is coordinating decentralized decisions while maintaining global efficiency and resilience. As enterprises adjust their network configurations in response to changing demand, outcomes depend on collaboration or competition, leading to global or individual goals. (2) Another challenge is when to open or close facilities as demand evolves. Since these actions involve high costs and lead times, poor decisions can lead to idle capacity or unmet demand. This is limited by the finite production and logistics resources in the DHN, which must be allocated among competing enterprises when demand exceeds supply. (3) Enterprises may collaborate at times but require isolation at others to protect their interests. This increases planning complexity, and handling these shifts together with enterprise entry and exit requires a framework that is both rigorous and flexible.

In this paper, we propose a dynamic embedding optimization model that integrates enterprises, facilities, transportation routes, and geographically distributed demand into a unified mixed-integer linear programming (MILP) model. The objective is to minimize total cost through joint facility allocation and hub-based transportation decisions under demand uncertainty. This paper makes the following contributions:

- We propose a unified framework to embed enterprises into a DHN with time-varying and geographically distributed demand. It captures demand changes, enterprise entry and exit, and resource reconfiguration in a PI-enabled network.
- We develop a MILP model that maps enterprises to the DHN, jointly optimizing facility allocation and hub-based routing.
- We introduce a mechanism to capture the range of enterprise behaviors, from resource isolation to collaboration, to analyze how these preferences affect cost and utilization.
- Numerical results show when collaboration is most beneficial and how the network can support flexible and scalable supply chain operations.

The remainder of this manuscript is organized as follows. Section 2 presents the problem and its mathematical formulation. Section 3 presents the experimental results. Finally, Section 4 concludes the paper and outlines directions for future research.

2 Problem Description and Formulation

First, the topologies and key attributes of the three main actors, DHN, enterprise networks, and customers, are formally stated. Next, the interactions among these actors are described, and the problem under study is precisely defined. Finally, a mathematical model (MILP) is developed to capture these dynamics and support the subsequent analysis.

2.1 Description of the Three Actors

We model the DHN as an undirected, hyperconnected graph divided into non-overlapping geographic zones (denoted by set Z). Each zone contains production facilities and transportation hubs. Facilities provide production capacity, while hubs serve only to connect different zones and do not perform production activities. The total number of nodes in the DHN is $\bigcup_{z\in Z} |F_z| + |H_z|$, where $|F_z|$ and $|H_z|$ represent the number of facilities and hubs in zone z, respectively. We assume that a route exists between adjacent nodes, and when multiple paths are available, the shortest one is selected. Transportation along these routes incurs fixed costs, with additional costs for connections between hubs.

We represent an enterprise network as an undirected graph $G = (N, L, \vec{\vartheta})$, where N and L denote the sets of nodes and links of this enterprise, respectively. The parameter $\vec{\vartheta}$ means the security or exclusivity preferences, indicating whether this enterprise willing to share DHN node resources with other enterprises. Specifically, for each node, $\vartheta = 1$ indicates a willingness to share access to DHN facilities, thereby enabling collaborative utilization of resources. In contrast, $\vartheta = 0$ enforces an isolation constraint, requiring exclusive access to the corresponding DHN node and prohibiting resource sharing with other enterprises.

Customers are geographically dispersed across multiple zones. Each zone exhibits varying levels of demand for different types of products (represented by the set P) in the planning time horizon. For a given product type $p \in P$, its demand d_{zp} may be distributed among different customer groups located in various zones. For example, the demand of product type $p = 1$ may be 60% in zone $z = 1$ and the remaining 40% zone $z = 2$ during a period.

2.2 *Interactions of the Three Actors*

Customer demand is fulfilled by enterprises that rent production and logistics resources from the DHN. Unlike traditional centralized systems, our framework allows enterprises to use distributed resources across multiple zones. Transportation hubs further support efficient logistics and timely delivery across regions. Customer demand evolves over the time horizon in terms of product types, quantities, and geographic distribution. These changes drive the opening and closing of facilities within the DHN and influence enterprise allocation decisions.

Figure 1 illustrates an example allocation solution. The system has three geographical zones with varying customer demands (red tier). The DHN comprises 20 nodes (blue tier) and three transportation hubs (green tier), with all nodes interconnected and each zone's nodes linked to its hub. Two subfigures show the network's service dynamics from period t to $t + 1$. In period t, demand is distributed as follows: product A has 60% demand in zone 1 and 40% in zone 3; product B has 100% demand in zone 2; and product C has 100% demand in zone 1. Demand evolves in period $t + 1$, where product A shifts to 45% in zone 1 and 55% in zone 3, while products B and C remain unchanged. In addition, a new product, product D, appears with 100% demand in zone 2. How to embed enterprise networks on the DHN and adapt the network efficiently to changing and uncertain demand over time are our key concerns.

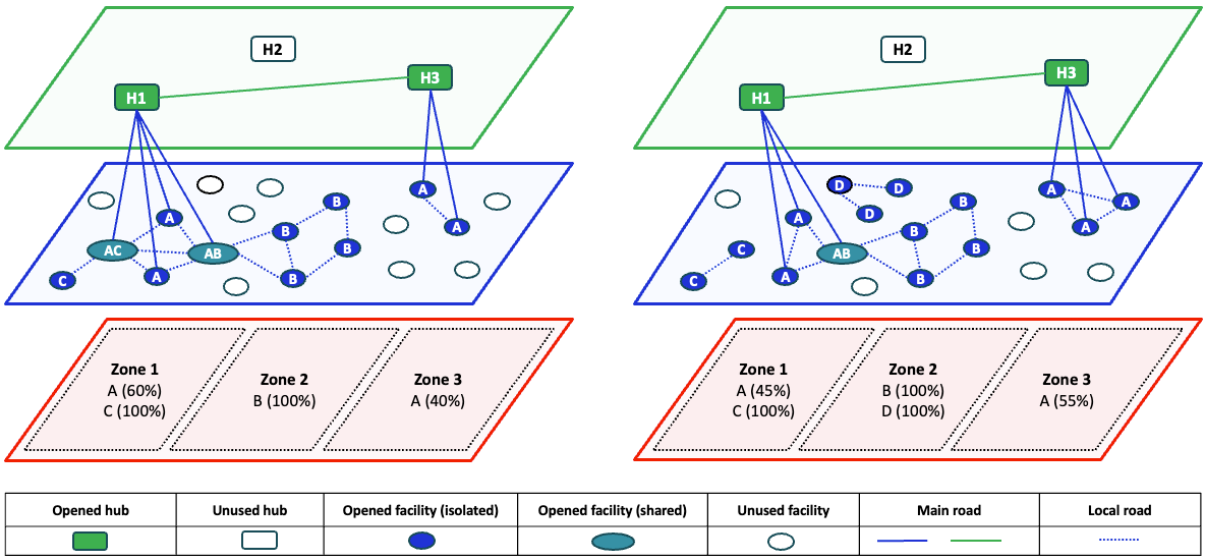


Figure 1. Example of a possible solution from period t (left subfigure) to t + 1 (right subfigure)

Figure 2, building on Figure 1, visualizes this possible solution for enterprise A. It rents four facilities in Zone 1 and two facilities in Zone 3 and uses hubs in both zones to enable product transfers across zones to satisfy customer demand.

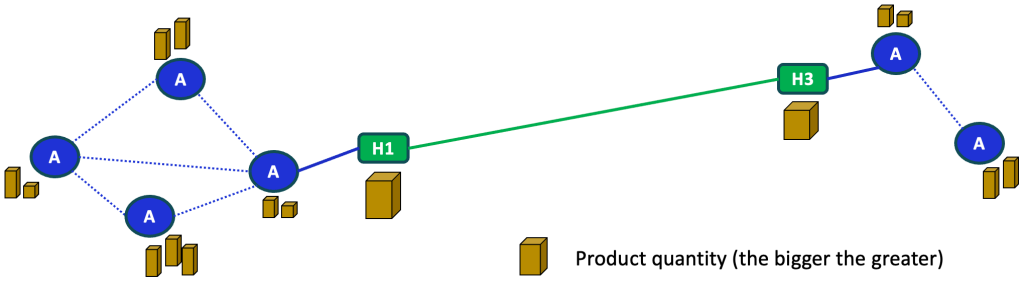


Figure 2. Example of a possible solution in period t (the view of the enterprise for product A)

2.3 Dynamics and Constraints of the Studied Problem

Enterprise networks must decide when and where to request production and transportation resources from the DHN to meet dynamic customer demand. Efficient embedding requires several constraints: each enterprise network node must be assigned to exactly one DHN node, and predefined sharing preferences, based on customer requirements or security concerns, determine node compatibility, represented by a binary parameter. At each allocation round, sufficient facility resources must be available to accommodate all accepted enterprise networks. Networks serving the same product across different zones may use transportation hubs, such as rail-connected main roads, to reduce costs and improve efficiency. The objective is to minimize the total system cost, including hub rental, facility rental, and transportation costs.

This problem is further complicated by dynamic changes: demand evolves over time, enterprise networks may enter or exit the system, and previously embedded networks may be reallocated. Continuous monitoring and updating are necessary to maintain accurate mappings. Key factors influencing network state include: (i) mapping an enterprise network request consumes nodes and links, reducing future availability; (ii) network departures release these resources; (iii) isolation constraints may fragment the network into disconnected subgraphs; and (iv) reallocation can simultaneously free and reassign resources across the network.

2.4 MILP for the Studied Problem

To facilitate reading comprehension, all sets, parameters, and decision variables in this problem are listed in Table 1.

Table 1: Notations

Notation	***Description***
Sets	
$\boldsymbol{Z}$	Set of geographic zones, indexed by $z, z' \in Z$.
$\boldsymbol{H_z}$	Set of hub nodes of zone z, indexed by $h, h' \in H_z$.
$\boldsymbol{F_z}$	Set of facility node of zone z, indexed by $f, f' \in F_z$, where $\bigcup_{z \in Z} F_z = \emptyset$.
$\boldsymbol{P}$	Set of types of products, indexed by $p \in P$.
Parameters	
$\boldsymbol{d_{zp}}$	Demand quantity for product p in zone z.
$\boldsymbol{\vartheta_{pp'}}$	Binary competition parameter, = 0 if the enterprises for product p and p' are competitors, = 1 otherwise.
$\boldsymbol{c_h^{hub}}$	Fixed leasing cost for activating hub h.
$\boldsymbol{c_f^{facility}}$	Fixed leasing cost for activating production facility f.
$\boldsymbol{c_{zz'}^{transport}}$	Unit transportation cost between physical nodes z and z'.
$\boldsymbol{m_{hz}^{hub}}$	Maximum transportation throughput capacity of hub h in zone z.
$\boldsymbol{m_{fz}^{facility}}$	Maximum production capacity of facility f in zone z.
Decision variables	
$\boldsymbol{X_{hz}}$	Binary variable, = 1 if hub h in zone z is leased, = 0 otherwise.
$\boldsymbol{Y_{fz}}$	Binary variable, = 1 if facility f in zone z is leased, = 0 otherwise.
$\boldsymbol{K_{fp}}$	Binary variable, = 1 if facility f fabricates product p, = 0 otherwise.
$\boldsymbol{Q_{fp}}$	Production quantity of product p at facility f.
$\boldsymbol{R_{zz'p}}$	Production quantity of product p transported from zone z to zone z'.

We establish a MILP model (denoted as M1) for the problem under study. The objective function is to minimize the total system cost, composed of three primary elements: the costs required to lease and activate transportation hubs, the costs for leasing production facilities, and the costs associated with transporting products between different geographic zones.

$$\textbf{M1}: \quad Min\ f = \sum_{z\in Z}\sum_{h\in H_z} c_h^{hub}\cdot X_{hz} + \sum_{z\in Z}\sum_{f\in F_z} c_f^{facility}\cdot Y_{fz} + \sum_{p\in P}\sum_{z\in Z}\sum_{z'\in Z, z\neq z'} c_{zz'}^{transport}\cdot R_{zz'f}$$

Subject to:

$$\sum_{p\in P}\sum_{z'\in Z, z\neq z'}(R_{zz'f} + R_{z'zf}) \leq \sum_{h\in H_z} m_{hz}^{hub}\cdot X_{hz}, \qquad \forall z\in Z \tag{1}$$

$$\sum_{p\in P} Q_{fp} \leq m_{fz}^{facility}\cdot Y_{fz}, \qquad \forall z\in Z, \forall f\in F_z \tag{2}$$

Constraint (1) ensures that the total amount of products moving through a zone's hubs does not exceed the maximum limit of the rented hubs. Similarly, Constraint (2) ensures that a facility does not produce more than its maximum manufacturing limit.

$$\sum_{f\in F_z} Q_{fp} + \sum_{z'\in Z, z\neq z'} R_{z'zf} - \sum_{z'\in Z, z\neq z'} R_{zz'f} = d_{zp}, \qquad \forall z\in Z, \forall p\in P \tag{3}$$

Constraint (3) balances the products in each zone. It states that the amount produced locally plus the amount brought in from other zones, minus the amount sent out to other zones, must exactly equal the local customer demand. This ensures no shortages and no wasted products.

$$K_{fp} + K_{fp'} \leq 1 + \vartheta_{pp'}, \qquad \forall z\in Z, \forall f\in F_z, \forall p, p'\in P, p\neq p' \tag{4}$$

$$Q_{fp} \leq m_{fz}^{facility}\cdot K_{fp}, \qquad \forall z\in Z, \forall f\in F_z, \forall p\in P \tag{5}$$

$$K_{fp} \leq Y_{fz}, \qquad \forall z\in Z, \forall f\in F_z, \forall p\in P \tag{6}$$

Constraint (4) ensures that competing products are not made in the same facility. Constraint (5) ensures that production only happens if a product is officially assigned to a facility. Constraint (6) prevents any product from being assigned to a facility that has not been rented.

$$X_{hz}, Y_{fz}, K_{fp} \in \{0,1\}, \qquad \forall z\in Z, \forall h\in H_z, \forall f\in F_z, \forall p\in P \tag{7}$$

$$Q_{fp}, R_{zz'f} \geq 0, \qquad \forall z, z'\in Z, z\neq z', \forall f\in F_z, \forall p\in P \tag{8}$$

Finally, Constraints (7) and (8) define the ranges of decision variables.

3 Numerical Experiments

This section presents the experimental design and results used to validate the proposed model. First, the parameter settings are described. Then, a case study is conducted to demonstrate the effectiveness of the proposed framework and to provide managerial insights. All models are solved using GUROBI 12.03, and all experiments are conducted on a personal laptop equipped with a 13th Gen Intel (R) Core (TM) i5-13500H (2.60 GHz) and 32 GB of RAM.

The case study includes three geographic zones (i.e., $|Z| = 3$). Each zone contains 20 nodes, resulting in 60 nodes in the DHN, specifically, $\bigcup_{z\in Z} |F_z| + |H_z| = 60$, where hubs and facilities contain 20% and 80%, respectively. This infrastructure serves geographically dispersed customer demand for four heterogeneous products (i.e., $|P| = 4$), each associated with a distinct enterprise. Other related inputs are derived by open-source industry data[1][2].

[1] https://m.50yc.com/
[2] https://www.zztender.com/announce/info?announceId=060b4732-46de-48c1-a542-507f5eb2af01

3.1 Global and Local Optimization

We compare two strategies reflecting collaborative versus competitive preferences. The first is a globally coordinated optimization, representing the proposed model. The second is a local, independent approach, where enterprises strictly or partially isolate resources, optimizing facility configurations and routing decisions independently without sharing nodes. Global optimization model (i.e., M1) is presented in Section 2, while local optimization model (denoted by M2) for a product p is as follows:

$$\mathbf{M2}: \quad minf_p = \sum_{z\in Z}\sum_{h\in H_z} c_h^{hub}\cdot X_{hz}^p + \sum_{z\in Z}\sum_{f\in F_z} c_f^{facility}\cdot Y_{fz}^p + \sum_{z\in Z}\sum_{z'\in Z, z\neq z'} c_{zz'}^{transport}\cdot R_{zz'p}$$

s. t., (1) - (8),

$$\vartheta_{pp'} = 0, \qquad \forall p, p' \in P, p \neq p' \tag{9}$$

where the binary decision variables for all preceding products in the sequence $\{1, \ldots, p-1\}$ are treated as fixed parameters. We assume a sequential decision-making priority where product 1 is optimized first, followed chronologically by products 2, 3, and 4, ensuring that any resources occupied by earlier products are locked and unavailable for subsequent enterprises.

The simulation results in Figure 3 show a clear performance gap between global and local optimization strategies within the DHN. As shown in Figure 3(A), the global approach achieves a total cost of 1.12e^6 CNY, about 32% lower than the 1.65e^6 CNY under local optimization. This difference is mainly driven by facility costs. Under local optimization, enterprises prioritize individual service and resource isolation, leading to redundant facility activation and higher costs. In contrast, the global approach exploits resource sharing to meet demand with a more efficient network structure.

Figure 3(B) further explains these differences by comparing the number of leased resources and capacity utilization. Under global optimization, the network operates with only 2 hubs and 7 facilities, yet achieves high efficiency, with facility utilization reaching 89.29%. In contrast, local optimization leads to a fragmented network with 5 hubs and 10 facilities, resulting in much lower utilization. Hub utilization drops from 35.71% to 14.04%, and facility utilization decreases to 63.41%.

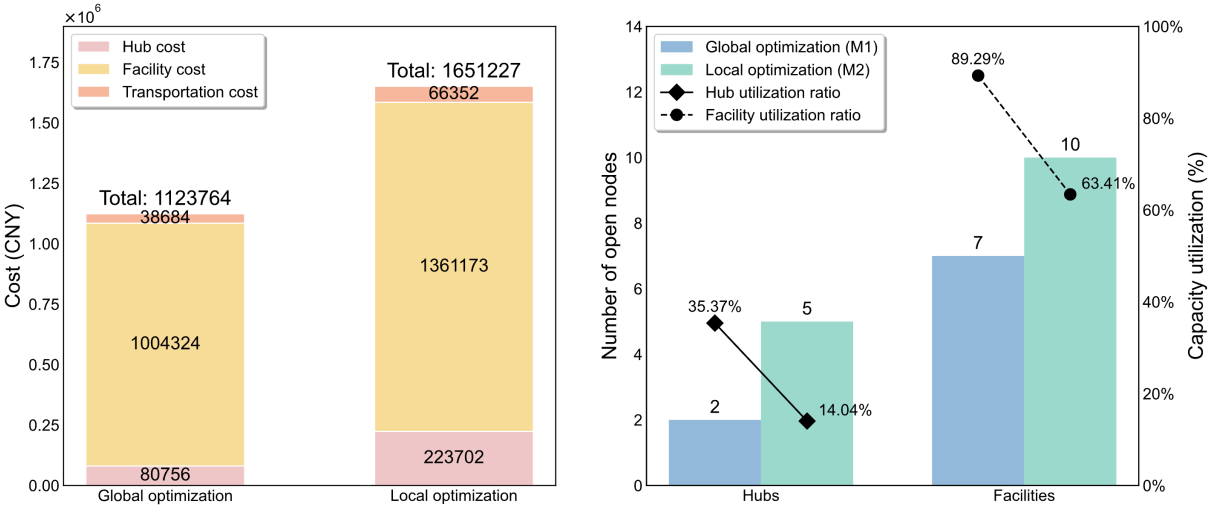


Figure 3: Cost and utilization comparison between global and local optimization strategies

3.2 Impact of Demand Violation

To evaluate the adaptability of the proposed DHN for *the next period*, we conduct further analysis across four distinct demand scenarios (S1-S4). S1 serves as the benchmark (baseline). Other scenarios represent varying degrees of demand shifts: S2 (zone 1 increases 60% but zone 3 decreases 60%), S3 (zone 2 increases 60% but zone 1 decreases 60%), and S4 (zone 2 decreases to zero, zone 3 increases to double). The results, comparing S1 with S2-S4, are shown in Figure 4. Notation (a/b) tracks the re-organization of nodes (facilities and transportation hubs), where $\pm a$ denotes the number of *newly* opened/closed nodes, while $\pm a$ denotes the number of opened/closed nodes *associated the previous period*. The ratios mean the usage percentage of facility or transportation capacity at associated nodes.

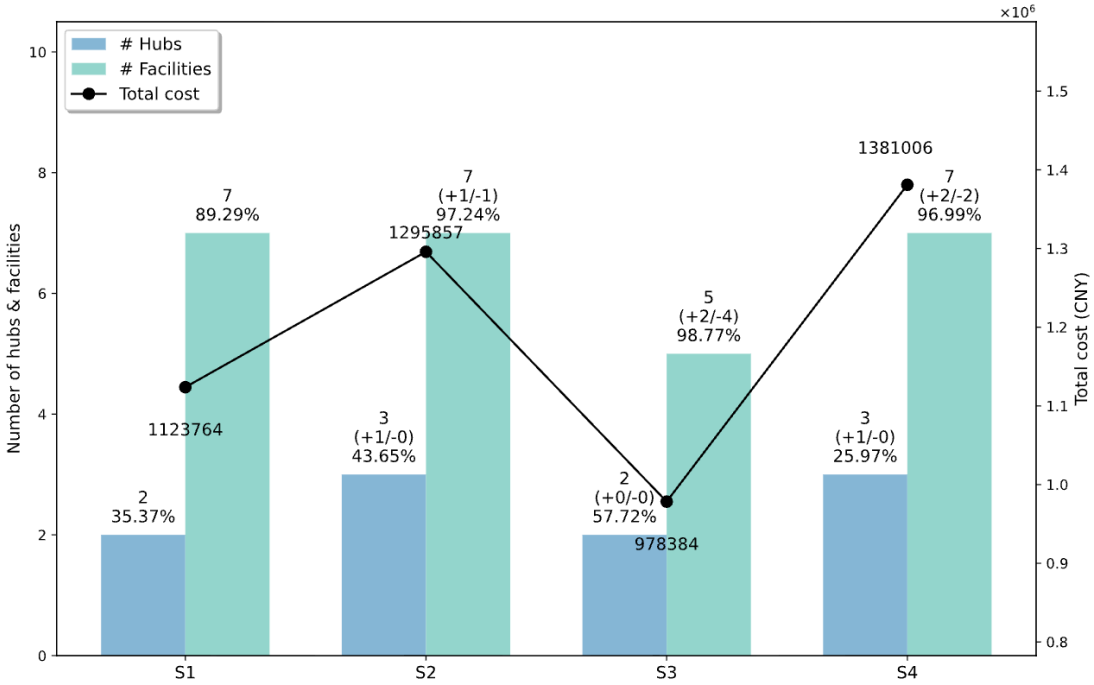


Figure 4: Impact of demand violation on different costs

In S2, the network responds to the demand shift by activating an additional hub (+1), bringing the total cost to 1.29e^6 CNY. Interestingly, in S3, the network achieves its highest efficiency; despite the 60% increase in zone 1, the system consolidates operations by deactivating four facilities while only opening two new ones (+2/-4). This aggressive reconfiguration reduces the total facility count to 5 and brings the total cost down to a 0.98e^6 CNY, validating the DHN's ability to downscale to eliminate idle capacity.

However, the extreme volatility of S4 tests the limits of the network's resilience. To support the 200% demand increasement in zone 3 while compensating for the exit of zone 2, the network maintains a high facility count of 7 and increases its hub capacity to 3 (+1/-0). This necessity for rapid scaling and the high setup costs associated with such a drastic shift result in the highest total expenditure of 1.38e^6 CNY. In summary, we confirm that the model can balance the "sharing-isolation tension" by dynamically adjusting facilities and transportation routes, ensuring the DHN remains service-feasible even faced with demand ripples.

3.3 Impact of Transportation Cost

We further conduct a sensitivity analysis on the transportation cost multiplier α, comparing three strategies: Global optimization (M1), Global optimization (M1 with $\vartheta = 0$), and Local optimization (M2). The results shown in the total cost curves are presented in Figure 5.

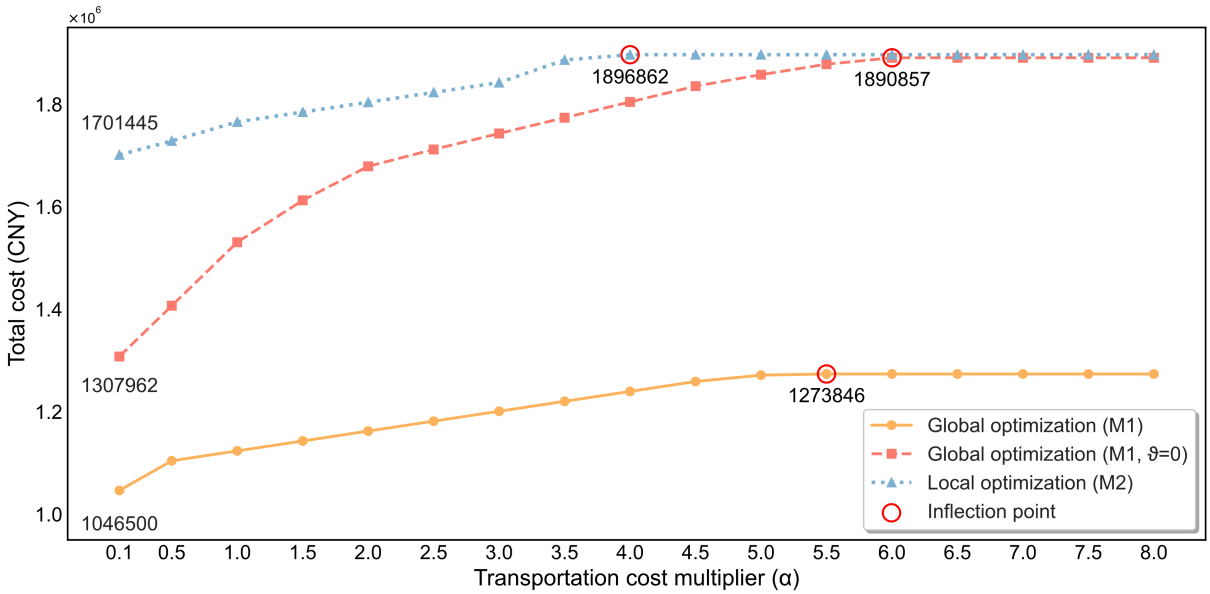


Figure 5: Impact of transportation cost and sharing preferences on total cost

Across all α values, Global optimization (M1) consistently achieves the lowest total cost, starting at 1.05e^6 CNY at $\alpha = 0.1$. In contrast, Global optimization (M1 with $\vartheta = 0$) incurs higher costs due to redundant production lines, with the disparity most pronounced at lower α but persisting as transportation costs rise. Local optimization (M2) remains the most expensive, highlighting that decentralized, non-sharing decisions impose substantial economic burdens.

All three curves exhibit that costs rise steeply with increasing α before plateauing. For Global optimization (M1), the inflection occurs at $\alpha = 5.5$ (1.27e^6 CNY), while Global optimization (M1 with $\vartheta = 0$) and Local optimization (M2) reach their ceilings earlier or at higher costs. These plateaus indicate the network's structural limits, where facilities and routes are fully reconfigured to mitigate rising transportation costs.

These findings underscore that Global optimization (M1 with $\vartheta = 0$) and Local optimization (M2) are incompatible with hyperconnected network efficiency. Resource sharing is not merely beneficial, it is essential. Without it, even globally optimized systems converge toward the high-cost profile of fragmented, local networks. This highlights the DHN's critical role in managing the balance between sharing and isolation to maintain cost-effective, service-feasible operations.

4 Conclusion and Future Research

This paper addresses embedding heterogeneous enterprises into a DHN to satisfy evolving, geographically dispersed demand. We model DHN, enterprise, and customer interactions within a unified MILP model. Results show that competitive, resource-isolated strategies can lead to underutilized capacity, higher costs, and reduced responsiveness. Conversely, collaborative approaches enhance resilience, cost efficiency, and adaptability to variable demand. Besides, our model can treat evolving demand and propose reasonable re-organization solutions. The study underscores the operational and strategic benefits of coordinated decision-making and examines the behavioral dimension of enterprise participation in the DHN.

Regarding practical insights for managers, first, leveraging shared resources and coordinated planning can significantly reduce costs and improve service reliability across the network. Second, anticipating changes in demand and proactively reconfiguring facilities and routes helps prevent capacity bottlenecks and operational disruptions. Third, understanding the

dynamic nature of enterprise relationships, balancing collaboration and competitive protection, is crucial for sustaining operational performance.

Several directions remain for future research. First, we plan to extend the framework to include stochastic demand modeling, helping the DHN better handle probabilistic forecasts and extreme market changes. Second, the behavioral side of enterprise participation could be explored using game-theoretic methods to model strategic negotiation when sharing capacity under limited resources. Finally, future work will examine integrating product type, proportion fluctuation, which is essential to develop a sustainable and emerging PI.

References

- Ballot E., F. Fontane (2008): Rendement et efficience du transport: un nouvel indicateur de performance. Revue Française de Gestion Industrielle, v27, no2, 41-55.
- Ballot E., B. Montreuil, C. Thivierge (2014): Functional design of physical internet facilities: A simulation-based approach. Progress in Material Handling Research, v13, 1-15.
- Fahimnia B., T.G. Reisi, H.S. Mamani (2015): Competitive supply chain network design: An overview of models and algorithms. Transportation Research Part E: Logistics and Transportation Review, v78, 12-32.
- Govindan K., M. Fattahi, E. Keyvanshokooh (2017): Supply chain network design under uncertainty: A comprehensive review and future research directions. European Journal of Operational Research, v259, no2, 524-541.
- Ivanov D., A. Dolgui, B. Sokolov (2019): The impact of digital technology and Industry 4.0 on the ripple effect and supply chain risk analytics for resilience. International Journal of Production Research, v57, no3, 829-846.
- Jayaswal S., N. Vidyarthi (2013): Integrated production-distribution planning with facility location and inventory considerations. International Journal of Production Economics, v144, no1, 22-34.
- Kim S., J. Lee, H. Park (2025): Adaptive facility management in decentralized hyperconnected logistics. Journal of Intelligent Manufacturing, v36, no4, 112-128.
- Liu Y., X. Chen (2025): A unified framework for real-time network reconfiguration in the Physical Internet. Transportation Research Part C: Emerging Technologies, v160, 104-119.
- Melo M.T., S. Nickel, F. Saldanha-da-Gama (2009): Facility location and supply chain management – A review. European Journal of Operational Research, v196, no2, 401-412.
- Montreuil B. (2011): Toward a Physical Internet. Service Orientation in Holonic and Multi-Agent Manufacturing and Robotics, v335, 71-87.
- Montreuil B., R. Meller, E. Ballot (2013): Physical Internet Foundations. Progress in Material Handling Research, v14, 25-40.
- Pan S., M. Nigrelli, E. Ballot, R. Sarraj, Y. Yang (2017): Perspectives of inventory and transport in Physical Internet. Computers & Industrial Engineering, v111, 481-490.
- Schulte J., M. Weber, D. Schmidt (2023): Stochastic modeling of enterprise participation in open logistics ecosystems. International Journal of Logistics Management, v34, no1, 88-105.
- Soosay C., P. Hyland (2015): A decade of supply chain collaboration and networking: A review. International Journal of Physical Distribution & Logistics Management, v45, no1/2, 158-177.
- Treiblmaier H., G. Rejeb, A. Rejeb (2020): The role of blockchain in the Physical Internet. Journal of Enterprise Information Management, v33, no6, 1225-1245.
- Wang L., Q. Zhang, R. Wu (2024): Quantitative models for dynamic facility location under fluctuating demand. European Journal of Operational Research, v312, no2, 550-565.
- Zhang X., F. Tao, A. Nee (2026): Hyperconnected production networks: Real-time asset sharing and optimization. Journal of Manufacturing Systems, v75, 210-225.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

PIMM: A Maturity Model for Advancing the Physical Internet from Concept to Practice

Tadashi Mizutani[1], Yuki Kurakake[1], Kei Yazaki[1], Maki Aoki[1], Tomohiro Okuzumi[2] and Satoshi Kagino[2]

1. Nomura Research Institute, Tokyo, Japan
2. Japan Physical Internet Center, Tokyo, Japan

Corresponding author: t-mizutani@nri.co.jp

Abstract: *Severe labor shortages in the logistics industry have made the transition to collaborative ecosystems via the Physical Internet (PI) essential. However, the absence of objective standards to measure the maturity of these initiatives hinders reliable partner matchmaking. To bridge this gap, this study develops the Physical Internet Maturity Model (PIMM) by synthesizing three established frameworks: CMMI, the Global SCM Scorecard, and the Integrative Framework for Collaborative Governance. PIMM translates the abstract PI concept into a concrete evaluation tool comprising 17 items across five dimensions. Integrating multi-stakeholder feedback, the model uniquely expands beyond technical aspects to address fundamental organizational challenges like human resource development. Furthermore, to ensure practical viability, an operational framework based on the 4W1H methodology is established. Featuring a hybrid assessment that combines document reviews with interviews, this approach resolves the implementation trilemma of reliability, system adoption, and sustainability. It mitigates subjective bias and accurately captures full operational utilization rather than mere system implementation. An initial trial with five companies demonstrated PIMM's practical efficacy, with 80% expecting it to facilitate future partner matching. Ultimately, PIMM provides a scalable, actionable foundation for globally interconnected logistics, paving the way for universal PI adoption.*

Keywords: *Physical Internet, maturity model, assessment tool, industry adoption*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☐ *PI Networks,* ☐ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☒ *Access and Adoption,* ☒ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

1 Motivation

1.1 Background

Japan is currently facing a critical logistics crisis driven by a severe shortage of truck drivers, a declining working population, and recent labor law reforms. Consequently, logistics—traditionally viewed as a highly competitive domain in Japan—is increasingly recognized as a cooperative sector. Initiatives aimed at realizing the Physical Internet (PI) are now transitioning from the research and demonstration phase to real-world application. Following the formulation of the Japanese government's PI Roadmap in 2022, numerous industry-specific working groups have been established, thereby accelerating collaborative logistics efforts. However, a significant challenge remains: the absence of a standardized metric to measure and compare the

maturity or progress of these diverse initiatives. Without such a benchmark, companies seeking to engage in collaborative logistics struggle to identify relevant best practices. To overcome this barrier and effectively guide these cooperative efforts, an objective evaluation framework is essential.

1.2 Objective

To address this gap, this study develops a pioneering Physical Internet Maturity Model (PIMM). The development of PIMM was assigned to a working group formed within the Physical Internet Realization Consortium, established by the Japan Physical Internet Center (JPIC), with its activities commencing in July 2025. The authors of this research actively participated as core members of the working group.

The primary objective of this model is to enhance the efficiency and resilience of logistics, ultimately contributing to sustainable supply chains. PIMM conceptualizes the developmental stages of PI implementation across five distinct maturity levels. While this model originates from Japanese initiatives, Japan's extreme labor shortage serves as a leading indicator of critical challenges that other regions may soon face. Consequently, the need for a common language and objective evaluation metrics is a universal issue. Therefore, this study aims to serve as a scalable, foundational step toward the global applicability of PIMM, promoting the worldwide adoption and implementation of the PI.

2 Development Context and Scope of This Paper

The “Physical Internet Roadmap,” established by the Japanese government in 2022, outlines the long-term strategic direction for logistics and categorizes its developmental phases. It designates the period up to fiscal year (FY) 2030 as the “Takeoff Phase,” focusing on intra-industry and intra-regional collaborative logistics (see Figure 1). The subsequent period, from FY 2031 to FY 2035, is defined as the “Acceleration Phase,” which targets broader cross-industry and cross-regional collaboration. Consequently, the end of FY 2030 serves as a critical milestone, marking the transition from closed, industry-specific initiatives to open, cross-industry networks. To align the practical deployment of PIMM with this 2030 milestone, its development and deployment are structured into three strategic phases:

Phase 1: Focuses on establishing the framework and operational design of the maturity model.

Phase 2: Aims to achieve broad industry recognition and integration of PIMM by FY 2028.

Phase 3: Targets comprehensive dissemination across the industrial sector by FY 2030, aligning with the conclusion of the Takeoff Phase. Successfully executing these domestic phases will establish a proven, scalable foundation for PIMM’s subsequent global adaptation.

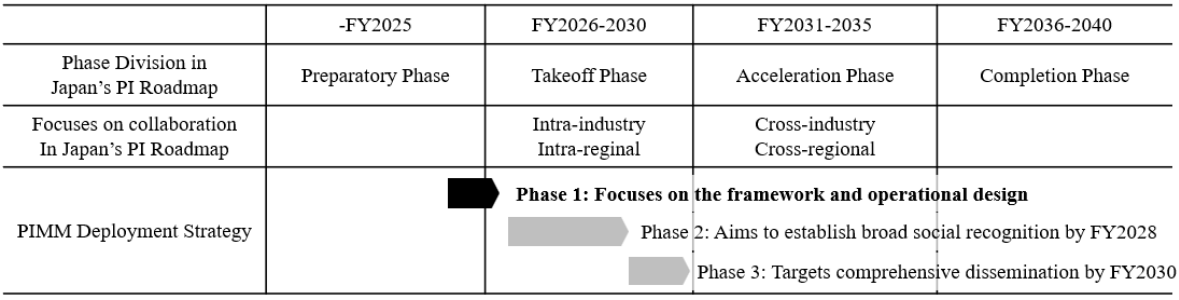


Figure 1: Alignment between Japan's PI Roadmap and PIMM Deployment Strategy

The scope of this paper is limited to the outcomes of Phase 1. Specifically, it details the framework and operational designs intended for PIMM to serve as a reliable and practical compass for enterprises. Subsequent activities, including broader dissemination strategies and the collection of large-scale empirical data (Phases 2 and 3), are beyond the scope of this study and remain areas for future research. Based on these foundational designs, PIMM Version 1 is scheduled to be officially released in the first half of 2026.

3 Framework Design

3.1 Framework Design Approach

Rather than developing PIMM from scratch, this study first analyzed prominent existing frameworks for supply chain maturity. For instance, widely referenced models such as the Gartner's Demand-Driven Value Networks (DDVN) maturity model and the ASCM's Supply Chain Operations Reference (SCOR) model provide robust methods for assessing and improving processes within a company's own supply chain or its established network of partners. However, these models were not designed to evaluate the key tenets of the Physical Internet: the open and dynamic sharing of physical assets (like containers and warehouse space) and logistics services among non-contracted, cross-industry stakeholders. They fall short in measuring the readiness for an open, collaborative ecosystem.

To bridge this critical gap, this study synthesizes and restructures three established, well-validated frameworks from different fields (see Table 1). The first foundational model is the Capability Maturity Model Integration (CMMI), widely recognized for software development maturity. However, because CMMI does not cover supply chain and logistics operations, the second framework incorporated is the "Global SCM Scorecard" developed by JILS (a Japanese industry association) and the Tokyo Institute of Technology in Japan. While this scorecard provides an excellent method for benchmarking internal supply chains, it is insufficient for evaluating collaborative logistics among non-contracted entities. Therefore, the third reference incorporated is the "Integrated Framework for Collaborative Governance," which evaluates multi-stakeholder coordination. PIMM was constructed by harmonizing the strengths of these three frameworks. To ensure long-term strategic consistency, the PI roadmaps developed by ALICE and the Japanese government were also integrated.

Table 1: Framework Design Approach: Analysis of Existing Models

Existing framework	Key strength	Structural blind spot
CMMI（Capability Maturity Model Integration）	◯ Structured maturity levels	× Ignores physical logistics operations
Global SCM Scorecard	◯ Evaluation of internal supply chain & logistics operations	× Fails to account for 3rd-party non-contracted entities
Integrated Framework for Collaborative Governance	◯ Multi-stakeholder problem-solving	× Lacks specific operational mechanics

3.2 Definition and Logic of Evaluation Dimensions

3.2.1 Logic Constituting the Core Evaluation Dimensions

PIMM comprises five interrelated evaluation dimensions: physical space, business processes, digitalization, ecosystem, and human resource development (see Figure 2). These dimensions were derived from the synthesis of the aforementioned frameworks and subsequently established through multi-stakeholder co-creation (as detailed in Section 3.5).

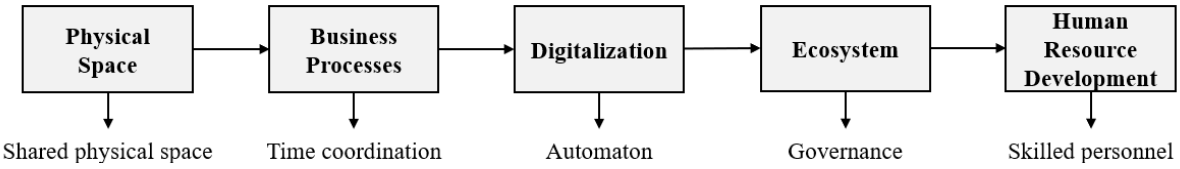


Figure 2: The Chained Structure of PIMM's 5 Evaluation Dimensions

These dimensions follow a sequential logic. First, multiple stakeholders must share a limited **physical space**. To effectively utilize this shared space, the second dimension, **business processes**, is required to synchronize the timing of space usage among stakeholders. Third, as cargo volume and complexity increase, manual coordination reaches its limits, necessitating **digitalization** for large-scale, automated synchronization. Fourth, an **ecosystem** with robust rules and governance is essential to transform short-term pilot projects into sustainable, long-term relationships. Finally, **human resource development** is crucial for sustaining and expanding multi-company collaboration.

This chained structure, particularly the inclusion of 'Physical Space' and 'Ecosystem' as core dimensions, is what makes PIMM uniquely tailored for the Physical Internet. Unlike traditional collaboration models, which prioritize information and process flows, PIMM fundamentally evaluates the standardization and sharing of physical assets and the governance for an open, multi-access network—both of which are central principles of the PI.

3.2.2 Granular Breakdown of the Evaluation Items

These five dimensions are further subdivided into 17 specific evaluation items (see Figure 3).

- **Physical Space**: Divided into the demand side ("containers") and the supply side ("equipment"). Truck beds are excluded from the supply side because they accommodate various containers, and the PI fundamentally prioritizes standardized modular containers over the generic carriers that transport them. Both are evaluated across three criteria: *standardization*, *openness*, and *sharing*, yielding six items in total.
- **Business Processes**: Expanded to reflect the logic of space and divided into four items: *standardization*, *openness*, *horizontal collaboration*, and *vertical integration*.
- **Digitalization**: Evaluates technological maturity as an essential enabler for operational synchronization, covering the following three items: *containers*, *equipment*, and *business processes*.
- **Ecosystem**: Comprises three items: *organization*, *governance* (including profit and responsibility distribution), and t*arget setting and evaluation*.
- **Human Resource Development**: Maintained as a single evaluation item.

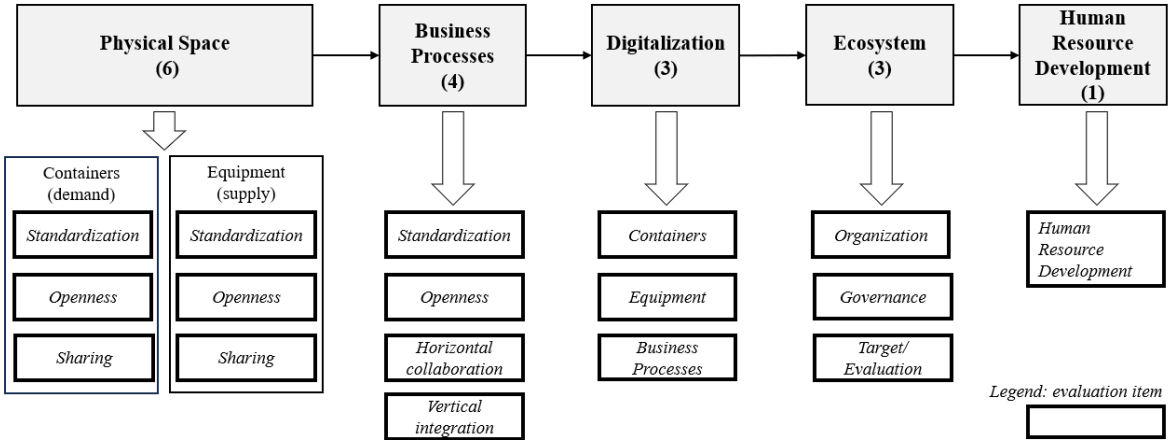


Figure 3: Granular Breakdown of PIMM's 17 Evaluation Items

3.3 Nuanced Level Definitions

3.3.1 Policy for Level Definitions in PIMM

To avoid the ambiguity that arises from applying uniform qualitative definitions to diverse tangible and intangible elements, PIMM establishes specific, milestone-based criteria for each evaluation item. Here are two examples to illustrate this. The first example shown in Table 2 concerns a tangible element, such as the standardization of containers: the criteria progress from having no unified rules (Level 1) to utilizing industry common standards that consider the combined efficiency of multiple container types (Level 5).

Table 2: Level definition examples for the "Standardization" under "Containers" within the "Physical Space" dimension

	Level 1	Level 2	Level 3	Level 4	Level 5
Qualitative State	Unestablished	Unified across internal warehouses/businesses	Unified with specific external companies	Standard specifications (single type)	Standard specifications (multiple types)
Description	No unified rules for container sizes and specifications	Adopting unified standards across multiple internal warehouses and business units	Unifying container standards with specific trading partners	Using containers compliant with common industry standards. Standardization is limited to a single type (e.g., folding containers only, or pallets only)	Using containers compliant with common industry standards. The efficiency of combining multiple types of containers, such as folding containers and pallets, is taken into consideration.

The second example shown in Table 3 involves an intangible element: digitalization progresses explicitly from Manual (Level 1) to Digitization (Level 2), System Integration (Level 3), Visualization and Analysis (Level 4), and AI-driven Decision Support (Level 5).

Table 3: Level definition examples for the "Containers" within the "Digitalization" dimension

	Level 1	Level 2	Level 3	Level 4	Level 5
Qualitative State	Manual	Digitization	System Integration	Visualization, Analysis & Improvement	AI-driven Decision Support
Description	Paper-based or manual management / Unmanaged	ID assignment per shipping unit(e.g., pallets, roll cages). Standardized identifiers like GS1 SSCC are desirable.	Container management system linked with ERP/WMS, enabling the tracking of the exact quantity, location, and specifications of containers	Dashboard visualization and analysis of utilization rates, turnover rates, and retention status, enabling human identification of areas for improvement	AI analyzes real-time and historical data to automatically detect and alert inefficiencies or problems in the supply chain process

3.3.2 Definitions of Minimum and Maximum Levels

Level 1 is not merely an "unmanaged" state; it explicitly requires awareness and a willingness to initiate PI implementation, distinguishing it from Level 0 (unaware). Conversely, Level 5 is aligned with the ideal future vision outlined in Japan's PI Roadmap for FY 2030, serving as a highly concrete and ambitious benchmark for the ultimate state of maturity.

While PIMM Version 1 sets FY 2030 as its target year for these level definitions, future updated versions may redefine the target year to FY 2035, aligning with the subsequent "Acceleration Phase."

3.4 Score Calculation Method

The calculated maturity score quantitatively reflects a company's progress toward PI integration. Rather than employing a simple average, specific weightings—established through the multi-stakeholder co-creation process (see Section 3.5)—are applied to the 17 evaluation items to reflect their proportional importance within the broader PI ecosystem. The final comprehensive maturity level is rounded to an integer (ranging from Levels 1 to 5). This deliberate rounding ensures that companies aim for visible, qualitative phase shifts (a full level increase) rather than marginal numerical improvements.

3.5 Refinement through Co-creation

PIMM was iteratively refined through multi-stakeholder feedback, incorporating approximately 70 constructive insights from external experts, academia, and industry practitioners. This co-creation process highlighted that PI cannot be achieved solely through process synchronization and technological implementation. Consequently, "organization" and "human resource development" were formally integrated into the model. This addition emphasizes that a robust internal foundation—spanning logistics, manufacturing, sales, and management—is crucial, thereby elevating PIMM from a specialized logistics tool to a holistic organizational framework for collaborative logistics.

4 Operational Design

4.1 The Implementation Trilemma and 4W1H Approach

A theoretically precise model is ineffective without a practical operational process. The design of this process revealed what is termed the **Implementation Trilemma**: a critical balancing act between three often-competing goals.

- **Assessment Reliability**: The need for a credible and rigorous evaluation process.
- **System Adoption**: The need for a simple and accessible process to encourage widespread participation.
- **Operational Sustainability**: The need for a resource-efficient process that can be maintained long-term.

A strict process enhances reliability but hinders adoption, while a simple process boosts adoption but may compromise reliability. The entire operational framework of PIMM, structured around the "4W1H" methodology, is designed to systematically resolve this trilemma.

Table 4: Key Operational Design Challenges within the 4W1H Framework

4W1H Framework	Key Operational Design Challenges
Why (Incentives to Apply)	**Designing incentives** that encourage companies to apply voluntarily and continuously.
What (Assesment Target)	**Defining the assessment scope** by balancing the promotion of participation (**Adoption**) with assessing practical effectiveness (**Reliability**).
Who (Assesment System)	**Building an assessment framework** that combines the scalability to handle increasing applications (**Sustainability**) with highly dependable assessments (**Reliability**).
How (Assessment Method)	**Designing an efficient assessment process** that eliminates individual subjectivity and dependency.
When (Assessment Frequency and Cycle)	**Designing an assessment cycle** that prevents opportunity loss for companies while managing the administrative workload.

4.2 System Design for Creating Incentives for Applicant Companies (Why)

PIMM is designed to incentivize voluntary and continuous utilization by providing two primary value propositions.

4.2.1 Continuous Improvement Using Assessment Results

The primary value for a participating company is the ability to collectively evaluate its current baseline and obtain clear guidelines for future improvement actions. Beyond a simple pass/fail outcome, companies receive a detailed radar chart comparing their scores against industry averages, accompanied by actionable advice for reaching the next maturity level (see Figure 4). This continuous improvement cycle serves as a strong incentive for participation.

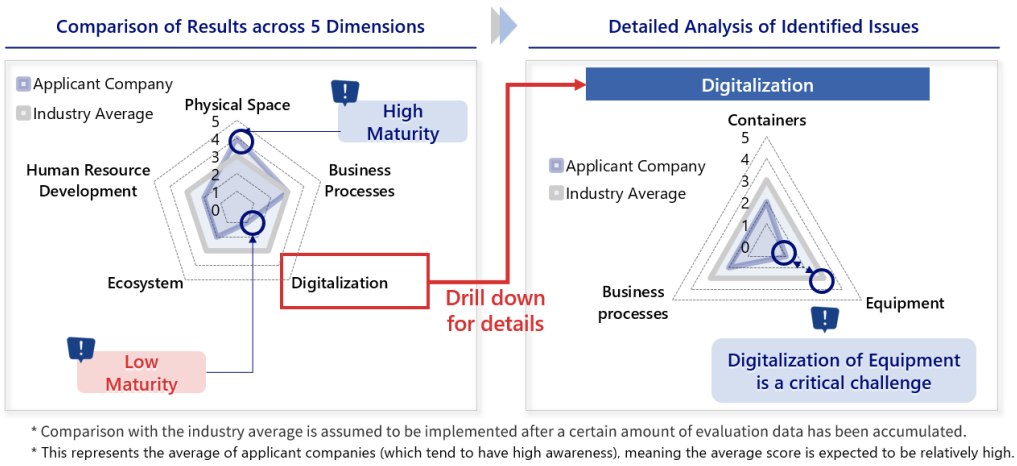


Figure 4：Visualization of Current Status and Identified Issues

4.2.2 Formation of a Community by Certificate Holders

The second value is membership in an exclusive community of certified organizations. Certification grants access to a network of "PIMM Certificate Holders," which facilitates the sharing of best practices and enables reliable partner matchmaking among companies that share a mutual understanding of standardization and digitalization.

4.3 Design of the Assessment Target (What)

The next consideration is defining"**What**" qualifies as an eligible evaluation target. To ensure assessment reliability and maintain the framework's credibility, evaluations are strictly limited to initiatives that are practically deployed and expected to operate continuously. This deliberate focus excludes transient Proof of Concept (PoC) projects that could otherwise dilute the value of certification and hinder meaningful system adoption.

4.4 Design of the Assessment System (Who)

Addressing "**Who**" conducts the assessment is critical for sustaining the system. To ensure operational sustainability while managing an increasing volume of applications, PIMM employs a decentralized assessment system. Appraisers—initially comprising academic experts and eventually incorporating personnel from highly mature certified companies—collaborate with the governing organization to deliver reliable evaluations grounded in current practical knowledge.

4.5 Design of the Assessment Method (How)

Addressing "**How**" the assessment is conducted is crucial for maintaining consistency. To mitigate the subjective variations inherent in a decentralized system without overburdening resources, PIMM utilizes a two-stage assessment process. Initially, appraisers conduct individual primary evaluations, explicitly declaring their respective confidence levels and areas of expertise. Subsequent committee consultations then prioritize discussions on borderline cases and low-confidence evaluations, thereby ensuring both objectivity and operational efficiency.

4.6 Design of Assessment Frequency and Cycle (When)

The final operational consideration is "**When**" to accept applications and conduct assessments. To balance administrative efficiency with the timely improvement needs of participating companies, PIMM adopts a "year-round application, quarterly review" cycle. Although applications are accepted continuously, formal assessment meetings (consultations) are consolidated every three months. This mechanism effectively manages the appraisers' workload while maintaining rigorous assessment quality, thereby ensuring long-term operational sustainability.

5 Initial Trial for Practicality Verification

An initial trial of PIMM was conducted from January to March 2026 with five candidate companies, which applied for the PI Award organized by JPIC. These five companies comprised one shipper, three logistics service providers, and one inventory visibility platform. This trial revealed critical practical insights into the framework's real-world deployment.

5.1 Issues in the Application Process

Applicants struggled to distinguish quantitatively between mere system implementation and full operational utilization. Furthermore, they found it challenging to provide objective evidence that adequately captured the comprehensive scope of their collaborative initiatives.

5.2 Issues in the Assessment Process

Document reviews alone often failed to capture practical on-site innovations, leading to the potential underestimation of an applicant's actual maturity level. This limitation, coupled with the evidence-gathering challenges highlighted in Section 5.1, reaffirmed the necessity of a hybrid assessment approach that combines document reviews with online applicant interviews. To further mitigate subjective bias and improve efficiency, future iterations will explore the integration of artificial intelligence into the evaluation process.

5.3 Intention to Cooperate in Future Operations

Post-trial survey results indicated that 80% of participating companies expect PIMM to facilitate future partner matching, and 100% expressed an intention to cooperate in future operations. These findings demonstrate PIMM's potential efficacy as a foundational platform for industry co-creation.

6 Conclusion and Future Work

6.1 Conclusion

The primary contribution of the Physical Internet Maturity Model (PIMM) is the translation of the abstract Physical Internet (PI) concept into a concrete, actionable implementation tool. This study establishes the model’s uniqueness and practical utility through its comprehensive framework and operational designs. By integrating multi-stakeholder feedback, PIMM expands beyond purely technical issues to encompass fundamental organizational challenges within collaborative logistics, such as organization and human resource development, thereby designing the 17-item framework to be globally adaptable. Furthermore, the 4W1H-based operational framework effectively resolves the implementation trilemma—balancing assessment reliability, system adoption, and operational sustainability—through a decentralized two-stage assessment process, a hybrid evaluation approach (combining document reviews and online interviews), an online quarterly review cycle, and community-driven value provision.

6.2 Future Work

Future priorities initially include the broader dissemination of PIMM Version 1 leading up to the 2030 “Takeoff Phase” of Japan’s PI Roadmap. Looking further ahead, updating the model to align its target year with the FY 2035 “Acceleration Phase” will be necessary. Concurrently, the integration of artificial intelligence into the assessment process to enhance evaluation objectivity, and the empirical verification of the framework's global applicability will be pursued. Continuous refinement through stakeholder feedback remains essential for establishing PIMM as a universally recognized global standard.

References

- CMMI Institute's homepage, https://cmmiinstitute.com/, 2026/03/30.
- CMMI Institute: CMMI Model Quick Reference Guide Japanese, https://cmmiinstitute.com/resource-files/public/cmmi-model-quick-reference-guide-japanese, 2026/3/30.
- Japan Institute of Logistics Systems: Global SCM Scorecard (in Japanese), https://www1.logistics.or.jp/data/scm/gsc/, 2026/3/30.
- Emerson K, Nabatchi T, Balogh S (2012): An Integrative Framework for Collaborative Governance, Journal of Public Administration Research and Theory, 22(1):1, Pages 1–29, https://doi.org/10.1093/jopart/mur011
- Physical Internet Realization Council (2022): Physical Internet Roadmap (in Japanese), https://www.meti.go.jp/shingikai/mono_info_service/physical_internet/pdf/20220308_1.pdf, 2026/03/30.
- Alliance for Logistics Innovation through Collaboration in Europe (2020): Roadmap to the Physical Internet, https://www.etp-logistics.eu/wp-content/uploads/2022/11/Roadmap-to-Physical-Intenet-Executive-Version_Final-web.pdf, 2026/03/30.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE
LAMIH UMR CNRS 8201

A Generalized Requirement Engineering Application for PI-Hubs: Integrating Triple Bottom Line and Physical Internet Fundamental

Monica-Juliana Perez[1], Tarik Chargui[1] and Damien Trentesaux[1]

1. LAMIH, UMR CNRS 8201, University Polytechnic Hauts-de-France, 59313, Mont Houy, France

Corresponding author: Monica.PerezMorales@uphf.fr

Abstract: *This research addresses the critical architectural challenge of designing sustainable and resilient trimodal nodes within the Physical Internet (PI) paradigm. While global logistics networks are rapidly evolving toward hyperconnectivity, the internal governance of trimodal hubs remains hindered by centralized oversight models that struggle with multimodal stochasticity and often ignore the social pillar of the Triple Bottom Line (TBL). To resolve this, a formal framework of nine System Requirements (SR) is proposed, specifically engineered for a self-organizing, decentralized hub architecture. By utilizing a case study from the Hauts-de-France region, the functional interactions between external carriers, hub operators, and the environment were modeled to validate the framework. The results establish a decentralized architecture where operational throughput is strictly bounded by human-centric constraints, marking a significant step toward the realization of Industry 5.0 principles within hyperconnected logistics.*

Keywords: *Physical Internet, Multimodal PI-Hubs, Go-Green ANEMONA, Social Autonomy, Triple Bottom Line, Synchromodality, Requirements Engineering.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☒ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☐ *PI Networks,* ☐ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☐ *In-Person presentation*

1 Introduction and literature review

The global logistics landscape is currently confronting an inescapable paradigm shift, catalyzed by the convergence of mounting environmental pressures, stringent regulatory frameworks, and systemic operational fragmentation. At the vanguard of this transformation is the Physical Internet (PI), a concept that reimagines global supply chains not as mere physical transport networks, but as a decentralized ecosystem of interconnected physical data packets. In this paradigm, the movement of physical goods is governed by the structural principles of digital data networks—utilizing modularity, encapsulation, and open protocols to maximize asset utilization and connectivity.

However, while significant academic and industrial efforts have advanced the optimization of transport "arcs" (the movement between geographic points), the "intelligence" of the nodes—specifically Trimodal PI-Hubs—remains a critical architectural bottleneck. As argued by

Chargui et al. (2022), the transition from traditional cross-docking to PI-compliant hubs requires a fundamental reimagining of terminal design and management to handle the high-frequency, fragmented flows of PI -containers. This challenge is further compounded by the lack of universal visibility; (Fahim et al., 2021) highlight that current information architectures in maritime and inland ports often fail to provide the granular, item-level track-and-trace capabilities necessary for seamless PI-interconnectivity. Furthermore, while digital twin-based frameworks have been proposed by (Leung et al., 2022) to synchronize inbound order management, these models often operate under semi-centralized assumptions that struggle with the stochasticity of real-world trimodal transitions.

The operational constraints inherent in traditional, centralized logistics oversight increasingly struggle to accommodate the adaptive capacity necessitated by the high-velocity, stochastic demands of PI synchronization. In high-density trimodal environments, centralized architectures often suffer from computational tractability issues and information latency, where the time required to process global perturbations exceeds the window for effective response.

Furthermore, the emergence of Artificial Intelligence (AI) and decentralized technologies offers a pathway to reconcile these competing objectives, provided they are deployed through transparent and responsible mechanisms. Meena et al. (2025) argue that aligning technological advancement with sustainable practices requires AI transparency to achieve ethical operations within circular supply chains. This is echoed by Samuels, (2024), who observes that the transition toward Industry 5.0 emphasizes the necessity of human-machine collaboration that fosters both operational efficiency and social sustainability. In the context of a trimodal PI-hub, this implies that the system's self-organizing logic must not only synchronize flows but also act as a guardian of labor standards. Zambrano-Rey et al. (2026), propose a holonic structure providing the necessary flexibility to handle the computational tractability issues inherent in centralized logistics oversight, ensuring that the decisional logic remains consistent from a single asset level to the complex coordination of a regional hub.

The architectural evolution of PI-hubs must be framed within the critical transition from the operational efficiency of Industry 4.0 toward the human-centric imperatives of Industry 5.0. While the PI paradigm aims to establish economically and environmentally sustainable logistics ecosystems (Ben Neila et al., 2021), a fundamental imbalance persists in the integration of the three dimensions of the Triple Bottom Line (TBL). As Treiblmaier (2019) points out, incorporating TBL goals into supply chains is a complex task that requires the precise coordination of multiple stakeholders' interests; however, the social dimension is frequently underserved in digital logistics frameworks. Current optimization models for hyperconnected networks tend to treat labor as a static resource rather than a dynamic, self-governing entity, creating a "social blind spot" where safety, worker fatigue, and ethical labor standards are secondary to throughput. This gap is particularly precarious in high-density nodes, where the pace of automated synchronization can outpace human operational limits, necessitating a design philosophy that explicitly encodes social welfare into the system’s decisional structure.

The global transition toward the PI is currently hindered by an architectural bottleneck: while transport "arcs" are becoming hyperconnected, the internal governance of trimodal nodes remains trapped in centralized models. These traditional oversight systems struggle with the high-velocity stochasticity of multimodal arrivals and frequently ignore the social pillar of the TBL, treating labor as a static resource rather than a self-governing entity. To resolve this, the primary objective of this research is to move beyond descriptive models by establishing a standardized functional specification for a self-organizing PI-hub. The core contribution of this

paper is the formalization of nine System Requirements (SR) liked by hard-code sustainability goals—particularly social autonomy—into the hub's decentralized operational logic. By structuring these requirements across four strategic pillars, this work provides a verifiable roadmap for designing the next generation of resilient and ethically grounded PI-enabled infrastructures.

2 Functional Design Principles for Sustainable PI-Hubs

This section establishes the formal architectural foundation required to transform trimodal nodes into sustainable and autonomous PI-Hubs. The functional requirements are standardized across four distinct pillars, derived from a synthesis of the foundational principles of the PI as defined by Montreuil et al., (2012) showed in Figure 1. The purpose of this integration is to reconfigure the original PI foundations into four categories that specifically address the operational requirements of high-density nodes

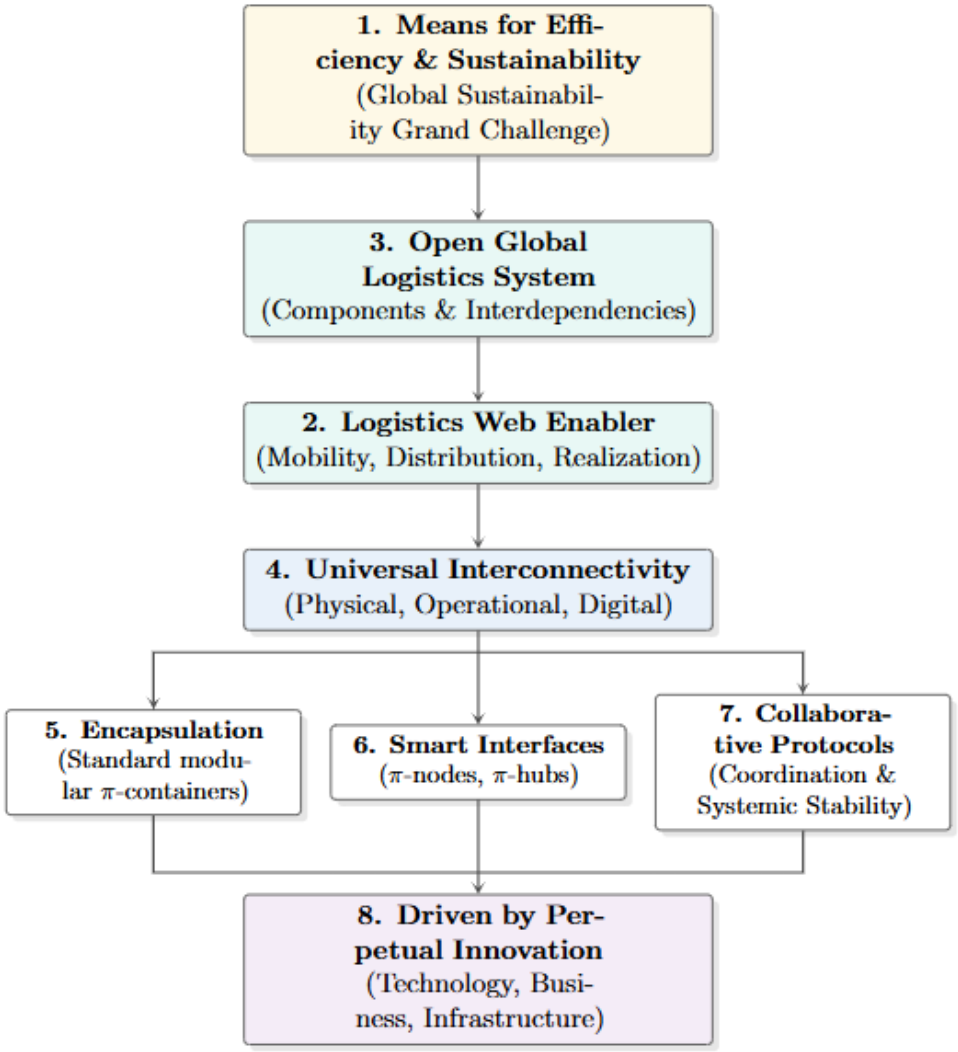


Figure 1:Model of the 8 Foundations of the PI paradigm

Pillar I: Encapsulation and Product Intelligence

This pillar defines the transition of cargo from a passive object to an active, self-aware entity. Intelligence is treated as a functional prerequisite for decentralization; without it, autonomous

routing within the trimodal node becomes impossible. This principle aligns with the fifth PI foundation regarding encapsulation, ensuring that the arrival of a PI-container at the hub triggers the product's own decisional logic for autonomous state monitoring and geolocation tracking. Furthermore, in accordance with the sixth PI foundation concerning standard smart interfaces, each smart PI-container utilizes a digital representative (or "smart tag") acting as its agent in the Digital Internet. This guarantees continuous identification, routing, and traceability, which are essential for enabling distributed handling and storage automation within the hub environment. By decoupling the container's intelligence from proprietary local databases, the system ensures that the cargo remains visible and accountable throughout its entire transit cycle.

Pillar II: Interconnectivity and Synchronization

The second pillar focuses on the systemic challenge of trimodal integration, establishing the requirements for harmonizing different transport modes into a synchronized operational flow. By viewing the hub through a Lean perspective, storage is identified as a non-value-added activity—essentially operational waste—that the system seeks to minimize through active synchronization and "just-in-time" transshipments. This necessitates an architecture capable of supporting "Synchromodality," where modal shifts are seamless and mode-agnostic. In accordance with the fourth PI foundation of universal interconnectivity and the second foundation defining the PI as a logistics web enabler, the goal is total fluidity. Within this framework, any physical entity must be able to flow transparently through interconnected physical, digital, and operational layers. By prioritizing real-time interconnectivity over static buffering, the hub evolves from a traditional storage node into a high-velocity transition point that can dynamically recalibrate sequences in response to external network variability (Montreuil, 2011; Montreuil et al., 2012).

Pillar III: Decentralized Coordination and Systemic Resilience

The third pillar establishes the hub as an adaptive system capable of maintaining operational continuity under stress. By conceiving the environment as a "disruption source," the hub requires a resilience framework founded on decentralized coordination rather than centralized restarts. In alignment with the third PI foundation of an open global logistics system, the hub is designed as a set of autonomous components whose well-defined interdependencies allow the system to meet its objectives even during local failures. Furthermore, the seventh PI foundation regarding standard collaborative protocols ensures network reliability. These protocols guarantee that if a specific asset or path becomes non-operational, PI-container traffic is rerouted automatically, protecting the node’s robustness against unforeseen stochastic events.

Pillar IV: Integral Sustainability

The final pillar formalizes the integration of the TBL by internalizing labor limits and safety protocols as hard architectural constraints. Within this framework, the system’s decisional logic is forced to evaluate all potential assignments through a multi-objective lens, effectively preventing the sub-optimization of human welfare in favor of pure operational throughput. This structure responds directly to the first PI foundation, which defines the paradigm as a means to achieve global efficiency by addressing the "Grand Challenge" of economic, environmental, and social sustainability (Montreuil, 2011). By embedding specific TBL Key Performance Indicators (KPIs) into the core logic, the framework ensures that every operational flow remains strictly bounded by ethical standards and social responsibility.

3 Proposed Requirements for PI-Hubs

This section presented the proposing a requirement framework for multimodal hubs. Nine critical System Requirements (SR)—structured across four pillars established in the previous section—are identified to govern the behavior of autonomous entities. By formalizing these requirements, the transition is made from a conceptual use case to a standardized functional specification that serves as the architectural foundation for the hub's intelligence. This standardization is for guaranteeing that decentralization does not compromise systemic resilience or ethical labor standards during real-time operational synchronization.

3.1 Pillar I: Encapsulation and Product Intelligence

This pillar defines the transition of cargo from a passive object to an active, self-aware entity (Sallez et al., 2015). Within the PI paradigm, intelligence is a functional prerequisite for decentralization; without it, autonomous routing becomes impossible.

- **SR-01: Autonomous Monitoring**

The operation of a trimodal hub necessitates a continuous and persistent synchronization between the physical PI-container and its digital counterpart. This requirement dictates that each cargo unit must independently maintain its state information—including geolocation, volumetric occupancy, and destination—throughout the entire transit cycle. By formalizing this synchronization, the system ensures that critical data remains intrinsically linked to the physical unit, thereby eliminating information silos during the modal shifts between barge, rail, and road. This autonomy in information management is what allows the container to remain "visible" and "accountable" across the different transport networks of the without depending on proprietary local databases

- **SR-02: Decisional Autonomy and Negotiation Capability**

Beyond information mirroring, the system requires that each container holon possesses the agency to act as a proactive participant in the hub’s operations. This requirement empowers the cargo unit to evaluate transport offers against its internal objectives. Specifically, when the hub experiences stochastic perturbations (e.g., train delays or congestion), the container holon must be capable of autonomously triggering negotiation protocols with available resources. This enables the entity to independently resolve trade-offs between environmental impact and temporal constraints, ensuring that routing decisions are optimized for sustainability at the local level without necessitating a total system reconfiguration.

3.2 Pillar II: Interconnectivity and Synchronization

The second pillar addresses the systemic challenge of trimodal integration. It defines the requirements for harmonizing disparate transport modes—road, rail, and water—into a synchronized operational flow. In a PI environment, the hub must evolve from a static buffer into a high-velocity transition point through decentralized coordination.

- **SR-03: Protocol Standardization and Functional Abstraction**

To mitigate the operational fragmentation typical of conventional logistics nodes, the architecture requires a standardized framework for resource interaction. This requirement dictates that all hub assets—regardless of their physical nature or manufacturer—must be characterized through a functional abstraction of their capabilities. By decoupling specific hardware constraints from the decision-making logic, the system ensures that transport demands are matched with available capacity based on service performance rather than proprietary interfaces. This universal interoperability is for maintaining a mode-agnostic flow, allowing the regional network to function as a single, integrated infrastructure.

- **SR-04: Dynamic Modal Synchronization (Synchromodality)**

A sustainable PI-hub must resolve the inherent misalignment between stochastic arrival patterns and scheduled departure windows. This requirement establishes the need for real-time synchronization of diverse transport flows, utilizing continuous data streams to align the terminal's internal operations with external network variability. The system must possess the agility to dynamically recalibrate resource allocation and transit sequences in response to fluctuations in vessel, train, or vehicle availability. This "synchromodal" requirement directly minimizes idling intervals and unutilized capacity, thereby fulfilling the environmental mandate of the TBL objectives by reducing the carbon intensity of terminal operations.

3.3 Pillar III: Decentralized Coordination and Systemic Resilience

The third pillar establishes the requirements for maintaining operational continuity under stress. In a trimodal environment, disruptions—such as resource failures or stochastic congestion—must be managed through local adaptation rather than global restarts.

- **SR-05: Distributed Optimization and Localized Processing**

The system architecture must avoid single points of failure by ensuring that the computational and decisional burden is distributed across the network. This requirement dictates that the logic for task allocation and resource matching must reside at the edge of the system. By enabling localized processing, the hub ensures that a localized disruption does not propagate into a system-wide failure.

- **SR-06: Self-Organization and Autonomous Reconfiguration**

To ensure resilience, the system must possess the capacity for autonomous reconfiguration in the event of resource unavailability. This requirement establishes that when a specific asset becomes non-operational (e.g., equipment failure or labor unavailability), the affected transport flows must possess the logic to re-negotiate their schedules and routing locally. This self-healing property ensures that the hub can bypass bottlenecks dynamically.

3.4 Pillar IV: Integral Sustainability

The final pillar formalizes the integration of the TBL. It ensures that economic efficiency does not compromise environmental targets or social standards, embedding these dimensions into the system’s core decision-making logic.

- **SR-07: Dynamic Objective Weighting**

The architecture requires a mechanism to adjust operational priorities in real-time. This requirement dictates that the criteria governing transport negotiations must be sensitive to shifting external factors, such as fluctuating energy costs or regional environmental regulations. By allowing the weighting of sustainability objectives to be dynamic, the system can transition between "High-Efficiency" and "Low-Carbon" modes as dictated by the current operational context.

- **SR-08: Multi-Objective Decision Support**

Every autonomous negotiation within the hub must be evaluated through a multi-objective lens. This requirement mandates that no single dimension of the TBL (economic, environmental, or social) can be optimized in total isolation. Decisions regarding modal shifts must simultaneously consider the cost, the carbon intensity, and the service-level agreements.

- **SR-09: Social Autonomy**

Unique to this framework is the requirement for social labor and safety standards to act as a hard constraint on optimization. This requirement introduces a social decision logic, where a proposed operational optimization is automatically invalidated if it violates

predefined human-centric standards, such as maximum labor fatigue limits or safety protocols.

The nine system requirements (SR-01 to SR-09) constitute a comprehensive roadmap for the architectural design of sustainable trimodal hubs. Unlike traditional centralized models, this decentralized specification shifts the operational burden from a single point of failure to a resilient network of autonomous entities.

By prioritizing functional abstraction and dynamic synchronization, the framework ensures that the hub remains mode-agnostic, facilitating the seamless integration of barge, rail, and road transport. Furthermore, the inclusion of the social autonomy mechanism (SR-09) ensures that the pursuit of operational efficiency remains ethically bounded by social labor standards, directly addressing the critical human-centric gap in current PI research.

4 Application and discussion

To implement the methodology this work is using as case study a trimodal hub system located in north of France region. This center acts as a node integrating waterway, railway, and roadway transportation. From a functional analysis perspective, this system represents a high-density scenario with an average daily flow of 452 TEU. The spatial configuration of the terminal is showed in Figure 2. This section is going to show a use case stay and after the requirements based on the study

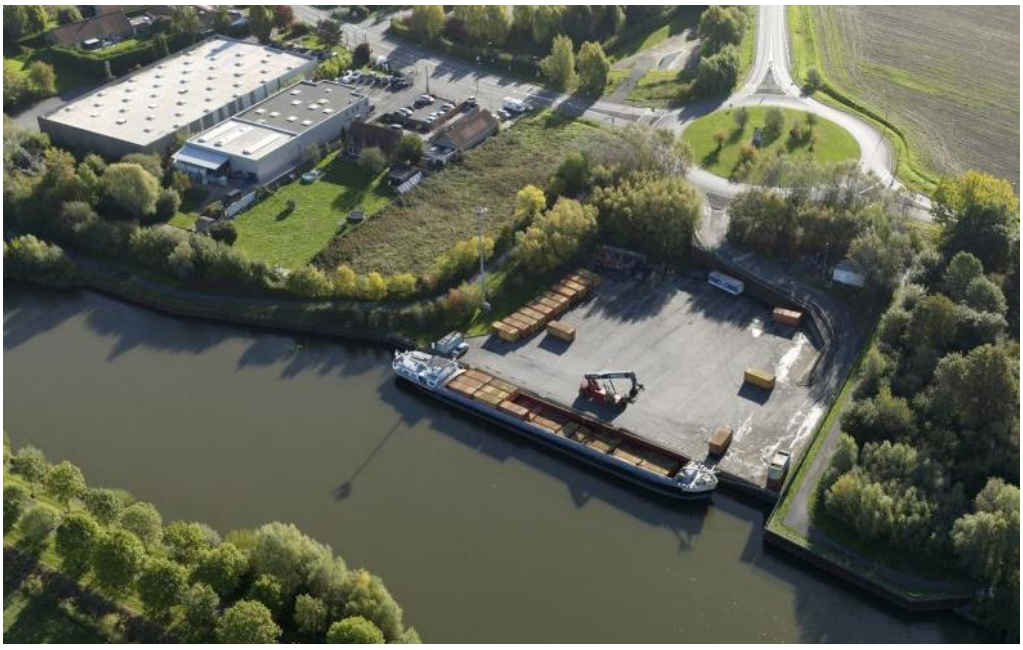


Figure 2: Aerial view of the Trimodal Hub. Source: Adapted from Ports de Lille (2024).

4.1 Operational Mapping and Actor Interaction

The functional scope of the trimodal PI-hub is structured through a comprehensive Use Case (UC) model showed in Figure 3. Each UC is to decompose different partitions of the terminal's complexity into specialized operational packages. This decomposition ensures that every physical movement within the hub is governed by a decentralized decision-making layer, adhering to the holonic principles of autonomy and cooperation. The interaction between the hub and its environment is mediated by three primary actors:

- **Carrier (External):** Represents the multimodal transport entities (Road, Rail, and Water) that initiate inbound flows and receive outbound shipments.
- **Hub Operator:** Acts as the "Policy Setter," responsible for defining the weight distribution for the TBL objectives, thus balancing economic throughput with social and environmental mandates.
- **Environment:** Functions as the "Disruption Source," introducing stochastic events—such as vessel delays or equipment failures—that trigger the system's resilience mechanisms.

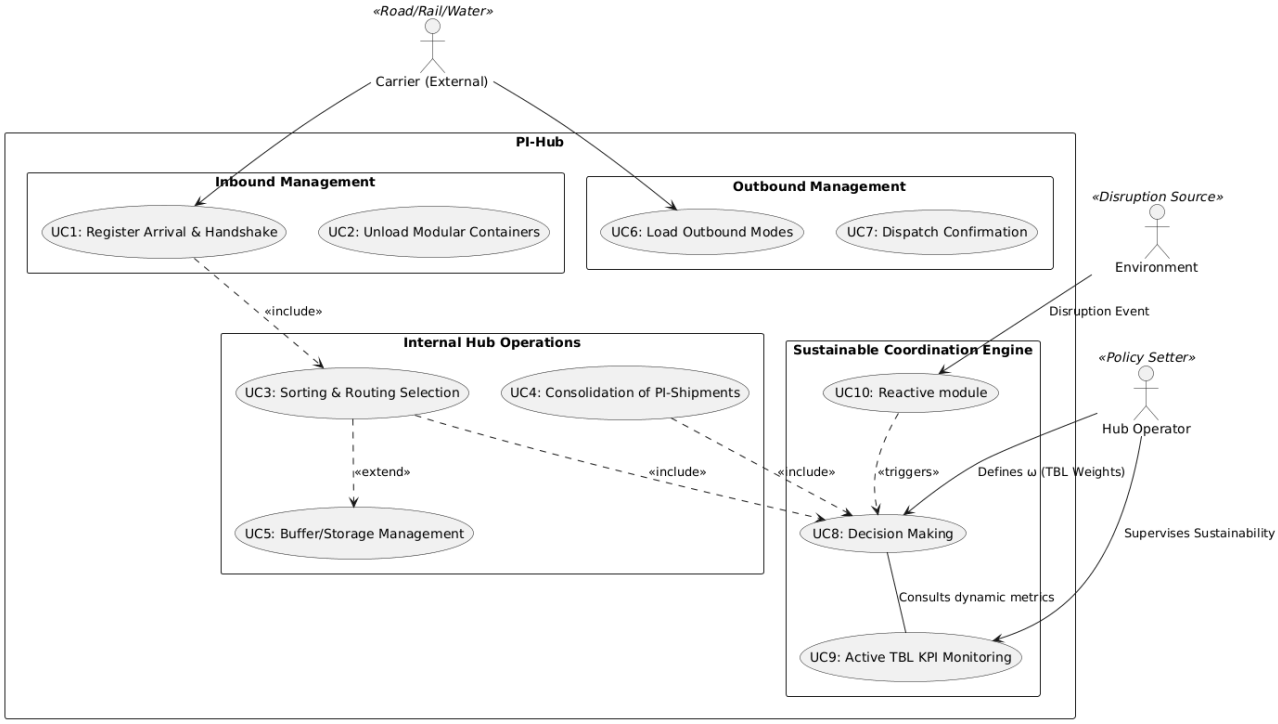


Figure 3: Use case case study trimodal hub system

The internal logic of the trimodal hub is governed by four strategic interactions that ensure the seamless movement of PI-containers. The process begins with Inbound Synchronization, where the initial registration and handshake (UC1) triggers an immediate, dynamic calculation of the container's internal path (UC3) based on the terminal's real-time state. This routing, alongside the consolidation of PI-shipments (UC4), is managed by a Sustainable Decision-Making engine (UC8)—the system's "decisional brain"—which continuously evaluates potential assignments against active TBL metrics (UC9). To maintain Systemic Resilience, any environmental perturbations trigger a reactive module (UC10) that prioritizes the rapid re-establishment of flow through automated re-calculations. Finally, the architecture adopts a Lean Perspective by treating storage as a non-value-added activity (UC5); consequently, the system seeks to minimize buffering through "just-in-time" transshipment and active synchronization, effectively reducing operational waste within the high-density terminal environment.

4.2 Implementation of PI Requirements

The requirements identified center on the coordination of sorting processes within a central grid where cargo units must undergo modal shifts efficiently. The primary goal is to minimize inefficiencies related to empty container movements and reduce CO_2 emissions through

enhanced cargo consolidation. The formalization of the nine system requirements (SR-01 to SR-09) serves as a design proposal for the engineering of high-density trimodal hubs based on PI paradigm. By transitioning from a centralized oversight model to one governed by Autonomous State Synchronization (SR-01) and Decisional Autonomy (SR-02), the architecture mitigates the "information myopia" inherent in traditional logistics nodes. This shift ensures that operational intelligence is distributed at the system's edge, allowing individual entities to manage localized data and agency in real-time. Consequently, the architecture verifies that decentralization is a functional necessity for maintaining systemic fluidity and transparency, particularly when navigating the stochastic perturbations typical of multimodal regional networks.

The systemic synergy between Functional Abstraction (SR-03) and Dynamic Temporal Synchronization (SR-04) further proves the framework's capacity for true synchromodality. By decoupling decision-making logic from proprietary hardware constraints, the architecture facilitates a "buffer-less" operational flow where the arrival of stochastic barge or rail units triggers an immediate, autonomous recalibration of road departure schedules. This capability directly reduces terminal dwell times and unutilized capacity, aligning the hub's performance with the environmental mandates of the TBL. Furthermore, systemic resilience emerges as an inherent property of Distributed Optimization (SR-05) and Self-Organization (SR-06). The resulting "self-healing" mechanism ensures that the global throughput remains optimized even during infrastructure failures or labor disruptions, providing a robust adaptability that is unattainable in rigid, top-down management schemes. Ultimately, the integration of Social Autonomy (SR-09) elevates this framework into an ethical engineering standard. By formalizing social labor and safety standards as hard architectural constraints, the system ensures that no degree of economic or environmental optimization can be achieved through the violation of human-centric protocols.

5 Conclusions and future work

This research establishes a formal bridge between the theoretical foundations of the PI and the functional engineering of sustainable trimodal nodes. By transitioning from a centralized oversight model to a decentralized, product-led architecture, this work successfully mitigates the "information myopia" and computational bottlenecks inherent in traditional logistics nodes. The primary contribution is the formalization of nine System Requirements (SR) structured across four strategic pillars: Encapsulation, Interconnectivity, Decentralization, and Integral Sustainability.

A significant advancement of this framework is the architectural integration of the TBL ensuring that social and environmental constraints are not merely secondary performance indicators but fundamental drivers of the hub's operational logic. Specifically, the inclusion of Social Autonomy (SR-09) as a architectural constraint ensures that the pursuit of logistics efficiency remains ethically bounded by human-centric standards, marking a definitive step toward the realization of Industry 5.0 principles within the PI paradigm. Moving forward, the primary trajectory for research involves the functional translation of these requirements into a dynamic simulation environment. While this work focuses on the formal derivation of the system's governing logic, subsequent research will deploy an Agent-Based Model (ABM) to perform extensive performance benchmarking. This next phase will evaluate the emergent behaviors of the system under the stochastic perturbations identified in this study, such as equipment failures and variable modal arrival times, to further validate the resilience of decentralized PI-Hub governance.

6 References

Barbosa, J. (2015). *Self-organized and evolvable holonic architecture for manufacturing control* [Theses, Université de Valenciennes et du Hainaut Cambrésis]. https://theses.hal.science/tel-01137643

Ben Neila, S., Rejeb, A., & Németh, P. (2021). The Interplay between the Physical Internet and Logistics: A Literature Review and Future Research Directions. *Acta Technica Jaurinensis*, *15*(1), 22–35. https://doi.org/10.14513/actatechjaur.00638

Chargui, T., Ladier, A.-L., Bekrar, A., Pan, S., & Trentesaux, D. (2022). Towards designing and operating physical internet cross-docks: Problem specifications and research perspectives. *Omega*, *111*, 102641. https://doi.org/10.1016/j.omega.2022.102641

Fahim, P. B. M., An, R., Rezaei, J., Pang, Y., Montreuil, B., & Tavasszy, L. (2021). An information architecture to enable track-and-trace capability in Physical Internet ports. *Computers in Industry*, *129*, 103443. https://doi.org/10.1016/j.compind.2021.103443

Giret, A., Trentesaux, D., Salido, M. A., Garcia, E., & Adam, E. (2017). A holonic multi-agent methodology to design sustainable intelligent manufacturing control systems. *Journal of Cleaner Production*, *167*, 1370–1386. https://doi.org/10.1016/j.jclepro.2017.03.079

Leung, E. K. H., Lee, C. K. H., & Ouyang, Z. (2022). From traditional warehouses to Physical Internet hubs: A digital twin-based inbound synchronization framework for PI-order management. *International Journal of Production Economics*, *244*, 108353. https://doi.org/10.1016/j.ijpe.2021.108353

Meena, R., Sahoo, S., Malik, A., Kumar, S., & Nguyen, M. (2025). Artificial intelligence and circular supply chains: Framework for applications and deployment from the triple bottom line model perspective. *Annals of Operations Research*, *354*(1), 71–101. https://doi.org/10.1007/s10479-025-06510-1

Montreuil, B., Meller, R. D., & Ballot, E. (2012). Physical Internet Foundations. *IFAC Proceedings Volumes, 14th IFAC Symposium on Information Control Problems in Manufacturing*, *45*(6), 26–30. https://doi.org/10.3182/20120523-3-RO-2023.00444

Sallez, Y., Montreuil, B., & Ballot, E. (2015). On the Activeness of Physical Internet Containers. In T. Borangiu, A. Thomas, & D. Trentesaux (Eds.), *Service Orientation in Holonic and Multi-agent Manufacturing* (pp. 259–269). Springer International Publishing. https://doi.org/10.1007/978-3-319-15159-5_24

Samuels, A. (2024). Human-centric technologies in sustainable supply chain management: A systematic review of the evolution from Industry 4.0 to 5.0. *International Journal of Business Ecosystem & Strategy (2687-2293)*, *6*(4), 285–302. https://doi.org/10.36096/ijbes.v6i4.539

Treiblmaier, H. (2019). Combining Blockchain Technology and the Physical Internet to Achieve Triple Bottom Line Sustainability: A Comprehensive Research Agenda for Modern Logistics and Supply Chain Management. *Logistics*, *3*(1), 10. https://doi.org/10.3390/logistics3010010

Zambrano-Rey, G., Perez, M.-J., Chargui, T., Sallez, Y., & Trentesaux, D. (2026). Design and control of physical internet multimodal hubs based on the go-green ANEMONA methodology. *International Journal of Logistics Research and Applications*, 1–39. https://doi.org/10.1080/13675567.2025.2608073

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE

Towards Competitive Inland Waterway Logistics: Virtual Models for Modal Shift Assessment

José Luis López[1] and David Ciprés[1]
1. Instituto Tecnológico de Aragón, Zaragoza, España
Corresponding author: jlopez@ita.es

Abstract: *Inland Waterway Transport (IWT) holds substantial potential to enhance urban and peri-urban freight sustainability, yet adoption remains limited due to uncertainties in operational, economic, and systemic performance. This paper presents a virtual modelling and simulation framework to assess IWT competitiveness within multimodal logistics networks. The approach combines agent-based modelling, route optimisation, and Monte Carlo simulation to evaluate modal shift potential under uncertain conditions. Competitiveness is measured across environmental, efficiency, reliability, safety, cost, and scalability dimensions. Applied to urban logistics in Ghent and Caen, the framework identifies utilisation thresholds, key drivers for modal shift, and the enabling role of collaboration and emerging Small Flexible Automated and Zero-emission Vessels (SFAZ). The results inform decision-making for more sustainable, Physical Internet-aligned urban logistics systems.*

Keywords: *Inland Waterway Transport, Modal Shift, Agent-Based Simulation, Urban Logistics, Physical Internet.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan: ☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories), ☐ Transportation Equipment, ☒ PI Networks, ☐ System of Logistics Networks, ☐ Vertical Supply Consolidation, ☐ Horizontal Supply Chain Alignment, ☐ Logistics/Commercial Data Platform, ☒ Access and Adoption, ☐ Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

1 Introduction

Urban freight systems are facing increasing pressure to reduce emissions, congestion, and energy consumption under tightening regulatory frameworks and growing spatial constraints. In this context, Inland Waterway Transport (IWT) is widely recognised for its superior energy efficiency and environmental performance compared to road-based solutions. Despite these advantages, its integration into urban and peri-urban logistics systems remains limited. This gap is not primarily driven by technological constraints, but rather by structural and organisational barriers, including fragmented demand, insufficient coordination among stakeholders, and limited understanding of system-level performance under realistic operating conditions.

Recent developments in Small Flexible Automated and Zero-emission Vessels (SFAZ) offer new opportunities to overcome some of these limitations. By combining automation, electrification, and smaller vessel sizes—such as Class I and II barges—SFAZ systems can access waterways and urban areas that are not feasible for conventional diesel vessels, enabling new operational concepts and extending the reach of IWT within cities. However, their economic and operational viability is not guaranteed. The competitiveness of SFAZ-enabled

IWT depends critically on demand density, asset utilisation, and the degree of integration within multimodal logistics networks. This creates a key challenge for decision-makers: under which conditions should SFAZ systems be deployed, and how do they perform relative to conventional road transport under uncertainty?

This paper proposes an integrated simulation-based framework that combines agent-based modelling, optimisation, and uncertainty analysis. The contribution is threefold:

- A multi-layer framework linking operations, evaluation criteria, and system performance.
- A systematic assessment of modal shift conditions under both deterministic and stochastic settings.
- Empirical insights from two contrasting case studies, identifying threshold conditions and coordination requirements for scalable IWT adoption.

2 Background and related work

IWT has increasingly been recognized as a viable and sustainable alternative to road freight, particularly in urban and peri-urban logistics where road congestion, environmental impacts, and energy consumption pose significant challenges. IWT offers inherent advantages in terms of energy efficiency, lower emissions, and capacity to transport large volumes of goods with reduced spatial requirements compared to conventional road transport (Calderón-Rivera et al., 2024). Despite these benefits, the adoption of IWT in integrated logistics chains remains limited, largely due to infrastructural constraints, limited operational flexibility, and the complexity of coordinating multimodal networks.

The integration of IWT into multimodal logistics systems has been studied, with research emphasizing both policy-driven and operational perspectives. Santén et al. (2021) provide a series of case studies demonstrating how targeted interventions, such as improved terminal infrastructure and incentivization schemes, can facilitate modal shift from road to waterways. Complementing these empirical studies, Beil et al. (2026) analyze the transferability of policy measures across regions using multicriteria evaluation, highlighting the importance of context-specific strategies to overcome barriers to modal shift. Collectively, these studies underscore the necessity of aligning policy frameworks, infrastructure development, and operational practices to enable effective integration of IWT into broader logistics networks.

Simulation approaches, particularly discrete event and agent-based models, have emerged as essential tools for evaluating IWT operations and potential modal shifts. Shobayo et al. (2024) employ discrete event simulation to quantify operational performance and logistical interactions along the Rhine-Alpine corridor, providing insights into bottleneck identification and throughput optimization. Similarly, Benedecti et al. (2024) develop an agent-based model of an inland container terminal and associated waterway services, enabling detailed analysis of terminal operations, vessel scheduling, and modal interaction. These modeling techniques allow researchers and practitioners to assess the operational feasibility of IWT integration under varying demand scenarios, providing a quantitative basis for decision-making in multimodal transport planning.

In addition to operational modeling, the evaluation of performance and sustainability outcomes requires robust Key Performance Indicator (KPI) frameworks. Duldner-Borca et al. (2023) propose a KPI framework tailored to IWT, capturing economic, logistical, and environmental dimensions. The application of such frameworks facilitates systematic assessment of both short-term operational performance and long-term sustainability objectives, offering a standardized basis for comparing interventions and supporting strategic planning.

Despite these advances, several research gaps remain. Existing studies often focus on specific corridors or terminals, limiting the generalizability of findings across diverse geographic and operational contexts. Moreover, while simulation models provide valuable insights, they require high-quality data inputs and calibration to accurately reflect real-world dynamics. Finally, the integration of KPI frameworks with simulation outputs remains underexplored, limiting the ability to evaluate the combined effects of operational and policy measures on multimodal performance. In addition, the role of emerging vessel concepts—such as SFAZ vessels—is not yet sufficiently addressed in the literature. In particular, there is limited evidence on how their operational characteristics (e.g., automation, modularity, and high-capacity utilisation) influence system-level performance, scalability, and modal shift potential under varying demand and uncertainty conditions. Addressing these gaps provides a clear pathway for the present study, which aims to combine multi-agent simulation, modal shift evaluation, and KPI-based performance assessment to inform the sustainable integration of IWT in urban and peri-urban logistics.

3 Methodological framework

The proposed framework is designed to assess the competitiveness of IWT with particular emphasis on SFAZ vessels, within complex multimodal logistics systems. The approach explicitly captures both operational dynamics and system-level performance, enabling the evaluation of modal shift potential under deterministic and uncertain conditions.

The system is defined over an integrated multimodal network comprising inland waterways and road infrastructure, where freight flows are represented across first-, main-, and last-mile operations. Within this setting, vehicles, vessels, and infrastructure interact dynamically, shaping overall system performance.

The framework follows a structured analytical logic that links operational processes to decision-relevant outcomes (Figure 1). Input data describing demand, infrastructure, and vehicle characteristics are translated into a network-based representation of the transport system. This representation feeds an agent-based simulation model, which generates time-dependent operational outputs. These outputs are subsequently transformed into performance indicators, enabling the comparative assessment of transport configurations and the evaluation of modal competitiveness. This sequential structure ensures consistency between micro-level dynamics and macro-level insights.

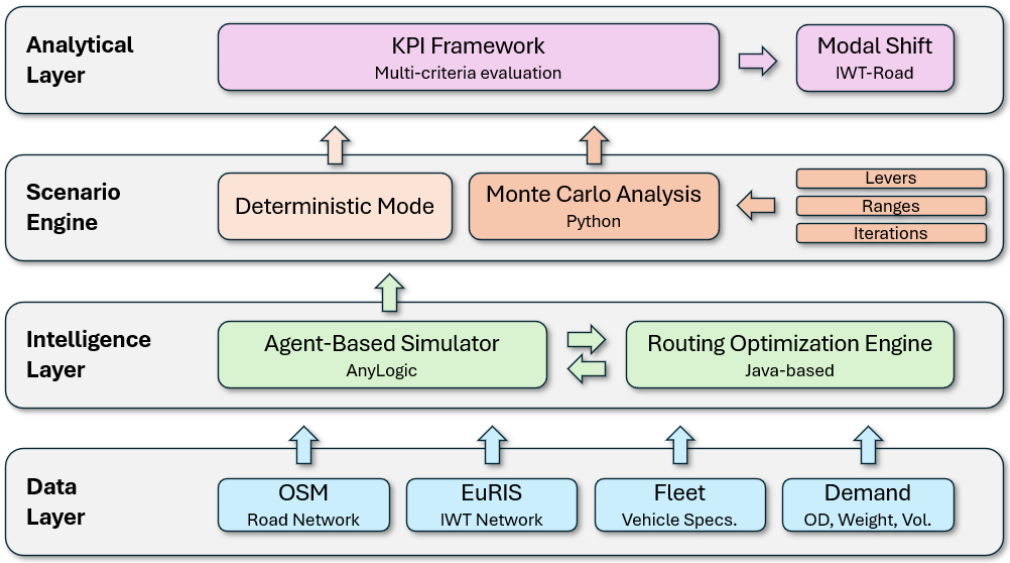


Figure 1: Simulation framework for modal shift assessment

3.1 Multimodal Network Modeling

The representation of the transport network constitutes a critical component of the modelling framework, as it defines the spatial and operational environment within which logistics activities take place. In this study, the transport system is modelled as a multimodal graph integrating both inland waterway and road networks.

The inland waterway network is represented through a graph-based structure in which nodes correspond to physical infrastructure elements such as locks, bridges, and junctions, while edges represent navigable fairways. Each fairway is characterised by a set of attributes, including length, permitted speed, and structural constraints such as bridge clearance or water depth. This level of detail allows the model to capture the operational limitations that directly influence route feasibility and travel time. Additional data from the European River Information Services (EuRIS) are incorporated to ensure consistency with real-world navigation conditions, including lock operations, fairway classifications, and regulatory constraints.

The road network is derived from OpenStreetMap data, providing a detailed and flexible representation of urban and regional transport infrastructure. The integration of road and waterway layers enables the construction of a seamless multimodal network in which different transport modes coexist and interact. This is particularly important for modelling first- and last-mile operations, which often determine the overall competitiveness of inland waterway solutions.

The combination of these data sources enables a high level of realism while maintaining flexibility for application across different geographical contexts.

3.2 Route Optimisation

Efficient routing is a key determinant of transport performance, particularly in inland navigation, where physical and operational constraints can significantly influence route selection. To address this, the framework incorporates a dedicated IWT routing algorithm based on a graph optimisation approach.

The algorithm identifies feasible and optimal routes between origin and destination nodes by evaluating multiple criteria, such as travel distance or time, while simultaneously accounting for navigational constraints. These constraints include structural limitations, such as bridge clearance, as well as operational factors like lock passage times. If a given fairway does not satisfy the requirements of a specific vessel—for instance due to insufficient clearance—it is excluded from the feasible solution space.

This approach extends traditional shortest-path algorithms by incorporating multi-objective optimisation and constraint-based filtering, enabling a more realistic representation of inland navigation. Moreover, the algorithm is dynamically integrated into the simulation environment, allowing each vessel agent to compute routes in simulation time based on its specific characteristics, such as size, speed, or air draft.

As a result, routing decisions are not static but adapt to both network conditions and vehicle attributes, reflecting the operational complexity of real-world logistics systems.

3.3 Agent-Based Simulation Model

The core of the proposed framework is a multimodal agent-based simulation model designed to reproduce the behaviour and interactions of logistics actors within a dynamic transport environment. In this model, vehicles, vessels, and infrastructure elements are represented as autonomous agents, each governed by its own objectives, constraints, and decision rules.

This modelling approach allows for the explicit representation of decentralised decision-making processes and heterogeneous actor behaviour. Vehicle agents, including SFAZ vessels, conventional barges, and trucks, operate independently while interacting with shared infrastructure such as locks, terminals, and road networks. Their behaviour includes routing decisions, energy management, cargo handling, and responses to network conditions.

Infrastructure elements are also modelled with operational characteristics, such as lock schedules, bridge constraints, and terminal capacities, which influence system dynamics. The interaction between vehicle and infrastructure agents generates emergent behaviours, such as congestion at bottlenecks or variations in service reliability.

The integration of real-world data sources ensures that the simulation environment reflects realistic spatial and operational conditions. Furthermore, the model includes dynamic data collection and visualisation capabilities (Figure 2), enabling both simulation-time monitoring and detailed post-simulation analysis. Performance indicators such as travel times, energy consumption, and utilisation rates are continuously recorded, providing the basis for KPI computation.

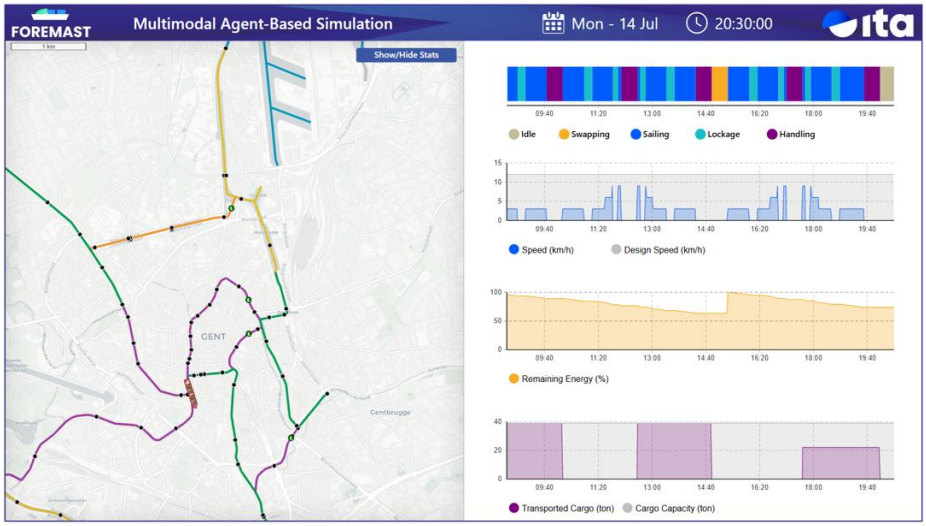


Figure 2: Multimodal Agent-Based Simulation Model

3.4 Monte Carlo Simulation and Uncertainty Modelling

Freight transport systems are inherently subject to uncertainty, arising from fluctuations in demand, variability in travel times, operational disruptions, and changing market conditions. Capturing this uncertainty is essential for a robust assessment of modal competitiveness. To this end, the framework adopts a Monte Carlo simulation approach in which key input parameters are treated as stochastic variables rather than fixed values. These parameters include demand levels, shipment sizes, loading factors, travel times, handling durations, and degrees of collaboration among stakeholders.

Each simulation run represents a distinct realisation of the logistics system under a specific set of conditions. By performing a large number of simulations, the framework generates distributions of performance outcomes rather than single deterministic results. This enables the identification of robust trends, sensitivities, and threshold effects that are critical for understanding modal shift dynamics. In particular, the Monte Carlo approach allows the analysis to move beyond average-case scenarios and instead explore the range of conditions under which inland waterways become competitive. This is especially relevant in the context of emerging technologies such as SFAZ vessels, whose performance and adoption depend on multiple uncertain factors.

3.5 KPI-Based Performance Framework

A fundamental element of the proposed methodology is the use of KPIs to systematically evaluate the performance of IWT and, in particular, SFAZ vessels. In this context, KPIs are defined as measurable indicators that reflect how effectively a transport system achieves its operational, environmental, and economic objectives.

This KPI-based framework offers three main advantages. First, it translates qualitative dimensions—such as efficiency, reliability, and safety—into quantifiable metrics that enable consistent comparison across scenarios and transport modes. Second, it establishes a common evaluation language across stakeholders and modelling components. Third, it allows the impact of technological innovations (e.g., automation or alternative propulsion) to be explicitly assessed in terms of measurable performance outcomes.

Rather than identifying a single optimal solution, the framework captures trade-offs between competing objectives. For example, environmental improvements may increase capital costs, while higher service frequency can reduce efficiency through lower utilisation. This enables a more nuanced understanding of modal competitiveness.

KPIs are structured into six categories: environmental impact, operational efficiency, reliability, safety, cost-effectiveness, and scalability. Environmental indicators include emissions and energy use (tank-to-wake and well-to-wake). Efficiency metrics capture travel time, speed, load factors, and utilisation. Reliability reflects service consistency and variability, while safety is assessed through operational risk indicators. Cost-effectiveness integrates CAPEX and OPEX, and scalability evaluates system performance under increasing demand. By linking KPIs directly to simulation outputs, the framework ensures that performance assessment is grounded in realistic operational dynamics.

4 Case Studies and Results

This section applies the simulation framework introduced in the previous section to evaluate the potential for modal shift from road transport to IWT using SFAZ vessels. The analysis is conducted across two representative case studies—Ghent and Caen—which reflect distinct operational environments, demand structures, and levels of system maturity.

4.1 Ghent: Urban Construction Logistics

The Ghent case study focuses on construction logistics associated with the renovation of the Royal Opera House, located in the historic city centre. This context provides a particularly relevant testbed for assessing the role of IWT in dense urban environments, where logistics operations are constrained by strict environmental regulations and spatial proximity to navigable canals. The analysis considers the transport of 30,000 tons of bulk materials and demolition waste over the project duration (24 months), characterised by high volumes, repetitive flows, and relatively stable demand patterns.

Two transport configurations are compared. The road-based system relies on conventional diesel trucks with limited payload capacity (12t), high labour dependency, and flexible routing. In contrast, the IWT-based system utilises a single electric automated SFAZ vessel, which offers a substantially higher payload per trip (39t), lower marginal energy costs, and reduced operational labour requirements, though requiring higher capital investment. For both modes, key operational and economic parameters are explicitly considered in the simulation. These include vehicle capacity, energy consumption and associated costs, fleet composition, labour requirements, capital expenditure, and CO_2 emissions.

The baseline analysis contrasts two extreme configurations: a fully road-based logistics chain and a fully waterborne system relying on a single SFAZ vessel. This comparison provides a clear benchmark for evaluating the structural differences between both modes. The results highlight a fundamental trade-off: while IWT routes are slightly shorter due to the geometric efficiency of the canal network, they exhibit significantly longer travel times. This is primarily due to lower cruising speeds and operational constraints such as lock passages. However, these temporal disadvantages are offset by substantial gains in payload efficiency. The SFAZ vessel, with a capacity more than three times that of the trucks, drastically reduces the number of required trips. As a result, the total number of loaded movements decreases from more than 2,500 truck trips to fewer than 800 vessel trips over the project duration. This consolidation effect translates directly into improved system efficiency and reduced pressure on urban infrastructure. From an economic perspective, the SFAZ configuration achieves a cost reduction of approximately 34% compared to the road-based system. This advantage is driven by a combination of higher payload utilisation, lower energy costs, and reduced labour intensity. However, the cost structure differs significantly between modes (Figure 3). Road transport remains labour-dominated, whereas the SFAZ system exhibits a more capital-intensive profile, reflecting the higher upfront investment associated with automated vessels.

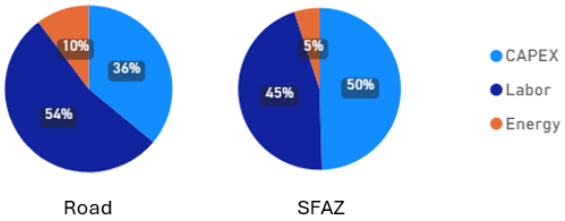


Figure 3: Cost breakdown for road and SFAZ transport modes

Environmental performance shows an even stronger contrast. The road-based system generates 60 tons of CO_2 emissions due to diesel consumption, whereas the fully electric SFAZ operates with zero tank-to-wake emissions. Even when accounting for upstream electricity generation, the overall carbon footprint remains significantly lower.

While the baseline comparison provides useful insights, it represents two limiting cases that rarely occur in practice. To capture more realistic adoption dynamics, a progressive modal shift analysis was conducted. The results reveal a non-linear relationship between SFAZ adoption and system performance (Figure 4). At low adoption levels, the introduction of IWT increases total system cost due to underutilisation of the vessel and the presence of fixed operational overheads. However, as adoption increases, these inefficiencies are progressively mitigated through improved consolidation and higher utilisation rates.

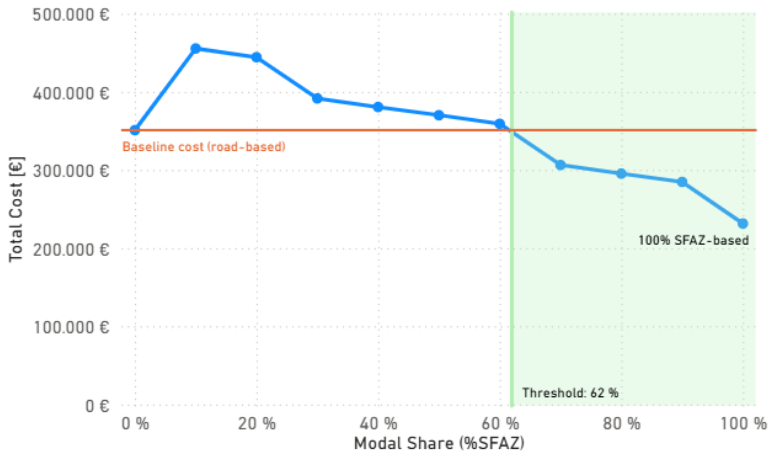


Figure 4: Total transport cost across different SFAZ adoption rates

A critical threshold emerges at approximately 60–65% modal share. Beyond this point, the SFAZ system becomes more cost-effective than the road-only baseline. The central implication is that utilisation, rather than technology alone, is the primary determinant of competitiveness. Without sufficient demand concentration, the economic advantages of inland navigation cannot be fully realised.

4.2 Caen: Urban Distribution under Uncertainty

The Caen case represents a prospective logistics system within the Archipel urban redevelopment project, where the canal network provides a direct connection between the port and the city centre, yet demand patterns and operational conditions remain uncertain. To address this variability, a Monte Carlo simulation framework is adopted, in which key input parameters—including daily transported tonnage, vehicle speeds, handling times, energy prices, and potential carbon taxation—are treated as stochastic variables. Two transport options are evaluated: conventional diesel trucks (12 t capacity) and a catamaran-like SFAZ vessel (7.5 t capacity), allowing assessment of how asset characteristics and operational profiles influence system performance under uncertainty.

Scenario generation reflects realistic operational and market variability. Daily demand is sampled uniformly between 5 and 40 t/day, truck speeds vary between 10–30 km/h to capture congestion effects, while SFAZ vessel speeds range from 8–12 km/h, reflecting expected operational envelopes. Handling times at terminals are drawn from a triangular distribution of 15–45 minutes, representing typical variability with occasional extremes. Energy costs fluctuate according to market conditions, with fuel prices sampled between 0.80–2.80 €/L (normal distribution, 15% standard deviation) and electricity prices between 0.05–0.40 €/kWh (normal, 20% standard deviation). Carbon taxes are modelled discretely at 0, 50, or 100 €/t CO_2, with assigned probabilities reflecting plausible policy pathways.

This approach produces distributions of KPIs—including cost, energy consumption, emissions, and utilisation—rather than single deterministic outcomes, enabling a robust evaluation of the relative competitiveness of road and inland waterway transport under a range of operational, economic, and policy conditions.

The stochastic analysis reveals that road transport remains structurally more cost-efficient on average. However, the cost distributions of trucks and SFAZ vessels exhibit a significant overlap, indicating that inland navigation becomes competitive under a non-negligible subset of conditions (Figure 5). These favourable conditions are primarily associated with high demand density, efficient vessel utilisation, and operational environments where road transport is constrained by congestion or high energy costs. This result reinforces the idea that inland navigation is not universally competitive, but rather context dependent.

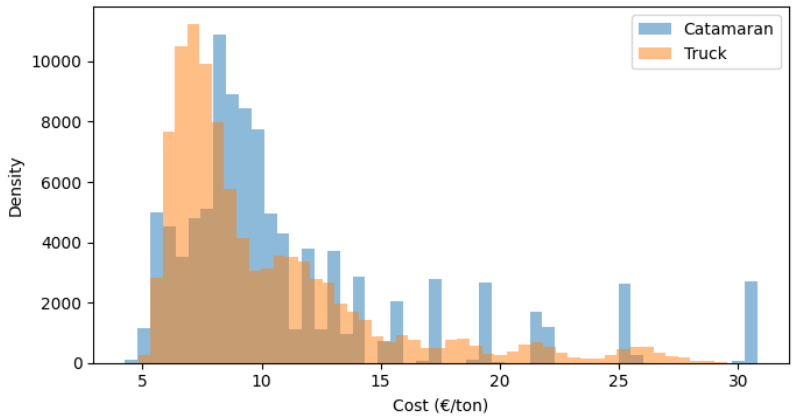


Figure 5: Histogram Comparing Cost per Ton for Truck and SFAZ

The sensitivity analysis identifies three dominant factors shaping modal shift potential. First, demand density emerges as the most influential parameter. The probability of SFAZ competitiveness increases sharply within a specific operational window—approximately between 12 and 30 tons per day (Figure 6, left). Below this range, fixed costs dominate and utilisation remains too low. Above it, the need to deploy additional vessels increases system costs.

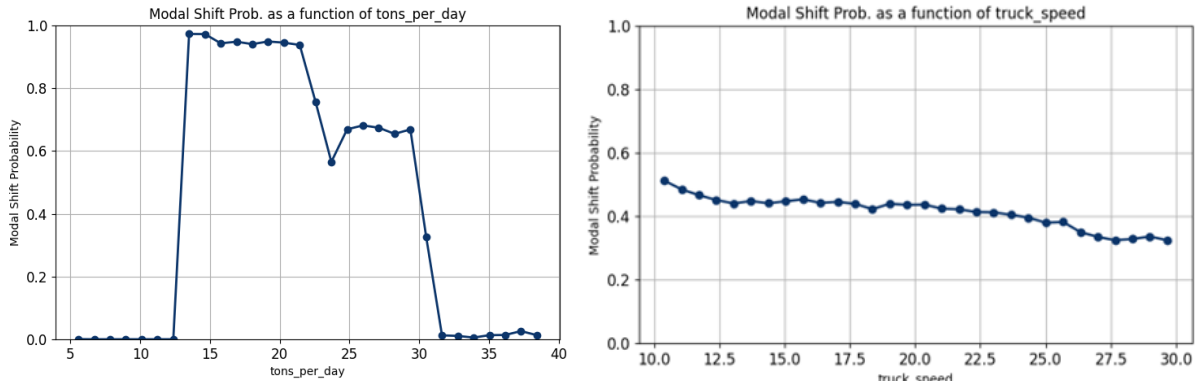


Figure 6: Probability of Shifting to SFAZ as a Function of Demand (Left) and Road Speed (Right)

Second, congestion plays a significant role. Reduced road speeds increase the relative attractiveness of inland navigation (Figure 6, right), highlighting its resilience to surface-level disruptions. Third, energy prices exert a secondary but consistent influence. Higher diesel prices and lower electricity costs both favour the SFAZ, although their impact is less pronounced than that of demand and congestion.

However, the system exhibits strong non-linearities. Small increases in demand or travel time can trigger the need for additional vessels, leading to sharp increases in cost. This sensitivity underscores the importance of careful service design and demand consolidation in real-world implementations.

5 Conclusions

The insights from both case studies are underpinned by the simulation-based methodological framework presented in this work. By integrating multimodal network representations, agent-based modelling, routing optimisation, and Monte Carlo uncertainty analysis, the framework enables a nuanced assessment of transport system performance under realistic operational and stochastic conditions. The explicit consideration of vehicle characteristics, demand variability, terminal handling times, and energy prices allows for the generation of KPI distributions rather than single deterministic outcomes, supporting probabilistic evaluation of modal competitiveness. This methodological approach proves particularly effective in identifying critical thresholds, utilisation patterns, and sensitivity to scale effects, providing robust guidance for both strategic planning and policy assessment.

The Ghent and Caen case studies provide complementary evidence on the conditions under which SFAZ-enabled IWT becomes competitive within urban logistics systems. In structured, high-demand environments such as Ghent, inland navigation outperforms road transport once a critical deployment threshold is reached (~60% modal share), generating both cost savings and substantial emission reductions. In contrast, the Caen case demonstrates that competitiveness in emerging contexts is inherently probabilistic, shaped by demand density, operational constraints, and external factors such as congestion and energy prices. The identification of a bounded operational window further highlights the sensitivity of system performance to scale effects.

Across both cases, a consistent insight emerges: effective capacity use and system integration dominate technological performance as drivers of competitiveness. While SFAZ vessels offer clear advantages in energy efficiency and payload, these benefits materialise only under conditions of sufficient demand concentration and coordinated operations.

This shifts the focus from technology adoption to system design. In particular, demand aggregation, flow synchronisation, and shared asset utilisation become critical to achieving viable load factors. Without such coordination, inland navigation systems risk persistent underperformance.

From a systemic perspective, these findings align closely with Physical Internet principles. The pooling of flows and shared use of capital-intensive assets enable higher efficiency levels while reducing entry barriers for individual actors. In this context, SFAZ vessels function not only as transport units, but as modular components within interconnected logistics networks.

Overall, the results demonstrate that large-scale modal shift towards inland waterways requires a dual transition: technological and organisational. While automation and electrification are necessary enablers, competitiveness ultimately depends on coordination mechanisms that align demand, optimise asset deployment, and support multimodal integration under uncertainty.

Acknowledgements

The work described in this article has been developed in the framework of the European Project "FOREMAST: Freight vOlumes transfer from Road to waterborne transport, using zero-EMission, Automated, Small, and flexible vessel protoTypes". This project has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No 101138261.

References

- Beil, D., Gupta, P., Dopler, S., Pani, A., Putz-Egger, L.-M., & Gupta, A. (2026). Freight modal shift to inland waterways: evaluating transferability of transport policies between Europe and Asia using bayesian best-worst method. Case Studies on Transport Policy, 24, 101777. https://doi.org/10.1016/j.cstp.2026.101777
- Benedecti, R. C., Silva, V. M. D., & Alves da Costa, G. A. (2024). Integrated simulation of an inland container terminal and waterway service for enhancing the maritime supply chain connectivity between Joinville and Itapoá Port. Latin American Transport Studies, 2, 100019. https://doi.org/10.1016/j.latran.2024.100019
- Calderón-Rivera, N., Bartusevičienė, I., & Ballini, F. (2024). Sustainable development of inland waterways transport: a review. Journal of Shipping and Trade, 9(1), 3. https://doi.org/10.1186/s41072-023-00162-9
- Duldner-Borca, B., van Hassel, E., & Putz-Egger, L.-M. (2023). Validation of a Key Performance Indicator Framework Demonstrating Economic Benefits Gained through Resolving Nautical Bottlenecks on Selected Sections of the Danube. Sustainability, 15(24), 16742. https://doi.org/10.3390/su152416742
- Santén, V., Rogerson, S., Williamsson, J., & Woxenius, J. (2021). Modal shift to inland waterway transport. European Journal of Transport and Infrastructure Research, 21(4). https://doi.org/10.18757/ejtir.2021.21.4.5474
- Shobayo, P., Bedoya-Maya, F., van Hassel, E., Vanelslander, T., Christopoulou, E., & Majoor, I. (2024). Assessing Inland Waterway Transport (IWT) container logistics on the Rhine Alpine corridor: A discrete event simulation approach. Research in Transportation Economics, 107, 101475. https://doi.org/10.1016/j.retrec.2024.101475

Maria Matusiewicz, maria.matusiewicz@ug.edu.pl

Urban Ride-Sharing as a Systemic Mobility Solution. Towards a Physical Internet of City Transport

Abstract

Urban transport systems are increasingly approaching the limits of their capacity. Congested streets, a shortage of parking space, high emissions, and inefficient land use are no longer exceptional occurrences but structural features of contemporary metropolitan life. Paradoxically, the very tool that for decades symbolised mobility and individual freedom - the private car - has become one of the primary sources of systemic inefficiency. Statistically, it carries barely one person, consumes a disproportionate share of infrastructure, and generates costs that are borne collectively by the entire city, while still enjoying a privileged status in urban space. With growing digitalisation and pressure to reduce environmental footprints, attention is shifting towards models that do not rely on expanding infrastructure but on smarter utilisation of existing resources. A key direction involves integrating data, technology, and user behaviour within shared forms of mobility. One possible solution is an urban ride-sharing platform operating under principles similar to BlaBlaCar, but designed for short, intra-urban trips. Such a system could create a dynamic network of connections based on real-time travel needs - a kind of "Physical Internet" of passenger transport, where routes, vehicles, and users interconnect in real time. In this approach, mobility ceases to be an individual service and becomes a network of relationships and exchanges, enhancing the city's systemic efficiency. The paper aims to conceptualize an integrated urban ride-sharing model as a foundational layer of the Physical Internet in city transport systems. The proposed approach differs fundamentally from platforms such as Uber or BlaBlaCar, as it is conceived as a conceptual architecture for urban mobility rather than a standalone transport service. The paper additionally provides an illustrative scenario analysis of potential reductions in vehicle kilometres travelled, CO_2 emissions, and external transport costs under different participation levels.

1. Introduction and Literature Review

Existing ride-sharing solutions and Physical Internet (PI) research have largely evolved in parallel. While PI has been extensively explored in freight and logistics, its systematic application to passenger mobility, particularly intra-urban ride-sharing integrated with public transport, remains underdeveloped.

This study contributes with:

- conceptualisation of intra-urban ride-sharing as a PI-inspired mobility layer
- mapping PI principles to passenger mobility

• modular architecture of an urban ride-sharing platform
• illustrative estimation of external cost reductions

In most large cities, cars account for over 50% of daily trips, and the average vehicle carries fewer than 1.3 occupants. This utilisation rate exposes the deep inefficiency of a mobility model built on private transport. Each car requires multiple parking spaces — at home, at work, and in public areas — placing a heavy burden on urban land use. In many metropolitan areas, 30–50% of central space is dedicated to car movement and parking. Private cars also account for a substantial share of urban CO_2 emissions and noise pollution.

Cities around the world are trying in various ways to deal with road congestion. The development of electromobility and low-emission zones represents important steps toward climate neutrality, yet they do not address the core issue — the number of vehicles. Reducing their impact requires not only technological change but a fundamental shift in mobility behaviour and access models. Reducing the availability of parking spaces and increasing parking fees encourage the use of alternative urban transport modes such as public transport, cycling, and walking. Parking charges can contribute to a reduction in vehicle kilometres travelled (VKT) and promote modal shift, particularly when combined with measures that limit road capacity (Rye et al., 2023). An improved public transport offer, together with better cycling and walking infrastructure, provides real alternatives to car use and reduces the motivation to drive. Reviews of the literature confirm the effectiveness of intervention-based approaches aimed at changing travel behaviour and reducing car dependency. Cleland et al. (2023) demonstrate that *comprehensive* measures—combining infrastructure development, pricing policies, and promotional efforts—are the most effective. In a study by O'Hare et al. (2025), three policy scenarios for 2019–2040 were modelled, showing that the implementation of a Low Emission Zone (LEZ) in Brussels could lead to cumulative reductions in greenhouse gas (GHG) emissions of 3.9–5.7 Mt by 2040. This corresponds to the annual GHG emissions of between 83,000 and 121,000 petrol cars in Brussels. The study highlights the potential of the LEZ to accelerate the transition towards cleaner urban transport and to help the city achieve its air quality and climate goals. The research proved that Brussels LEZ could significantly reduce life-cycle GHG emissions from road transport and accelerate the uptake of electric vehicles. The study estimates that the implementation of the LEZ could result in a 50% reduction in annual GHG emissions by 2030. Investing in a robust mobility framework that promotes non-car modes of transport can further accelerate emission reductions. The findings suggest that prioritising non-car transport is effective in achieving faster decarbonisation. If 25% of drivers were to shift from large-segment vehicles to medium-segment ones starting in 2025, the estimated cumulative emission savings by 2040 would be equivalent to eliminating the GHG emissions of around 2,300 petrol cars in Brussels over the same period. Congestion charges for entering city centres typically reduce car traffic and improve air quality, as evidenced by cases such as Stockholm and London. However, these measures initially face significant public resistance (Börjesson, 2018). Road pricing is considered

as an option to manage demand, and reduce emissions. Hosford et al. (2021) made a review of congestion charging zones in selected cities, which is presented in table 1.

Table 1. Comparison of Urban Congestion Charging Systems in Selected Cities.

City	Date introduced	Objectives	Geographical context	Complementary measures	Revenue use	Transport & health impacts
Singapore	Electronic Road Pricing (Sept 1998), replaced Area Licensing Scheme (June 1975)	Reduce congestion and improve travel speeds.	93 charge points around city centre and expressways; Metro population: 6 million; 67% of peak trips by public transit.	Expanded rapid transit and light rail; increased bus service; higher downtown parking fees.	Not specified.	↓ congestion; ↑ shift to public transport.
London	Feb 2003; Western Extension (2007–2011)	Reduce congestion in central London.	Area: 21 km²; Metro population: 9 million; 83% of trips to central area by public transport pre-charge.	Expanded bus services, bus lanes, Low Emission Zone (2008), Ultra Low Emission Zone (2019).	Revenue reinvested in transport: public transit, pedestrian and cycling infrastructure.	↓ congestion; ↓ traffic volume; ↓ air pollution; ↓ crashes; ↑ life expectancy.
Stockholm	Trial in Jan 2006; permanent from Aug 2007	Reduce congestion, improve speeds, lower pollution.	Area: 36 km²; Metro population: 2.2 million; 77% of crossing trips by public transit pre-charge.	Expanded bus and rail capacity; park-and-ride system.	Revenue used for road and transport system improvements.	↓ traffic volume; ↑ public transport use; ↓ air pollution; ↑ life expectancy.
Milan	Area C (Jan 2012), replaced Ecopass (Jan 2008)	Reduce congestion and air pollution.	Area: 8 km²; Metro population: 3 million; 34% of crossing trips by public transport pre-charge.	Improved bus service; bicycle-sharing program.	Revenue reinvested in public transport and cycling initiatives.	↓ traffic volume; ↓ air pollution; ↓ road crashes.
Gothenburg	Jan 2013	Reduce congestion, improve the environment, and fund transport infrastructure.	36 charge points around city centre; Metro population: 1 million; 26% of crossing trips by public transport pre-charge.	More bus lanes; improved cycling infrastructure.	Revenue reinvested in transport improvements, including public transit and new road tunnel.	↓ Traffic volume; ↑ public transport use; ↓ air pollution.

Source: Hosford et al.2021.

The sharing economy is gaining strength under the pressure of resource scarcity. Shared mobility has become a frequent topic in academic research addressing the challenges of urban mobility. The most

common models—such as bikesharing and carsharing—have transformed the way urban communities travel. Shared mobility is a travel strategy that allows users short-term access to transport modes as needed, based on various business models.

According to Cohen and Shaheen (2016), public policy can have a significant impact on the success or failure of shared mobility and other transport innovations by developing equitable public policies, encouraging the use of alternative transport modes, and supporting market competitiveness.

Participants in a study by Habib and Lynn (2020) identified key factors supporting the development of shared mobility as regulatory frameworks for emerging technologies, standardisation in infrastructure design, rethinking spatial planning, and anticipating economic benefits and equity considerations.

Furuhata et al. (2013) provided one of the most comprehensive classifications of ridesharing systems, identifying six main categories ranging from traditional carpooling to emerging dynamic, real-time services. They highlighted persistent barriers (pricing, trust, matching, and institutional design) that limited large-scale adoption.

However, despite technological advances, these systems remained mostly isolated and unintegrated with urban public transport and mobility planning frameworks.

Recent research on dynamic ride-sharing has focused on algorithmic optimization of real-time matching under spatial and temporal constraints. Studies such as Alonso-Mora et al. (2017) demonstrate that large-scale ride-pooling systems can significantly reduce vehicle kilometres travelled through high-capacity matching algorithms. Similarly, Santi et al. (2014) highlight the theoretical potential of ride-sharing based on mobility data, while Agatz et al. (2012) provide a comprehensive overview of optimization approaches in dynamic ride-sharing systems.

However, these studies predominantly focus on improving the efficiency of individual platforms or services. In contrast, the approach proposed in this paper conceptualizes ride-sharing as a systemic mobility layer inspired by the Physical Internet, emphasizing interoperability, modularity, and integration with public transport rather than standalone service optimization.

The concept of the Physical Internet (PI) assumes transferring the idea of an open, interoperable Internet (for data) to physical movement — primarily of goods — through the standardization of loading modules, transit nodes, shared use of space and transport means, and dynamic routing (Treiblmaier et al., 2020).

Urban logistics and the Physical Internet are two key concepts aimed at transforming logistics toward sustainability — improving economic, environmental, and social efficiency. Urban logistics represents the final and most challenging segment of logistics and transport networks; it is also the most expensive part of the supply chain due to limited spatial resources combined with increasing consumption and growing demand for logistics services (Crainic & Montreuil, 2016).

For this reason, the Physical Internet can be particularly applicable in urban spaces, enabling the sharing of fleet and distribution network resources within limited city infrastructure. Urban logistics involves not

only the management of flows of goods, services, and information, but also of people (Taniguchi, 2014; Szołtysek, 2015; Farchi et al., 2021). Therefore, PI should also be leveraged to enhance the efficiency of these flows. This paper builds on those insights, extending the discussion toward systemic, city-level integration of coordinated ride-sharing as a sustainable mobility layer. The paper adopts a conceptual research design, aiming to integrate insights from urban logistics, shared mobility, and the Physical Internet into a coherent framework for systemic urban transport efficiency. The study focuses on developing a theoretical foundation and identifying potential areas for empirical validation.

2. The Conceptual Framework

Modern digital solutions now allow real-time integration of route, driver, and passenger data. On this basis, it becomes possible to design an urban ride-sharing application inspired by the logic of BlaBlaCar but tailored to short, intra-city distances. This paper explores the possible effects of implementing an app like this.

Unlike ride-hailing services such as Uber or Bolt, drivers in this model are not professional carriers. They simply share empty seats with passengers travelling in the same direction, splitting trip costs. This system would create a flexible mobility layer that complements, rather than competes with, public transport. Conceptually, it can be understood as a "Physical Internet of Transport" — a network in which trips, drivers, and passengers are dynamically matched, analogous to the flow of data in computer networks. Such an approach enables cities to use existing resources — vehicles, routes, and energy — more effectively, flexibly, and on demand, much like digital systems operate.

Unlike inter-city carpooling platforms such as BlaBlaCar, the proposed system prioritises short-distance, high-frequency trips, partial route overlaps, and explicit complementarity with public transport.

The key differences compared to BlaBlaCar are outlined below:

1) Intra-urban context

• Trips take place entirely within the city.

• Travel distances are short.

• Travel patterns exhibit high repeatability (e.g. daily commuting).

2) Short trips

• Typical distances range from 5 to 15 km.

• Low individual costs and limited acceptable detours.

• Higher sensitivity to delays compared to long-distance trips.

Table 2. Novelty of the proposed concept compared to existing models

Feature	Ride-hailing (e.g. Uber)	Inter-city carpooling (e.g. BlaBlaCar)	Proposed system
Driver type	Professional	Non-professional	Non-professional
Trip type	On-demand, point-to-	Long-distance	Short intra-urban

	point		
Matching logic	Direct trip matching	Route-based	Partial route matching
System role	Transport service	Platform	System-level mobility layer
Integration with public transport	Limited	None	Core feature
Pricing model	Market-based	Cost-sharing	Cost-sharing + public incentives

3) Partial route matching

• Passengers do not necessarily travel from door to door.

• Users may join or leave a shared trip at intermediate points along the route.

4) First-/last-mile function

• The system is not designed to replace public transport.

• Instead, it can complement public transport by extending its reach.

• A shared ride may therefore serve as a first- or last-mile connection to a public transport hub rather than as a full end-to-end journey.

5) Public transport as the backbone

• The core objective is to merge multiple individual trips (one vehicle-one traveller) into a smaller number of vehicles carrying more passengers.

• The proposed system is not intended to compete with public transport; rather, it aims to alleviate congestion and enhance the overall efficiency of the urban mobility system.

Ride matching in an intra-urban context is subject to stricter constraints than in inter-city carpooling. Matches must respect narrow time windows, limited acceptable detours, and directional compatibility rather than full route coincidence. Importantly, matching may prioritise proximity to public transport nodes, enabling shared rides to function as first- or last-mile connections rather than complete door-to-door services.

The objective of the proposed platform is to create a data-driven, adaptive network of shared trips that supplements public transport and reduces the need for individual car travel. Figure 1 presents the operating scheme of the app.

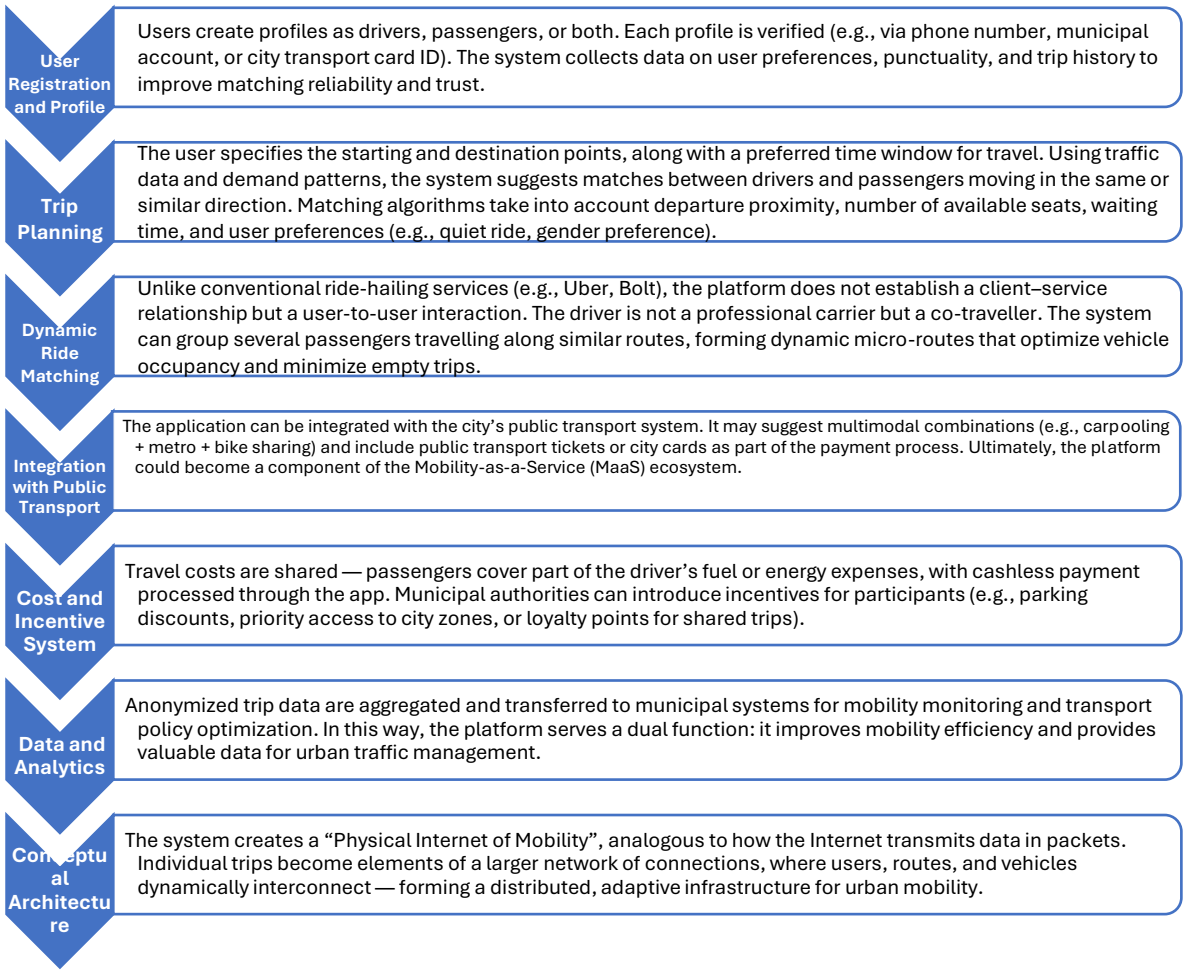


Figure 1. The app operating scheme.

The platform operates through a modular architecture in which user profiles, trip requests, and real-time location data interact dynamically. Trip-matching outputs are continuously updated based on incoming user requests and public transport data, while feedback and ratings feed back into the user trust and reputation module.

The application could be named SharePath, CivRide, FlowCity, CoRide, Let'sRide, OneDirection. The implementation of an urban carpooling system could generate a range of social, environmental, and economic benefits, presented in table 3.

Table 3. Potential Benefits.

Benefit	Description
Reduced traffic volumes	Fewer private vehicles on the road translate into smoother traffic flow and shorter travel times.
More efficient use of urban space	Reduced parking demand allows cities to reclaim space for pedestrians, greenery, and public functions.
Environmental benefits	Fewer trips mean lower emissions and decreased noise pollution.
Lower mobility costs	Shared rides reduce individual transport expenses.
Social impact	Shared trips foster interpersonal interaction and help build social capital.
Support for Sustainable Development Goals (SDG 11)	Promoting intelligent, accessible, and low-emission cities.

Together, these effects demonstrate that local-scale carpooling initiatives can contribute not only to transport efficiency but also to broader goals of social inclusion and urban sustainability. The Physical Internet is not used here as a literal transfer of logistics solutions, but as a design paradigm for organising passenger mobility in dense urban systems. Table 4 presents what PI principles are to be adopted in the proposed concept.

Table 4. Mapping PI principles to urban ride-sharing

PI principle	Application in proposed system
Modularity	Trips treated as modular units
Encapsulation	Each ride as a standardised mobility package
Open protocols	Data exchange with PT systems
Interoperability	Integration with MaaS platforms

The next section outlines a simplified estimation of potential benefits, focusing on the reduction of external transport costs.

3. The Benefits

To illustrate potential impacts, the study employs a simplified estimation of external cost reductions (CO_2 and noise) using standard indicators from transport economics and public health literature (social cost of carbon, SCC; DALY); Murray, Lopez, 2013). The assumed 30% adoption rate does not represent a forecast, but an illustrative scenario, broadly consistent with participation levels observed in mature shared-mobility systems. Lower participation rates would proportionally reduce estimated benefits but would not change their direction. The SCC value of €100/t CO_2 follows recent literature (Rennert et al., 2022) and is used for illustrative comparison rather than precise valuation. These serve as illustrative rather than empirical calculations, supporting the conceptual argument.

Two approaches may be used to estimate environmental and social externalities associated with urban transport.

1) Common Measure (CBA): how to "monetize" CO_2, noise and well-being

CO_2

The reduction in vehicle kilometres travelled (VKT) is estimated as:

$$\Delta VKT = p \cdot N_commute \cdot d_RT \cdot D \cdot (1 - 1/occ)$$

where:
p – participation rate
N_commute – number of daily commuters
d_RT – average round-trip distance (km)
D – number of commuting days per year
occ – average vehicle occupancy after ride-sharing

The term $(1 - 1/occ)$ reflects the reduction in the number of vehicles due to increased occupancy.

CO_2 emissions reduction is calculated as:

$$\Delta CO_2 = \Delta VKT \cdot FE$$

where FE is the fleet emission factor (tCO_2/km).

The monetary benefit is estimated using the social cost of carbon (SCC):

$$Benefit_CO_2 = \Delta CO_2 \cdot SCC$$

Noise

We assume the traffic decrease (20%)
The decrease of sound

$$\Delta L(db) \approx 10log_{10}(1 - x)$$
$$\Delta L(db) \approx 10log_{10}(1 - 0.2)$$
$$\Delta L(db) \approx -0.97\ dB$$

The number of people exposed to the noise on a given road part: N_{exp}
Monetization of noise reduction based

a) on willingness to pay (WTP)

$$Benefit\ noise\ reduction(€) = \sum_{road} N_{exp} \cdot |\Delta L(db)| \cdot WTP_{dB\ pers/year}$$

Scenario analysis

To illustrate the sensitivity of results to different participation levels, three scenarios are considered: low (10%), moderate (30%), and high (50%) adoption rates.

Assumptions remain consistent with the base case, with variation only in participation rate (p).

Table 5. Scenario analysis

Scenario	Participation rate (p)	ΔVKT (km/year)	CO_2 reduction (t/year)	CO_2 benefit (€ / year)
Low	0.1	27,500,000	3,300	330,000
Moderate	0.3	82,500,000	9,900	990,000
High	0.5	137,500,000	16,500	1,650,000

The results indicate a near-linear relationship between participation rate and system-level benefits, suggesting that even moderate adoption can generate substantial reductions in vehicle kilometres travelled and associated external costs.

b)or on national health costs using DALY

DALY – Disability-Adjusted Life Year (or: *lost healthy life year*).

It is a measure of the total burden of disease, used, for example, by the WHO, OECD, EEA and in analyses of the costs of noise, pollution, accidents.

1 DALY = 1 year of healthy life lost.

This may be due to premature death (mortality), or living with illness, pain, stress, loss of fitness (morbidity). DALY therefore combines:

DALY=YLL+YLD, where:

YLL (Years of Life Lost) – years of life lost due to premature death,

YLD (Years Lived with Disability) – years lived with reduced quality of life.

Exposure to noise causes stress, sleep disorders, heart disease, depression, etc. On this basis, the WHO calculates how much DALY the society loses annually due to noise (e.g. traffic noise). These lost DALYs can then be *monetized*:

Each such exposure - DALY reduction by function:

$$DALY_{effect} = N_{exp} \cdot L_{den} \cdot W_{effect}$$

$$N_{exp} - the\ number\ of\ exposed\ individuals$$

$$L_{den} - proportion\ affected\ population\ at\ noise\ level\ L_{den}$$

$$W_{effect} - health - weighting\ factor\ (e.g.\ 0.01\ DALY\ per\ person\ per\ year)$$

2) Multi-criteria scoring, without "monetization"

Multi-Criteria Scoring. It is a way to compare different options (e.g. transport scenarios, urban policies, emission reduction models) without having to convert everything into money (€). Instead of a monetary valuation, each criterion is given a score on a scale of 0-100 points, and then everything is weighted and summed up.

Objective - to evaluate and compare scenarios (e.g. transport policies) in terms of:

1. CO_2 emissions,
2. Noise
3. Public costs.

CO2

0 – no reduction, 100 – best scenario

$$S_{co2} = 100 \cdot \frac{\Delta\, CO_2}{\max s\, \Delta\, CO_2}$$

Noise

$$S_{noise} = 100 \cdot \frac{\sum_{road} N_{exp} \cdot |\Delta L(db)|}{\max s\ \sum_{road} N_{exp} \cdot |\Delta L(db)|}$$

Public costs = $weight\ S_{co2} + weight\ S_{noise}$

Data:

1. Traffic

$N_{commute}$, share of drivers, distance length d_RT, number of days D, current occupancy

FE

2. Noise

Reduction of congestion/per route (x)

Exposition map (Nexp) – how many people is exposed

WTPdb/person/year

4. Example

Assumptions:

ΔVKT = p · N_commute · d_RT · D · (1 − 1/occ)

= 0.3 · 220000 · 10 · 250 · (1 − 1/2)

= 82,500,000 km/year

p = 30% users (0.3), Ncommute = 220 000, drt (round-trip distance per day) = 10km, D=250 working days/year;

FE = 0.00012t/km; SCC = 100€/t CO2; 0.5 = two occupants per vehicle, instead of one.

Benefit_CO_2 = ΔVKT · FE · SCC

$Benefit\ CO_2$(€) = $82500000 km/year \cdot 0.00012 t/km \cdot 100$ = 990 000€/year

Traffic -20% per 100km roads, let's assume Nexp = 120000 ppl,

$\Delta L(db) \approx -0.97\ dB$

$If\quad WTPdB\ pers/year = 10$ €

$Benefit\ noise\ reduction$(€) = $120000 \cdot 0.97 \cdot 10$ = 1 164 000€/$year$

Of course, these potential benefits also imply a number of challenges and preconditions for successful implementation. The effective introduction of an urban ride-sharing platform requires addressing several

key issues, presented in table 6. Addressing these factors determines whether carpooling can evolve from a niche experiment into a stable component of the city's mobility ecosystem. Table 6 presents possible research directions, and the areas highlighted in bold indicate topics addressed by the author in another paper, which is currently under review.

Table 6. Potential challenges.

Challenge	Description
User trust and safety	The system must include mechanisms for user verification, ratings, and data protection.
Legal aspects	Clarification of the legal status of shared rides within transport and tax regulations is essential.
Integration with public transport	The platform should interoperate with city trip-planning systems, enabling multimodal routing and ticketing.
Institutional incentives	Local authorities can promote participation through parking discounts, access to bus lanes, or dedicated parking zones for shared vehicles.
Technological aspects	The development of real-time matching algorithms and mobility-pattern analytics is crucial to ensure system efficiency and scalability.

Table 7. Proposed research directions and areas covered in a related paper (in bold).

Research direction	Description
User behavior and participation motivations	• Studies on socio-behavioral acceptance; • Motivation models, the role of economic (savings), social (trust, community), and environmental (pro-ecological identity) factors. **• to what extent do incentives (e.g., access to bus lanes, parking discounts) increase participation?** **• how to balance between voluntariness and policies restricting individual car use?**
Systemic impact on the urban environment and economy	**• Analysis of CO_2 and noise reduction effects - simulations based on transport data and traffic models (e.g., SUMO, matsim);** **• Impact on the external costs of transport** • Life Cycle Assessment (LCA) - comparison of emissions from traditional individual and shared transport; • Implementation scenarios - would 10%, 30%, or 50% adoption of shared rides change the city's traffic structure and energy balance?
Integration with the public transport system	• Complementarity analysis - does carpooling act as a complement or a competitor to public transport? • Optimization of multimodal routes - how to combine carpooling with bike-sharing, micromobility, and public transit? • Modeling maas (Mobility-as-a-Service) networks - where should such a platform be positioned within the urban ecosystem?
Algorithms and data infrastructure	• analysis of efficiency, waiting times, and route optimization; • Data privacy and ethics - how to ensure transparency and security of sensitive information; • Integration with municipal open data - cooperation with transport operators, integration with traffic and weather data.

Economic models and regulatory frameworks	• Sharing economy vs. public transport - does carpooling require subsidies, or can it be financially self-sufficient? • Legal regulations - the status of shared rides (private individual vs. Carrier), liability, and insurance; • Urban policy scenarios - how cities can actively support system adoption (e.g., within SDG 11 and low-emission policies).
Social and spatial effects	• **Equity, accessibility** • **Spatial effects, land use** • Social trust, community effects
Scalability and local adaptation	• City-level variation • Pilot studies, real tests • EU integration and policy link

5. Discussion

The proposed concept of a city-level ride-sharing platform should be viewed not only as a technological solution but as a potential structural shift in how urban mobility is organised. The results of the simplified estimations suggest that even partial adoption of such a system could significantly reduce external transport costs related to CO_2 emissions and noise. This paper deliberately limits itself to a conceptual framework and simplified estimation. Future work will include agent-based simulations to test behavioural responses, detours, and rebound effects under different adoption scenarios. However, the feasibility of this transformation depends on several interrelated factors.

First, behavioral acceptance remains a critical barrier. Even if the technological infrastructure allows real-time matching, the willingness of users to share rides with strangers is culturally dependent and influenced by trust, privacy concerns, and perceptions of safety. Experience from existing carpooling platforms shows that strong community mechanisms (ratings, mutual verification) are essential to reach critical mass.

Second, while the system could reduce car ownership and traffic volumes, its interaction with public transport must be carefully managed. If not properly integrated, ride-sharing could attract users away from buses and trams rather than from private cars, undermining sustainability goals. A coordinated mobility ecosystem, supported by data exchange between the platform and municipal transport operators, is therefore necessary. This effect may be interpreted as a form of rebound or modal shift risk, where efficiency gains in one mode lead to unintended increases in overall mobility demand or substitution away from more sustainable modes. To mitigate this risk, the proposed system is explicitly designed as a complementary layer within a Mobility-as-a-Service (MaaS) ecosystem. This includes prioritizing first- and last-mile connections, aligning pricing mechanisms with public transport, and integrating routing logic that favors multimodal journeys over direct substitution. Such design principles aim to ensure that ride-sharing reduces private car dependency rather than competes with high-capacity public transport systems.

Third, the regulatory and institutional framework must evolve to support such hybrid models. Legal definitions of shared rides, liability, taxation, and insurance need clarification to distinguish non-commercial

cost-sharing from professional transport services. Municipal authorities could play a catalytic role by offering incentives such as dedicated lanes or parking privileges.

Fourth, data governance and ethics emerge as a decisive challenge. Real-time mobility data have high social value but also pose privacy and surveillance risks. Transparent data-sharing standards and user control mechanisms will be key to ensuring legitimacy and public trust.

Public authorities are envisioned not only as regulators, but as platform partners enabling data exchange and incentive schemes. Importantly, the proposed system is designed to scale incrementally through pilot implementations rather than full-city deployment.

From an economic perspective, the proposed system operates on a hybrid cost-sharing and incentive-based logic. Unlike ride-hailing platforms, where pricing reflects market-driven service provision, this model is based on partial cost recovery among users, reducing marginal travel costs without generating full commercial profit for drivers. The system's efficiency relies on achieving a critical mass of users, where increased participation leads to higher matching probability and improved vehicle occupancy.

Public authorities may play a catalytic role by introducing targeted incentives, such as parking benefits or access privileges, to support early adoption and correct market imbalances. At the same time, the system does not require full public subsidization, as its core mechanism is based on voluntary participation and shared cost structures. This hybrid model positions urban ride-sharing between private service provision and public transport, requiring coordinated governance rather than purely market-based regulation.

In summary, the proposed model holds potential as a complementary mobility layer that could make cities more efficient, equitable, and sustainable; however, its success will require not only technological readiness but also social acceptance, institutional alignment, and a coherent urban policy framework.

6. Conclusions

The concept presented in this paper illustrates that the transformation of urban mobility is not merely a question of technology, but of system design and collective coordination. A ride-sharing platform functioning as part of a "Physical Internet" of transport would require a rethinking of how cities organise access, ownership, and responsibility for mobility. Its successful deployment could mark a transition from fragmented, car-centric systems to networked, demand-responsive mobility ecosystems. Yet, this shift will only materialise if supported by coherent policy frameworks, transparent governance, and public engagement.

Shared mobility represents one of the most promising directions for the development of urban transportation. Instead of investing in ever-expanding road infrastructure, cities can improve mobility efficiency through intelligent management of existing resources.

The proposed concept of a city ride-sharing application is not merely a technological addition to the transport system, but rather a new way of organizing it - more flexible, economical, and aligned with the

principles of sustainable development. In the long term, it may contribute to a paradigm shift in urban mobility: from ownership-based to sharing-based, from individual to interdependent. This reflects the very essence of both the Physical Internet concept - a network of shared connections, resources, and interactions that enhances collective efficiency. Understanding and developing this synergy between physical mobility and digital connectivity should therefore become a key area of further exploration. Future research should focus on several interrelated aspects. It is essential to examine user behavior and participation motivations, including social, economic, and environmental factors that influence engagement in shared mobility systems. Equally important is the study of the broader systemic impacts - from reductions in CO_2 emissions and traffic congestion to potential effects on the urban economy and spatial organization (an initial attempt at such calculations has already been made by the author and is currently under review). Further attention should be given to how such systems can be effectively integrated with public transport networks and other mobility services, forming coherent multimodal ecosystems. Additionally, the development of algorithms and data infrastructures that ensure efficiency, transparency, and privacy will be crucial. Economic and legal frameworks must also be refined to balance innovation with fairness and safety. Finally, future studies should explore scalability and adaptation to local contexts, assessing how different cities and regions can best implement and support shared mobility within their sustainable development strategies.

References

1. Agatz, N., Erera, A., Savelsbergh, M., & Wang, X. (2012). Optimization for dynamic ride-sharing: A review. European Journal of Operational Research, 223(2), 295–303. https://doi.org/10.1016/j.ejor.2012.05.028

2. Alonso-Mora, J., Samaranayake, S., Wallar, A., Frazzoli, E., & Rus, D. (2017). On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. Proceedings of the National Academy of Sciences, 114(3), 462–467. https://doi.org/10.1073/pnas.1611675114.

3. Börjesson, M. (2018) Long-term effects of the Swedish congestion charges. Discussion paper. Swedish National Road and Transport Research Institute; OECD International Transport Forum.

4. Cleland, C. L., Jones, S., Moeinaddini, M., Weir, H., Kee, F., Barry, J., Longo, A., McKeown, G., Garcia, L., & Hunter, R. F. (2023). Complex interventions to reduce car use and change travel behaviour: An umbrella review. Journal of Transport & Health, 31, 101652. https://doi.org/10.1016/j.jth.2023.101652.

5. Cohen, A., Shaheen, S. A. (2016). Planning for shared mobility. University of California, Berkeley, Transportation Sustainability Research Center. Retrieved from https://www.researchgate.net/publication/310504165.

6. Crainic, T. G., Montreuil, B. (2016). Physical Internet Enabled Hyperconnected City Logistics. Transportation Research Procedia. 12. 383-398. 10.1016/j.trpro.2016.02.074.

7. Furuhata, M., Dessouky, M., Ordóñez, F., Brunet, M. E., Wang, X., & Koenig, S. (2013). Ridesharing: The state-of-the-art and future directions. Transportation Research Part B: Methodological, 57, 28–46. https://doi.org/10.1016/j.trb.2013.08.012

8. Habib, M. A., & Lynn, R. (2020). Planning for connected, autonomous and shared mobility: A synopsis of practitioners' perspectives. Procedia Computer Science, 170, 419–426. https://doi.org/10.1016/j.procs.2020.03.084

9. Hosford, K., Firth, C., Brauer, M., & Winters, M. (2021). The effects of road pricing on transportation and health equity: A scoping review. Transport reviews, 41, 766–787. https://doi.org/10.1080/01441647.2021.1898488.

10. Murray, C. J. L., & Lopez, A. D. (2013). Measuring the global burden of disease. The New England Journal of Medicine, 369(5), 448–457. https://doi.org/10.1056/NEJMra1201534.

11. O'Hare, K., Lee, K., & Bernard, Y. (2025). Life-cycle greenhouse gas impact of the low-emission zone in Brussels. The Real Urban Emissions Initiative.

12. Rennert, K., Errickson, F., Prest, B.C. et al. (2022) Comprehensive evidence implies a higher social cost of CO2. Nature 610, 687–692, https://doi.org/10.1038/s41586-022-05224-9.

13. Rye, T., Burns, J., Whitefield, A., Oxley-Glenister, K., Pinkard, J. (2023). Reducing car use through parking policies: An evidence review. Ansons Consulting Ltd. https://doi.org/10.7488/era/3776.

14. Santi, P., Resta, G., Szell, M., Sobolevsky, S., Strogatz, S. H., & Ratti, C. (2014). Quantifying the benefits of vehicle pooling with shareability networks. Proceedings of the National Academy of Sciences, 111(37), 13290–13294. https://doi.org/10.1073/pnas.1403657111

15. Szołtysek, J., Logistyczne aspekty zarządzania przepływami osób i ładunków w miastach, AE Katowice 2005, p. 105.

16. Taniguchi, E. (2014). Concepts of city logistics for sustainable and liveable cities. Procedia – Social and Behavioral Sciences, 151, 310–317. https://doi.org/10.1016/j.sbspro.2014.10.029.

17. Treiblmaier H., Mirkovski K., Lowry P.B., Zacharia, Z.G. (2020), "The physical internet as a new supply chain paradigm: a systematic literature review and a comprehensive framework". The International Journal of Logistics Management, Vol. 31 No. 2 pp. 239–287, doi: https://doi.org/10.1108/IJLM-11-2018-0284.

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE
AUTOSUP

Implementing the Physical Internet in Multimodal Hubs: The AUTOSUP Digital Twin-Based Decision-Support Approach

Aristea M. Zafeiropoulou[1], Konstantinos Christidis[2], Filippos Adamidis[3], Athanasios Karydis[4], Babis Magoutas[2], Antonis Mygiakis[1], Ilias Gkotsis[4]

1. Konnecta Systems, Athens, Greece
2. Frontier Innovations, Athens, Greece
3. Technical University of Munich, Munich, Germany
4. Inlecom Innovation, Athens, Greece

Corresponding author: aristea.zafeiropoulou@konnecta.io

Abstract: *The Physical Internet (PI) envisions logistics networks as open, modular, and interoperable systems enabling seamless multimodal freight flows. However, its operationalisation in real logistics hubs remains limited due to fragmented data environments, heterogeneous automation levels, and lack of integrated decision-support tools. This paper presents the AUTOSUP approach, which implements PI principles through a Digital Twin (DT)-based Decision Support System (DSS) applied in multimodal Living Hubs (L-Hubs). The framework integrates Information Flow modelling and Knowledge Graphs to capture interdependencies across assets, processes, and stakeholders, enabling consistent and explainable scenario evaluation. A PI readiness model is introduced, structuring automation adoption across three dimensions: automation, connectivity, and collaboration. The DSS supports the assessment of alternative operational and investment strategies through simulation and cost–benefit analysis. Early structured evidence from two European L-hubs shows measurable improvements in operational efficiency, resource utilisation, and environmental performance for one quantified use case per hub. The results highlight that coordinated information exchange and system-level integration are critical to scaling automation and enabling synchromodal, PI-compatible logistics operations.*

Keywords: *Physical Internet, Digital Twins, Decision Support System, Automation, Logistics Hubs, Information Flows.*

Physical Internet (PI) Roadmap Fitness*: ☒ PI Nodes (Customer Interfaces, Logistics Hubs, Deployment Centres, Factories), ☐ Transportation Equipment, ☒ PI Networks, ☒ System of Logistics Networks, ☐ Vertical Supply Consolidation, ☐ Horizontal Supply Chain Alignment, ☒ Logistics/Commercial Data Platform, ☒ Access and Adoption, ☒ Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

1 Introduction

The Physical Internet (PI) frames freight transport and logistics as an open, interoperable, and intelligence-enabled system in which goods move through modular networks of connected nodes and services (Montreuil, 2011; ALICE, 2022). In this perspective, logistics hubs are operational nodes where physical handling, information exchange, and coordination decisions shape the performance of wider multimodal transport chains. Despite the maturity of the PI concept, its operationalisation in real logistics environments remains limited.

A key reason is that automation and digitalisation are often introduced as local technical interventions rather than as elements of a coordinated network logic. Ports, terminals, rail yards, and inland logistics hubs increasingly deploy automated equipment, digital platforms, tracking systems, and optimisation tools, but these solutions often remain weakly connected across organisational and modal boundaries (Heilig et al., 2017; Villarreal et al., 2021). As a result, the PI objectives of interoperability, modularity, and network-level coordination are only partially realised.

Digital Twins (DTs) and Decision Support Systems (DSSs) offer a promising foundation for bridging this gap. DTs replicate assets, processes, and operational interactions and support what-if analysis prior to real-world deployment (Barricelli et al., 2019; Raptis et al., 2022). DSSs leverage these models to assess automation pathways, investment priorities, and coordination strategies under uncertainty (Lindgren et al., 2020). Existing work, however, has focused mainly on isolated hub simulation, local optimisation, or technology-specific control functions. Less emphasis has been placed on integrating Information Flows, interdependencies, and economic feasibility within a PI-aligned decision-support environment. The relation to synchromodality is also less developed in practical hub-oriented automation studies, even though dynamic coordination across modal options is closely linked to PI-compatible logistics operations (Labarthe et al., 2023).

This paper addresses these gaps through the AUTOSUP approach, which implements PI principles in multimodal Living Hubs (L-Hubs), through an integrated DT-based DSS. The approach combines Information Flow (IF) modelling, Knowledge Graph (KG) based representation of interdependencies, simulation-based DT models, and cost-benefit analysis to support informed and explainable automation decisions. The paper reports the conceptual design of this framework along with its implementation using a DSS architecture and early validation evidence from two European L-Hubs.

The paper addresses four questions: how PI principles can be operationalised in multimodal hubs through a DT-based decision-support approach; how Information Flows and interdependencies can be modelled to support interoperability and explainable automation decisions; what measurable operational, economic, and environmental benefits can be observed in current L-Hub validation; and to what extent the approach can facilitate synchromodal decision-making through enhanced coordination across modes and nodes.

Thus, the contribution of the paper is fourfold. First, it conceptualises multimodal hubs as PI-oriented decision nodes within a system of logistics networks. Second, it proposes an integrated modelling approach that couples IF, KGs, DTs, and DSS functionalities. Third, it applies a PI readiness perspective structured around automation, connectivity, and collaboration. Fourth, it reports validation insights from two L-Hubs, illustrating how the framework supports automation pathway assessment, bottleneck identification, and investment prioritisation.

2 Background and Related Work

The PI frames logistics as an open and interoperable system in which goods move through coordinated networks of nodes, services, and interfaces (Montreuil, 2011; Meller et al., 2012; ALICE, 2022). In this perspective, logistics hubs are central as they connect transport modes, infrastructures, and actors, and because local decisions at hub level affect the continuity and efficiency of wider freight systems. While the PI literature has established a strong conceptual basis around modularity, collaboration, and interoperability, fewer studies show how these principles can be translated into practical decision-support mechanisms in real multimodal hubs.

DTs and DSSs have emerged as important tools in logistics and transport. DTs replicate assets and processes to facilitate simulation, forecasting, and operational analysis (Barricelli et al., 2019; Raptis et al., 2022). DSSs support planning and evaluation under uncertainty with applications in fleet management, terminal scheduling, and investment analysis (Lindgren et al., 2020). However, existing work has often focused on local optimisation, isolated assets, or single modes of transport. Information exchange is typically treated as a background technical condition rather than as a modelled part of the decision process, and economic feasibility is often separated from operational simulation.

These limitations are directly relevant to multimodality and synchromodality. Dynamic coordination across modes requires visibility into both physical and informational states of the system, together with a way to compare alternative operational configurations under changing conditions (Labarthe et al., 2023). Existing approaches rarely combine explicit IF modelling, graph-based representation of interdependencies, microscopic multimodal simulation, integrated cost-benefit evaluation, and validation in real multimodal hubs. AUTOSUP addresses this gap through a PI-oriented DSS environment that links these elements within a common decision-support framework.

The originality of AUTOSUP lies in the integration of four elements seldom combined in hub-oriented studies. First, Information Flows are modelled as operational structures that affect coordination quality, rather than passive data channels. Second, a Knowledge Graph explicitly represents interdependencies and enables traceable scenario logic. Third, microscopic Digital Twins are used to compare AS-IS and TO-BE configurations in real multimodal hubs. Fourth, the simulation outputs are linked directly to cost-benefit and feasibility assessments. This synthesis enables the analysis of PI mechanisms in operational terms, through interface conditions, routing and handling configurations, and coordination choices across connected nodes.

3 AUTOSUP conceptual framework

3.1 Multimodal hubs as PI nodes and PI readiness

AUTOSUP conceptualises multimodal hubs as PI nodes within a system of logistics networks. In this view, ports, terminals, yards, and connected inland facilities are places where physical handling, information exchange, and coordination decisions jointly affect the performance of wider freight systems. Automation is therefore treated as a networked capability rather than as a local technical feature (Montreuil, 2011; ALICE, 2022).

The framework interprets PI readiness through three interdependent dimensions. The first pertains automation at hub level, including automated equipment, process automation functions, and local decision-support capabilities. The second concerns connectivity and information sharing, including the ability of actors and systems to exchange operational information that supports coordinated action. The third relates to collaboration, including the capacity to align decisions across terminals, networks, and modes. These dimensions are treated as complementary, since local automation gains are constrained when connectivity or collaboration conditions are weak. In operational terms, PI compatibility is expressed here through modular handling and routing options, standardised information exchanges between actors, open interfaces between modelling layers, and scenario-based decisions that account for network-level coordination.

3.2 Information flows and interdependencies in PI-oriented automation

The AUTOSUP framework places Information Flows at the core of automation analysis. In multimodal hubs, operational performance depends on the timely exchange of information concerning schedules, asset availability, disruptions, slot allocations, readiness states, and handling conditions. These flows are modelled explicitly, as they shape the feasibility and performance of automation pathways.

IF modelling identifies the entities involved in data exchange, the operational interdependencies among them, and the digital interactions required to support coordinated action. This enables the detection of bottlenecks, asymmetries, and vulnerabilities that are not visible through physical process modelling alone. In methodological terms, the IF layer provides the link between operational reality and the decision logic of the DSS (AUTOSUP, D2.3, 2025).

To support this representation, AUTOSUP employs a Knowledge Graph (KG), i.e., an Interdependency Graph, acting as the interdependency backbone of the DSS. The graph captures infrastructure assets, physical and information links, and logical dependencies between operations and actors. Its development follows a structured workflow (Figure 1) in which infrastructure nodes and physical connections are identified from the use cases, IF diagrams are captured as a separate logical layer, and a reproducible Neo4j graph is instantiated automatically from shared datasets (AUTOSUP, D2.3, 2025). This enables a traceable and updatable representation of how local changes propagate through the multimodal system.

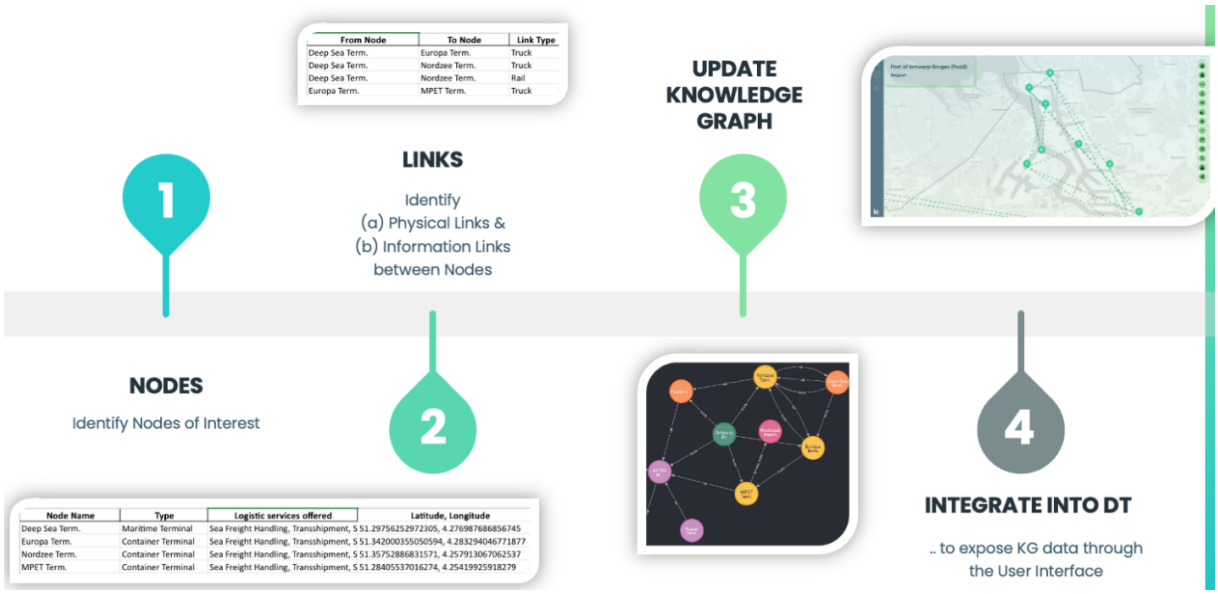


Figure 1: Interdependency Graph Pipeline

3.3 From conceptual PI principles to decision-support logic

AUTOSUP translates PI principles into a decision-support logic for real multimodal hubs. Modularity is represented through alternative process, routing, and handling configurations. Interoperability is represented through IF models and graph-based interfaces between assets, systems, and actors. Collaboration is represented through scenario evaluation that accounts for governance conditions, coordination requirements, and data-sharing responsibilities.

These dimensions are brought together in a DSS that combines the KG, DT models, IFs, optimisation components, and cost-benefit analysis in one analytical environment (AUTOSUP, D2.3, 2025). The framework therefore provides both a PI-oriented interpretation of hub maturity and an operational pathway for evaluating automation decisions in multimodal logistics systems.

4 Methodology

4.1 Research design

The methodology follows a simulation and case-based design implemented in two L-Hubs: Trieste and Antwerp-Bruges. The framework combines four layers: IF and KG modelling, DT-based scenario simulation, and DSS-based evaluation supported by cost-benefit analysis. This design allows the paper to examine both how PI-relevant dependencies can be modelled and how automation interventions behave in distinct multimodal environments.

4.2 Information Flows modelling

IF modelling identifies how operationally relevant information is generated, transferred, and used across multimodal freight operations. System boundaries and relevant entities are first defined. Digital flows are then mapped with particular attention to the exchanges that affect coordination quality, responsiveness to disruption, and automation reliability. These structures are integrated into the KG and the simulation environment, making the IF layer the link between operational reality and the decision logic of the DSS (AUTOSUP, D2.3, 2025).

4.3 Knowledge Graph and interdependency modelling

The KG represents the multimodal system through nodes representing operational entities and edges showing physical connections, process dependencies and information exchanges. It offers a consistent structure for integrating heterogeneous operational knowledge and supports explainability by making dependencies explicit. The graph acts as a knowledge layer supporting both the DT and DSS rather than as a standalone analytical artefact (AUTOSUP, D2.3, 2025).

4.4 Digital Twin modelling and scenario analysis

The Digital Twins in AUTOSUP are simulation-based representations of the L-Hubs and their surrounding transport interfaces. Their purpose is to compare existing operating conditions with alternative automation-enabled scenarios under realistic assumptions. In both L-Hubs, the DT models are built as agent-based microscopic models using SUMO, as it represents vehicle movements, network geometry, delays, throughput, fuel and energy consumption, and emissions in a manner that is comparable across use cases (AUTOSUP, D2.3, 2025; AUTOSUP, D2.1, 2025; AUTOSUP, D2.2, 2026). Specific steps in different processes, such as container loading and unloading, inspections, and so on, are represented as stops along the trip starting from one location within the L-Hub to another with durations that are representative of different operational conditions.

Each use case is depicted using an AS-IS baseline and several TO-BE scenarios. The first reflects current operational conditions, while the latter represent automation interventions, coordination improvements, or infrastructure changes. As an example, predictive maintenance in Trieste compares the current maintenance structure, which produces approximately 400,000 hours of annual wagon downtime, with an automation-enabled configuration combining a smart gate, an on-site repair workshop, and an enhanced Gorizia maintenance centre, reducing total downtime to about 72,500 hours (AUTOSUP, D2.1, 2025). In Antwerp-Bruges, the AS-IS and TO-BE scenarios compare conventional and autonomous truck operations on the internal port road, together with alternative operating windows and electric vehicle configurations (AUTOSUP, D2.2, 2026).

The simulations are informed by hub-specific operational knowledge, available historical data, and assumptions validated with the L-Hub stakeholders. In Antwerp-Bruges, baseline calibration used publicly available data of traffic counts and trip durations (Vectris, 2015),

vehicle kinematic characteristics, and public emissions data through an iterative process that seeks to match model outputs for trip delays with recorded data. In Trieste, the DT relied on operational maintenance parameters, train numbers, travel distances, faulty wagons, and repair times that were provided by the L-Hubs (AUTOSUP, D2.1, 2025; AUTOSUP, D2.2, 2026). Models were calibrated with respect to route assignment and temporal distribution of movements to ensure that trip durations and recorded delays were consistent with historical L-Hub data.

4.5 Decision Support System and evaluation pipeline

The DSS is designed as a human-in-the-loop platform that supports scenario definition, execution, interpretation, and multiple scenario comparison. Users interact with the platform through scenario configuration tools, graph-based representations of hub dependencies, dashboards, and map-based visualisations. The DSS supports informed judgement by making the consequences of alternative choices visible across operational, economic, and environmental dimensions (AUTOSUP, D2.3, 2025).

The evaluation pipeline connects IF, the KG, and the DT models with a cost-benefit and feasibility layer (Figure 2). SUMO simulations provide transport performance metrics, such as trip duration, delays, throughput, fuel and energy consumption, emissions, and noise. These outputs are then used as inputs to the cost-benefit model, which evaluates capital and operational costs, environmental benefits, productivity effects, and demand-related impacts over a defined planning horizon (AUTOSUP, D2.3, 2025; AUTOSUP, D3.2, 2025). This links operational simulation with strategic investment evaluation rather than treating the two as separate tasks.

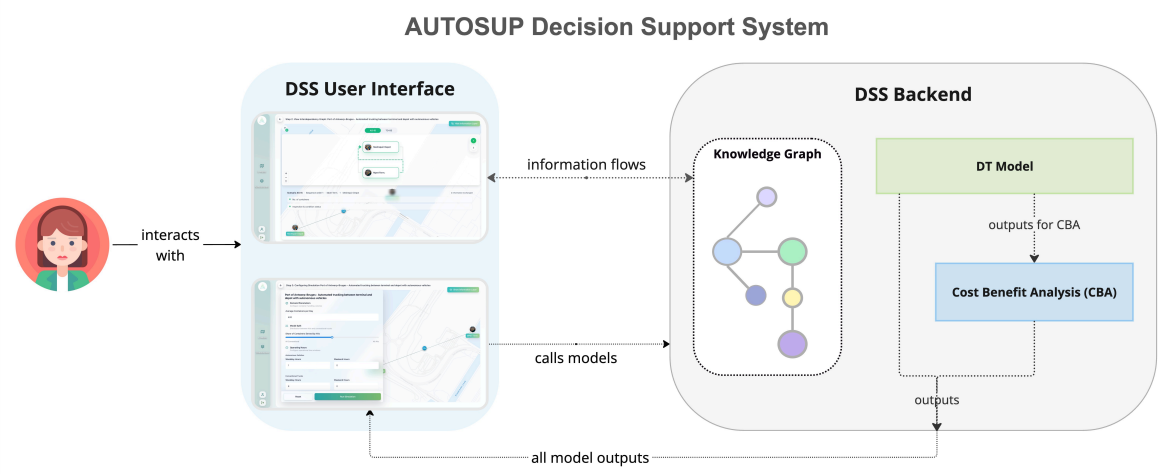


Figure 2: AUTOSUP model execution pipeline

4.6 Validation protocol and comparative logic

Validation is carried out through use cases in Trieste and Antwerp-Bruges. In each hub, the framework compares AS-IS and different TO-BE operational conditions. Specifically, it compares how the same modelling logic performs under different modal compositions, traffic conditions, and governance settings. This comparative design is important for PI-oriented research since it explores whether the same decision-support principles remain useful across heterogeneous multimodal environments. The current validation stage is designed to support comparative scenario assessment rather than statistical generalisation. For each use case, the AS-IS corresponds to the current operating logic of the L-Hub, while the TO-BE scenarios represent specific automation or coordination interventions discussed with hub stakeholders.

The purpose of the comparison is to examine whether the integrated modelling environment captures meaningful changes in performance under realistic assumptions. In this sense, the baseline is not a generic non-PI benchmark, but the current operational configuration against which a PI-oriented and automation-enabled alternative can be assessed.

5 L-Hubs and use cases

5.1 Comparative role of the two L-Hubs

The two L-Hubs (Trieste and Antwerp-Bruges) were selected as they represent different multimodal structures, operational constraints, and automation trajectories, while still sharing the need to coordinate freight flows across interconnected infrastructures and actors.

Trieste provides a rail-centred context in which terminal handling, wagon maintenance, and cross-border continuity are closely connected. The selected use cases focus on predictive rail maintenance, multimodal slot coordination, and automated cross-border Ro-Ro transhipment. The main automation questions consider reliability, asset availability, and the reduction of operational friction in intermodal rail processes. Antwerp-Bruges provides a denser port environment in which internal traffic, terminal coordination, and environmental performance are central. Its use cases focus on port automation coordination, ETA-based slot adjustment, and autonomous internal trucking. Automation questions focus on congestion, dwell time, traffic regularity, and local emissions performance.

Table 1 summarises all use cases, decision problems, DSS users, and KPI focus. Its role is to compare baseline and automation-oriented scenarios and support decisions on operational change, coordination strategy, and investment prioritisation.

Table 1: Comparative overview of the L-Hubs and selected use cases

L-Hub	Use case	Decision focus	KPI focus
Trieste	UC1 Predictive rail maintenance	Wagon downtime and maintenance coordination	Downtime, asset availability, maintenance cost
Trieste	UC2 Multimodal slot and last-mile routes coordination	Carrier selection, route planning, and slot coordination	Waiting time, slot adherence, coordination efficiency, service reliability
Trieste	UC3 Automated cross-border Ro-Ro transhipment	Alignment of rail, terminal, and maritime operations	Handling continuity, transhipment efficiency
Antwerp-Bruges	UC1 Port automation coordination	Congestion and operating window management	Congestion, dwell time, emissions
Antwerp-Bruges	UC2 ETA-based slot adjustment	Terminal and distribution-centre coordination	Slot adherence, throughput reliability
Antwerp-Bruges	UC3 Autonomous trucking and MagRail feasibility	Internal transport efficiency and emissions	Dwell time, fuel and energy use, emissions

In Trieste, the AUTOSUP framework is tested in use cases where automation is expected to improve rail reliability and maintenance efficiency, while in Antwerp-Bruges automation is expected to improve traffic coordination and environmental performance. This difference offers the basis for the comparative interpretation of the current validation results.

6 Results

6.1 Scope of current validation evidence

At the current stage of AUTOSUP, quantified validation results are available for one use case per L-Hub. In Trieste, the most mature evidence concerns the predictive rail maintenance use case. In Antwerp-Bruges, the most mature evidence concerns the autonomous trucking use case between MPET and Medrepair terminals. The remaining use cases, which are implemented in the framework, are still under assessment, and thus lack comparable quantitative detail at this point.

Table 2 showcases whether the combined use of IF, KGs, DTs, and DSS-based evaluation can support PI-aligned automation decisions in practice. These two use cases were selected for quantified reporting as they currently provide the most mature combination of calibrated DT outputs, scenario definition, and interpretable KPI evidence in each L-Hub.

Table 2: Current quantified validation results from the AUTOSUP L-Hubs

L-Hub	Use case	Baseline	TO-BE scenario	Main quantified result	Interpretation
Trieste	UC1, Predictive rail maintenance	Current maintenance structure with no smart gate and no dedicated on-site repair workshop	Smart gate, on-site Trieste workshop, enhanced Gorizia maintenance centre	Annual wagon downtime reduced from about 400,00(hours to about 72,500 hours saving about 330,000 hours	Reliability gains through earlier diagnosis, local repair capability, and shorter maintenance trave distance
Antwerp-Bruges	UC3.1, Autonomous trucking between MPET and Medrepair	Diesel trucks operating on the private road circuit	Autonomous truck scenarios, including a fully electric configuration	Fuel and emissions reductions of about 50 percent in the lower-scale autonomous scenario, approaching elimination in the fully electric scenario	Environmental and traffic gains linked to smoother internal circulation and electrified vehicle use

6.2 Trieste, predictive rail maintenance

Predictive rail maintenance in Trieste compares the existing maintenance structure with TO-BE configurations that combine a smart gate, an on-site repair workshop, and an enhanced maintenance centre in Gorizia. The baseline produces around 400,000 hours of annual wagon downtime. This drops to about 72,500 hours in the TO-BE scenario, corresponding to savings of about 330,000 hours. This result indicates a substantial improvement in asset availability and rail operational continuity. It also shows that the value of automation lies in the reorganisation of the maintenance process through earlier fault identification, shorter maintenance travel distance, and improved coordination between inspection, repair, and dispatch functions (AUTOSUP, D2.1, 2025; AUTOSUP, D3.4, 2026). At the current stage, the model evaluates the combined effect of the intervention package and does not isolate the marginal contribution of each component separately.

6.3 Antwerp-Bruges, autonomous trucking and internal port traffic

Autonomous trucking in Antwerp-Bruges compares conventional internal diesel truck operations with automation-oriented scenarios, including a fully electric variant. The current quantified outputs focus on fuel consumptions and emissions. Depending on the scenario, the model indicates reductions of approximately 50% in the lower-scale autonomous case and near elimination in the fully electric autonomous case. These results are associated with smoother internal circulation, fewer inefficient vehicle movements, and the replacement of diesel-based operations (AUTOSUP, D2.2, 2026; AUTOSUP, D3.3, 2026). Accordingly, the current result should be interpreted as the combined effect of improved coordination and vehicle substitution, rather than as a separated estimate of the PI logic alone.

6.4 Comparative interpretation

The two quantified use cases illustrate different forms of PI-oriented automation support. In Trieste, the main gain is reliability in rail-related operations through reduced maintenance downtime. In Antwerp-Bruges, the main gain is improved traffic coordination and lower environmental burden within the port area. This contrast reflects the different operational roles of the two hubs and confirms that the impact of automation depends on the structure of dependencies within each multimodal system.

The results suggest that automation generates stronger benefits when it is assessed as part of a coordinated operational configuration rather than as an isolated technical intervention. This is consistent with the AUTOSUP decision-support logic, in which Information Flows, interdependencies, simulation, and economic reasoning are treated as interconnected elements of automation pathway assessment.

6.5 Implications for PI and synchromodality

The current evidence indicates several implications for PI implementation. First, the benefits observed in both L-Hubs depend on coordination conditions rather than on automation alone. Second, explicit modelling of Information Flows and interdependencies is necessary for explainable and PI-compatible automation decisions. Third, although the current use cases do not implement full synchromodal control, the framework already supports the comparison of alternative operational configurations, timing conditions, and coordination patterns.

7 Discussion

The results show that PI-oriented automation in multimodal hubs depends on the joint modelling of automation functions, information exchange, and operational dependencies. In both L-Hubs, the benefits do not come from automation alone. They depend on how the intervention changes the structure of coordination across processes and actors.

The two cases illustrate different decision contexts. In Trieste, the main gain is improved rail reliability through reduced maintenance downtime. In Antwerp-Bruges, the benefit is improved internal traffic performance and lower environmental burden. This contrast underscores the usefulness of the AUTOSUP framework in diverse multimodal settings demonstrating that its value resides in the integration of modelling layers rather than reliance on any single automation technology.

The present validation is partial, as quantified evidence exists for one use case per L-Hub. Nevertheless, the results suffice to demonstrate the framework's operationality and that its capacity to enable meaningful comparisons between baseline and automation-oriented scenarios within real hub contexts. Although the current evidence does not demonstrate real-time synchromodal control, it lays a foundation for future synchromodal evaluations in PI-oriented hub systems.

8 Conclusions

This paper introduces the AUTOSUP framework for implementing PI principles in multimodal logistics hubs via IF modelling, KG-based interdependency representation, simulation-based DTs, and a human-in-the-loop DSS. It addresses a gap in existing work by integrating automation, interoperability, and economic feasibility within real multimodal environments.

Early evidence from Trieste and Antwerp-Bruges shows that the framework can support meaningful comparisons between baseline and automation-oriented scenarios. In Trieste,

predictive maintenance substantially reduced wagon downtime. In Antwerp-Bruges, autonomous trucking lowered fuel consumption and emissions. These results remain partial, yet they show the practical relevance of the approach, and support extension to the remaining project use cases.

Acknowledgements. This work is part of the AUTOSUP project (101147468), which has received funding by the European Union. Views and opinions expressed are, however, those of the author (s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

References

- ALICE (2022). Roadmap to the Physical Internet. Alliance for Logistics Innovation through Collaboration in Europe.
- AUTOSUP Project (2025). D2.1 Digital Twin Models of the Trieste L-HUB. Horizon Europe AUTOSUP, Project Deliverable.
- AUTOSUP Project (2026). D2.2 Digital Twin Models of the Antwerp-Bruges L-HUB. Horizon Europe AUTOSUP, Project Deliverable.
- AUTOSUP Project (2025). D2.3 DSS for operational automation in multimodal hubs. Horizon Europe AUTOSUP, Project Deliverable.
- AUTOSUP Project (2025). D3.2 Cost-Benefit Analysis. Horizon Europe AUTOSUP, Project Deliverable.
- AUTOSUP Project (2026). D3.3 Feasibility Studies for the Port of Antwerp-Bruges Living Hub. Horizon Europe AUTOSUP, Project Deliverable.
- AUTOSUP Project (2026). D3.4 Feasibility Studies for the Trieste Living Hub. Horizon Europe AUTOSUP, Project Deliverable.
- Barricelli, B.R., Casiraghi, E., and Fogli, D. (2019). A survey on digital twin: Definitions, characteristics, applications, and design implications. IEEE Access, 7, 167653–167671.
- Heilig, L., Schwarze, S., and Voß, S. (2017). An analysis of digital transformation in the history and future of modern ports. Computers in Industry, 89, 12–30.
- Labarthe, O., Klibi, W., and Montreuil, B. (2023). Urban synchromodality: Synergies between freight and people mobility in urban areas. Proceedings of the IPIC 2023 International Physical Internet Conference, Athens, Greece.
- Lindgren, P., Sörensen, C., and Jensen, T. (2020). Decision support systems for transport logistics: A literature review. Transportation Research Procedia, 46, 285–292.
- Meller, R.D., Ellis, K.P., and Loftis, B. (2012). From horizontal collaboration to the Physical Internet: Quantifying the effects on sustainability and profits when shifting to interconnected logistics systems. Final Research Report of the CELDi Physical Internet Project, Phase I.
- Montreuil, B. (2011). Towards a Physical Internet: Meeting the global logistics sustainability grand challenge. Logistics Research, 3(2–3), 71–87.
- Raptis, T.P., Passarella, A., and Conti, M. (2022). Digital twin empowered smart ports for sustainable logistics. IEEE Internet of Things Journal, 9(6), 4202–4214.
- Vectris (2015, Mobiliteit in Ringland, Ringgenootschap, https://ringland.be/sites/default/files/2021-06/rapport-mobiliteit-Ringland-Vectris.pdf.
- Villarreal, B., Garza-Reyes, J.A., and Kumar, V. (2021). Smart logistics: A literature review and future research agenda. Journal of Industrial Information Integration, 22, 100193.

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE
IMT Mines Alès École Mines-Télécom
IMT Mines Albi École Mines-Télécom

Enabling the Physical Internet through AI-Driven Federated Interoperability: An Ontology Matching Perspective

Leonardo Daou[1,2], Zhuoxun Zheng[3], Syed Juned Ali[3], Dominik Bork[3], Eva Petitdemange[2], Gregory Zacharewicz[1], Séverine Durieux[2], and Nicolas Daclin[1]
1. UPR SyCoIA, IMT Mines Alès, Alès, France
2. Centre Génie Industriel, IMT Mines Albi – University of Toulouse, Albi, France
3. Business Informatics Group, TU Wien, Austria
Corresponding author: leonardo.daou@mines-ales.fr

Abstract: *Semantic interoperability remains a key challenge in decentralized logistics systems, particularly within the context of the Physical Internet, where heterogeneous systems must dynamically exchange and interpret information. This paper investigates the integration of ontology matching within the translation module of a federated interoperability architecture. Building on the MILA framework, we evaluate multiple configurations under varying levels of additional contextual information and large language model (LLM) involvement. Experiments conducted on heterogeneous logistics ontologies show that LLM-driven approaches significantly improve semantic alignment between heterogeneous systems, particularly as model size increases. However, performance remains dependent on model capacity and the quality of ontology representations, especially in the presence of incomplete or weakly defined concepts. The results highlight a trade-off between efficiency and accuracy, as well as limitations affecting reliable real-time interoperability, motivating future work to focus on enhancing semantic representations and incorporating relational information to improve robustness and scalability.*

Keywords: *Physical Internet, Federated Interoperability, Ontology, Ontology Matching, Semantic Interoperability, Large Language Models.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☐ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☒ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☒ *Governance.*

Targeted Delivery Mode – s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☐ *In – Person presentation*

1 Introduction

Global challenges such as supply chain disruptions, climate change, and increasing demands for sustainable logistics require enhanced coordination among distributed and heterogeneous organizations. Modern logistics ecosystems involve autonomous stakeholders – such as carriers and service providers – that must dynamically exchange and interpret information. The Physical Internet (PI) emerges as a paradigm addressing these challenges by promoting open, modular, and hyperconnected logistics networks, enabling decentralized routing, shared infrastructures, and multimodal coordination (Ballot et al., 2021). However, realizing this vision requires seamless interoperability across highly heterogeneous and independently managed systems.

Despite the emergence of digital platforms, compatibility barriers persist due to heterogeneity in data models, operational processes, and semantic representations (Pan et al., 2021; Wu et al., 2025). Consequently, interoperability – defined as the ability of systems to exchange and meaningfully use information across heterogeneous environments – becomes essential to

overcome these barriers (ISO/TC 184/SC 5, 2011). These challenges span multiple layers, including data, service, process, and business levels, each introducing distinct constraints. While integrated and unified approaches respectively rely on standardization and shared representations, they often limit flexibility and scalability (Chen et al., 2008). In contrast, federated interoperability (FI) enables dynamic collaboration with minimal coupling while preserving organizational autonomy, making it particularly suitable for decentralized environments such as those envisioned by the PI (Daou et al., 2025).

In this context, the AI-driven federated interoperability architecture proposed by Daou et al. (2026) introduces a modular coordination layer composed of task reasoning, translation, resource memory, and workflow orchestration. It supports runtime adaptation through mechanisms such as short-lived ontologies, similarity-based discovery, and federated learning, enabling systems to interact while preserving autonomy and minimizing coupling. However, real-world adoption of PI remains limited due to interoperability gaps, heterogeneous systems, and governance challenges across independent actors (Pan et al., 2021; Wu et al., 2025).

Semantic interoperability is therefore a central challenge. Ontologies provide formal representations of domain knowledge, enabling shared understanding across systems (Grubic and Fan, 2010). However, independently developed ontologies introduce heterogeneity in structure, terminology, and abstraction levels, making direct interoperability difficult. Ontology matching addresses this issue by identifying correspondences between entities, enabling semantic alignment across systems (Euzenat and Shvaiko, 2013).

Within this architecture, the translation module plays a central role by transforming exchanged messages into representations interpretable by recipient systems. It relies on context enrichment mechanisms, including short-lived ontologies and adaptive reasoning, to ensure semantic consistency at runtime. In this setting, ontology matching becomes a key mechanism for aligning heterogeneous ontologies without requiring shared representations. However, enabling ontology matching in federated PI environments introduces significant challenges: interactions are dynamic, context-dependent, and short-lived, requiring alignment mechanisms capable of operating at runtime with limited information and without predefined mappings. Although recent approaches leverage artificial intelligence, particularly large language models, to support adaptive translation and context-aware reasoning (Taboada et al., 2025), their integration into operational interoperability architectures remains limited.

To address these challenges, this work investigates the following questions: (1) How can ontology matching techniques support dynamic, context-aware semantic alignment within the translation module of a federated interoperability architecture? (2) How do different matching strategies combine embeddings, heuristics, and large language models affect alignment performance? (3) To what extent does the proposed approach achieve accurate and reliable semantic alignment when evaluated against a reference mapping in federated environments where runtime execution and limited information are critical factors? As such, the analysis focuses on the trade-offs between matching performance, contextual enrichment, and large language model (LLM) involvement. Subsequently, we detail interoperability in logistics and Physical Internet domains, with a focus on the federated approach, and review related ontology matching approaches in the Related Work section. We then present the selected ontologies from the Physical Internet domain and a federated interoperability-aligned ontology matching approach along with its evaluated variations in the Methods section. The Results section reports performance, followed by an analysis of implications in the Discussion section. Finally, the Conclusion summarizes the contributions and outlines future research directions.

2 Related Work

2.1 Interoperability in Logistics and the Physical Internet

Research has increasingly emphasized interoperability, particularly semantic interoperability, to ensure consistent interpretation of exchanged information across systems in various fields such as supply chain management (Pan et al., 2021). This is critical in logistics and Physical Internet environments, where decentralized and dynamic interactions require flexible and scalable solutions (Ballot et al., 2021).

Interoperability approaches are commonly classified as integrated, unified, and federated. Integrated and unified approaches rely on tight coupling or shared representations, limiting adaptability, whereas federated interoperability enables collaboration with minimal coupling while preserving autonomy (Chen et al., 2008; Daou et al., 2025).

Early federated approaches leveraged multi-agent systems with short-lived ontologies for context-specific mediation (Zacharewicz et al., 2009), followed by distributed solutions such as Data Spine for heterogeneous data transformation (Deshmukh et al., 2021) and hypergraph- or embedding-based methods enabling dynamic mapping but with limited reusability (Labreche, 2023). More recently, large language models enable adaptive translation, reasoning, and semantic alignment without predefined mappings, while federated learning supports privacy-preserving evolution (Taboada et al., 2025; Daou et al., 2026). However, achieving consistent semantic interoperability in dynamic federated environments remains challenging without shared representations, highlighting the need for ontology-based approaches to support semantic interoperability in such environments.

2.2 Ontologies in Logistics, Transport, and Physical Internet

Ontologies are widely used in logistics, transport, and Physical Internet domains to support semantic interoperability by providing formal representations of domain knowledge (Grubic and Fan, 2010). They enable consistent interpretation of entities such as shipments, resources, and services. Applications include semantic service discovery in logistics (Daniele and Pires, 2013) and integration of multimodal transport data for decision-making (Moreno et al., 2015). However, independently developed ontologies introduce heterogeneity in structure and semantics, necessitating ontology matching techniques.

2.3 Ontology Matching Techniques

Ontology matching addresses semantic heterogeneity by identifying correspondences between entities in different ontologies, enabling data integration and interoperability (Euzenat and Shvaiko, 2013). Techniques include lexical, structural, semantic, and instance-based approaches, as well as hybrid methods that combine them to improve accuracy (Khan et al., 2023). Recent research focuses on embedding-based and neural approaches, where entities are represented in vector spaces capturing semantic similarity beyond lexical overlap (Hao et al., 2023). Advances in ontology embeddings, including geometric, sequence-based, and graph-based representations, further enhance semantic modeling (Chen et al., 2025). Neural and NLP-based methods incorporate both textual and structural information, while recent hybrid approaches integrate embeddings, Graph Neural Networks (GNNs), and large language models (LLMs) to improve contextual reasoning and matching performance (Lu and Yan, 2025). Despite these advances, most approaches remain designed for static environments with predefined mappings, limiting their applicability in dynamic settings such as logistics networks and Physical Internet systems.

2.4 Ontology Matching in Dynamic and Federated Environments

In the Physical Internet domain, particularly in logistics and transport systems, ontology matching faces challenges due to the dynamic, heterogeneous, and decentralized nature of these environments. Concepts may vary depending on context – such as geographic region, transport mode, or organizational practices – making consistent mappings difficult (Wu et al., 2025).

Additionally, continuously evolving logistics networks require matching mechanisms that can adapt dynamically. Recent research leverages artificial intelligence to address these challenges. Machine learning and deep learning approaches enable the automatic discovery of semantic correspondences, while embedding-based methods capture semantic similarity beyond lexical matching (Kulmanov et al., 2021; Li et al., 2021), improving robustness and scalability. They further demonstrate the practical effectiveness of adaptive and context-aware semantic alignment approaches in heterogeneous environments (Kulikov et al., 2025). These results suggest that adaptive semantic translation mechanisms are not purely conceptual, although their application to federated interoperability in Physical Internet environments remains limited.
In parallel, interoperability frameworks emphasize dynamic and context-aware semantic alignment. In federated environments, ontology matching must operate on demand without predefined mappings or centralized representations. Approaches based on short-lived ontologies enable context-specific alignment by focusing on relevant knowledge for each interaction (Zacharewicz et al., 2009), supporting flexibility while preserving autonomy. Despite these advances, integration into operational interoperability mechanisms remains limited, highlighting the need for solutions capable of real-time semantic alignment in decentralized Physical Internet systems.

3 Methods

3.1 Ontologies

Various ontologies alongside their reference mapping between each other concerning Physical Internet, supply chain, and logistics exist, but they are mostly proprietary. Hence, we chose two open-source ontologies, namely LogiServ (Daniele and Pires, 2013) and OneRecord (IATA, 2023). The first is composed of 378 entities, while the second is composed of 857 entities. Entities refer to all ontology elements, including classes (concepts defined in Section 2.4), properties (relations or attributes linking or describing classes and instances), and instances (individuals representing concrete occurrences of classes).
Furthermore, given that we are investigating the feasibility of using ontology matching to facilitate semantic interoperability in the translation module of the federated interoperability architecture (Daou et al., 2026), these two ontologies do not have a shared representation layer that they can be mapped to and are heterogeneous. However, in order to evaluate ontology matching, we must create a reference mapping. To remediate that, we picked subsets from both LogiServ and OneRecord that have a direct relation with the Physical Internet domain. This reduced the number of entities from 378 to 123 for the former, and from 857 to 384 for the latter. After selecting the appropriate entities, each term and its match were manually selected and this results in three cases: exact matches, matches with higher abstraction level, or no matches. An example regarding the second case is the relation between a truck and transportation means, with the latter being an abstraction of the former. The full details of their specific composition in terms of classes, properties, and instances, as well as the reference matches between the two ontologies are present in Table 1.

Table 1: Statistics about the filtered ontologies

Ontology	Ontology Type	Entity Counts	Class Counts	Property Counts	Instance Counts	Matched Entities	Unmatched Entities
Logiserv	Source	123	90	33	0	80	43
OneRecord	Target	384	91	252	41		

3.2 Ontology Matching

Federated interoperability requires plug-and-play functionalities that operate with minimal information while maintaining high matching performance. Unlike unified approaches, no shared representation layer is available, constraining the choice of ontology matching techniques for the translation module (Daou et al., 2026). For these reasons, we investigate the MILA (Minimizing LLM Prompts in Ontology Mapping) framework (Taboada et al., 2025) as its principles align with federated interoperability requirements. It is a LLM-based ontology matching approach that combines embedding-based retrieval with a retrieve-identify-prompt pipeline, designed to reduce reliance on costly LLM queries. This also aligns with the specifications of the translation module of the federated interoperability architecture (Daou et al., 2026) which leverages LLMs for flexibility and automation. Its objective is to align a source ontology with a target ontology, namely LogiServ and OneRecord, respectively. It starts by retrieving candidate correspondences by encoding source and target ontology entity labels into vector representations using SBERT (Reimers and Gurevych, 2019), a pre-trained embedding model, and then computes their semantic similarities through cosine similarity. Based on this, the most relevant k target candidates are selected for each entity in the source and target ontologies, with candidates having similarity scores below a predefined threshold t are discarded to improve efficiency. In our case, k and t were empirically chosen to be 5 and 0.4.

The framework then identifies high confidence bidirectional (HCB) matches through lightweight heuristics and only invokes the LLM for ambiguous cases. HCB matches refer to correspondence where two entities from different ontologies are mutually identified as the best match for each other, indicating strong semantic agreement. For example, if the entity “shipment” in a source ontology is matched as the top candidate to “consignment” in a target ontology, and “consignment” is also independently matched back as the top candidate to “shipment”, this reciprocal alignment constitutes a high-confidence bidirectional match.

MILA’s process is guided by a prioritized depth-first search (PDFS) strategy, which explores candidate correspondences in decreasing similarity order. For each source entity, these candidates are evaluated in decreasing similarity order until a valid match is identified. This enables early termination once a correspondence is confirmed, reducing unnecessary comparisons and limiting LLM queries to ambiguous cases. As a result, the approach ensures efficient exploration of candidate mappings while minimizing computational overhead and maintaining scalability (Taboada et al., 2025). In the context of the translation module (Daou et al., 2026), this reduces LLM usage and query complexity, allowing smaller models (1-12B parameters) to achieve satisfactory performance in real-time settings compared to much larger models with hundreds of millions to trillions of parameters.

Building on this framework, we define four evaluation scenarios reflecting different levels of contextual enrichment and LLM involvement:

- Scenario 1 (Baseline MILA): The original MILA configuration is applied, relying solely on entity labels for candidate retrieval and matching. High-confidence bidirectional matches are identified first, and the LLM is only invoked for ambiguous cases using the PDFS strategy. This represents the minimal-information, efficiency-oriented setting aligned with federated interoperability constraints.
- Scenario 2 (MILA and Definitions): The baseline approach is extended by enriching ontology entities with their textual definitions. These definitions are incorporated into the matching process to provide additional semantic context, while preserving the original MILA workflow (including HCB filtering and PDFS). This scenario evaluates the impact of additional contextual information on matching performance.
- Scenario 3 (Top-k LLM Selection with Definitions): When no high-confidence correspondence is identified, instead of iteratively exploring candidates via PDFS, the top-k candidates – augmented with their definitions – are directly provided to the LLM. The

LLM then selects the most appropriate correspondence or rejects all candidates. This variant increases reliance on the LLM while reducing search complexity.

- Scenario 4 (Full LLM-driven Matching): The high-confidence filtering step is removed entirely, and the LLM is systematically invoked with the top-k candidates and their definitions to determine the final match. This represents the most permissive configuration, maximizing LLM involvement and minimizing heuristic constraints.

Within this process, the LLM serves as the core decision-making component responsible for resolving semantic ambiguity among candidate correspondences. To implement this step, we employ two LLM families available through Ollama – DeepSeek-R1 and Gemma – each with three variants of increasing size: DeepSeek-R1 with 1.5B, 7B and 8B parameters, and Gemma 3 with 1B, 4B and 12B parameters. This setup enables us to analyze how model capacity influences matching performance and computational cost in the context of ontology matching. These variations allow us to analyze the trade-offs between efficiency, context availability, and LLM reliance, providing insights into how ontology matching mechanisms can be adapted for dynamic federated interoperability environments, where both information availability and computational constraints may vary.

3.3 Evaluation Metrics

All ontology matching systems were evaluated using standard information retrieval metrics, namely Precision (P), Recall (R), and F-measure. As shown in equation (1), precision assesses the correctness of the produced alignment by measuring the proportion of correspondences – matched entities between source and target ontologies – identified by the system (I) that are also present in the reference alignment (T). Recall, defined in equation (2), evaluates alignment completeness by measuring the proportion of reference correspondences successfully identified by the system. Subsequently, as given in equation (3), the F-measure combines Precision and Recall into a single metric, providing a balanced assessment of correctness and completeness.

$$P(I,T) = \frac{|I \cap T|}{|I|} \quad (1)$$

$$R(I,T) = \frac{|I \cap T|}{|T|} \quad (2)$$

$$F - Measure(I,T) = 2 \times \frac{P \times R}{P + R} \quad (3)$$

4 Results

We compare the performance of the different MILA-based scenarios across multiple model sizes belonging to the DeepSeek-R1 and Gemma 3 LLM families to evaluate the effectiveness of the proposed ontology matching configurations. The evaluation focuses on Precision, Recall, and F-measure, which provide a balanced assessment of alignment correctness and completeness. The results of this evaluation are summarized in Table 2. The three evaluation metrics values are identical due to the one-to-one matching evaluation setup.

Overall, Scenario 1 and Scenario 2 yield stable but low performance across all model sizes, with F-measure values consistently around 0.48. The inclusion of textual definitions alone does not significantly improve matching performance, indicating that additional semantic context is not effectively leveraged under strict MILA constraints. In contrast, Scenario 3 shows a clear improvement as model size increases. This trend is clearly visible in Figure 1, where both model families exhibit a steady rise in F-measure values as model capacity increases. Smaller models exhibit limited gains, while medium and large models achieve progressively higher performance. Notably, both DeepSeek-R1 and Gemma 3 reach an F-measure of 0.594 with larger configurations, demonstrating the importance of model capacity when increasing

reliance on LLM-based reasoning. Scenario 4 further reinforces this trend. As illustrated in Figure 1, performance diverges significantly depending on model size, whereby smaller models experience a drop in performance due to the removal of heuristic constraints, while larger models significantly outperform all other configurations. DeepSeek-R1 (8B) achieves the highest F-measure of 0.724, followed by Gemma 3 (12B) with 0.610. This confirms that relaxing heuristic constraints can be beneficial when sufficiently powerful LLMs are used.
These results highlight a trade-off between efficiency and performance: heuristic-driven approaches ensure stability with minimal computational cost, while LLM-driven approaches require larger models but yield superior semantic alignment performance.

Table 2: Results of the various MILA variations based on the Large Language Model used.

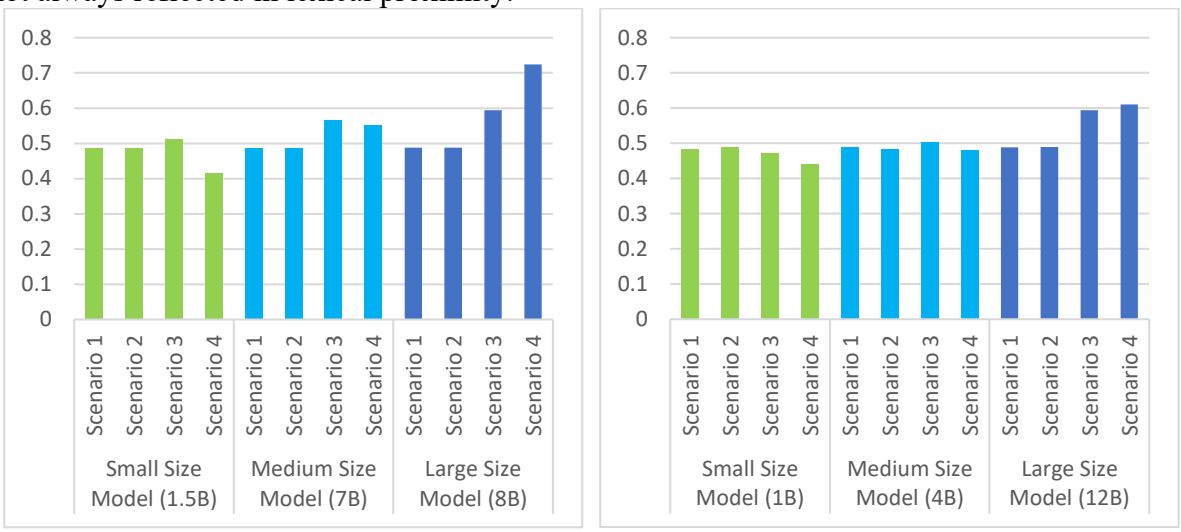
Model Size	Scenario	DeepSeek-R1			Gemma 3		
		P	R	F-measure	P	R	F-measure
Small (DP:1.5B and GE:1B)	1	0.488	0.488	0.488	0.484	0.484	0.484
	2	0.488	0.488	0.488	0.489	0.489	0.489
	3	0.512	0.512	0.512	0.472	0.472	0.472
	4	0.415	0.415	0.415	0.439	0.439	0.439
Medium (DP:7B and GE:4B)	1	0.488	0.488	0.488	0.488	0.488	0.488
	2	0.488	0.488	0.488	0.484	0.484	0.484
	3	0.565	0.565	0.565	0.504	0.504	0.504
	4	0.553	0.553	0.553	0.480	0.480	0.480
Large (DP:8B and GE:12B)	1	0.488	0.488	0.488	0.488	0.488	0.488
	2	0.488	0.488	0.488	0.489	0.489	0.489
	3	0.594	0.594	0.594	0.594	0.594	0.594
	4	**0.724**	**0.724**	**0.724**	*0.610*	*0.610*	*0.610*

Note: Scenarios 1 to 4 refer to the scenarios detailed in Section 3.2. The model sizes are denoted by DP for DeepSeek-R1 and GE for Gemma 3, while the accompanying number indicates the number of parameters (in billions). Bold values indicate the highest results, while italic values indicate the second highest.

5 Discussion

The experimental results demonstrate that LLM-based approaches can significantly enhance ontology matching performance in federated interoperability contexts, particularly when heuristic constraints are relaxed. This improvement is particularly relevant in the context of the Physical Internet, as semantic alignment between heterogeneous and autonomous systems is a prerequisite for enabling decentralized and scalable logistics interactions. However, this improvement is highly dependent on model capacity, as smaller models struggle to compensate for reduced structural guidance.
This behavior can be partly explained by the characteristics of the evaluated ontologies. A key characteristic of the selected ontologies is their heterogeneity, both in structure and terminology, as well as the fact that some terms lack definitions or contain semantically weak or ambiguous descriptions. These factors directly impact matching performance and highlight

the challenges of applying ontology matching in real-world federated environments. Closer inspection of the results reveals several limitations of the base MILA framework. The reliance on HCB matching and PDFS introduces early termination issues. In many cases, the matching sequence stops prematurely either because an incorrect correspondence satisfies the HCB condition or because a candidate is selected before reaching the true semantic match. This behavior is particularly problematic in heterogeneous ontologies, where semantic similarity is not always reflected in lexical proximity.



(a) (b)

Figure 1: F-measure values for three model sizes of two families of LLMs across four different scenario variations of MILA. (a) DeepSeek-R1 models F-measure values. (b) Gemma 3 models F-measure values.

To address this, MILA's constraints were relaxed by increasing LLM involvement. This allows the system to explore a broader set of candidates and recover from early misalignments. However, this approach requires sufficiently large LLMs to be effective. As observed in Scenarios 3 and 4, larger models are able to overcome these limitations and surpass baseline performance, demonstrating improved reasoning and semantic discrimination capabilities. Another important advantage of the approach is its ability to correctly identify cases where no valid correspondence exists. Even when candidate matches are proposed, the system can reject them appropriately, reducing false positives and improving overall robustness. Within the context of the translation module of the federated interoperability architecture (Daou et al., 2026), these results indicate that increasing LLM involvement can enhance the module's ability to perform adaptive and context-aware semantic alignment at runtime.

Despite these improvements, several challenges remain. Some correspondences produced by the system are semantically inconsistent, while other entities fail to retrieve meaningful candidates. These issues stem from limitations in ontology entity representation, particularly when definitions are missing, incomplete, or not semantically informative. In addition, the current embedding approach relies primarily on surface-level representations, which do not fully capture the relational structure of ontologies. Ontology entities are not isolated elements but are embedded within a network of relationships, including hierarchical, associative, and property-based connections. These relationships provide essential contextual information that is critical for understanding the true semantics of a concept. Ignoring this structure limits the effectiveness of matching approaches that rely solely on textual or embedding-based similarity. Furthermore, these limitations have direct implications for real-time interoperability scenarios. In federated environments, where translation must occur under strict time constraints, premature termination, incorrect correspondences, or missing candidates can propagate errors

through the translation pipeline, affecting downstream tasks and overall system reliability. Such limitations impact the feasibility of semantic interoperability in Physical Internet environments, where interactions are distributed and require reliable on-the-fly translation between systems. This highlights the need for more robust and context-aware matching mechanisms. Furthermore, while the proposed approach incorporates high-confidence filtering and candidate rejection mechanisms to reduce false positives, the results indicate that scalability and operational reliability remain strongly dependent on ontology quality, contextual information, and model capacity.

In light of these limitations, our future work focuses on enhancing both semantic representation and contextual understanding within the matching process. First, we plan to leverage large language models to generate missing definitions and enrich existing ones, using the contextual relationships that entities have with other elements of the ontology. By incorporating information from neighboring entities and their relationships, LLMs can produce more precise and semantically meaningful descriptions, improving the overall quality of ontology representations. Second, we aim to revisit the embedding approach to better capture structural and contextual information. We intend to explore graph-based embedding techniques, such as Graph Neural Networks (GNNs), to encode both the features of entities and their relational context within the ontology. This would enable the model to account for the topology of the ontology graph, leading to more robust and semantically informed matching. Collectively, these directions aim to improve the reliability and scalability of ontology matching in federated interoperability environments by reducing semantically inconsistent correspondences and enhancing contextual understanding, where semantic richness and adaptability are critical.

6 Conclusion

This work investigates the integration of ontology matching within the translation module of a federated interoperability architecture (Daou et al., 2026). By evaluating different MILA configurations under varying levels of LLM involvement and contextual enrichment, the results show that LLM-driven approaches improve semantic alignment, particularly with larger models. These results contribute to the development of the translation module by identifying key factors that influence effective semantic alignment, reinforcing the role of ontology matching as a core mechanism in federated interoperability. In the context of the Physical Internet, this highlights the importance of dynamic and scalable semantic alignment mechanisms to enable interoperability across decentralized logistics systems.

However, performance remains dependent on model capacity and the quality of ontology representations, especially in heterogeneous settings with incomplete or weak definitions. These findings highlight both the potential and limitations of ontology matching in dynamic, runtime interoperability scenarios. To address these limitations, our future work focuses on improving semantic representations and incorporating relational context through advanced embedding techniques to enhance robustness and scalability.

References

- Ballot E., B. Montreuil, Z.G. Zacharia (2021): Physical Internet: First results and next challenges. Journal of Business Logistics, v42, 101–107.
- Chen D., G. Doumeingts, F. Vernadat (2008): Architectures for enterprise integration and interoperability: Past, present and future. Computers in Industry, Elsevier, v59, no7, 647–659.
- Chen J., O. Mashkova, F. Zhapa-Camacho, R. Hoehndorf, Y. He, I. Horrocks (2025): Ontology embedding: A survey of methods, applications and resources. IEEE Transactions on Knowledge and Data Engineering, 1–20.
- Daniele L., L.F. Pires (2013): An ontological approach to logistics. In Enterprise Interoperability (eds M. Zelm, M. van Sinderen, L.F. Pires, G. Doumeingts). Wiley.

- Daou L., E. Petitdemange, G. Zacharewicz, N. Daclin, S. Durieux (2025): Structuring federated data interoperability: A multi-level framework. Procedia Computer Science, Elsevier, v274, 664–675
- Daou L., E. Petitdemange, G. Zacharewicz, N. Daclin, S. Durieux (2026): A novel AI-driven system architecture for federated interoperability. In I-ESA 2026: Proceedings of the I-ESA Conferences (accepted).
- Euzenat J., P. Shvaiko (2013): Ontology Matching, 2nd edition. Springer.
- Grubic T., I.-S. Fan (2010): Supply chain ontology: Review, analysis and synthesis. *Computers in Industry*, Elsevier, v61, no8, 776–786.
- Hao Z., W. Mayer, J. Xia, G. Li, L. Qin, Z. Feng (2023): Ontology alignment with semantic and structural embeddings. Journal of Web Semantics, Elsevier, v78, 100798.
- IATA (2023): ONE Record ontology. International Air Transport Association. https://onerecord.iata.org/ns/cargo (last accessed: March 31, 2026)
- ISO/TC 184/SC 5 (2011): Advanced automation technologies and their applications – Requirements for establishing manufacturing enterprise process interoperability – Part 1: Framework for enterprise interoperability. Standard, International Organization for
- Khan H., M. Saqib, H.A. Khattak, S.I. Ali, S. Lee (2023): Ontology alignment for accurate ontology matching: A survey. In Digital Health Transformation, Smart Ageing, and Managing Disability (ICOST 2023), Lecture Notes in Computer Science, Springer, v14237.
- Kulikov I., J. Yang, N. Zhukova, T. Man (2026): An approach for linking dynamic network information models based on ontology matching. Expert Systems with Applications, Elsevier, v298, 129622.
- Kulmanov M., F.Z. Smaili, X. Gao, R. Hoehndorf (2021): Semantic similarity and machine learning with ontologies. Briefings in Bioinformatics, v22, no4, bbaa199.
- Labreche M. (2023): Implémentation d'une interopérabilité fédérée supportée par la transformation automatisée à la volée de modèles de données hétérogènes : application aux problèmes d'appariement des schémas. Doctoral dissertation, École des Mines d'Albi-Carmaux.
- Li H., Y. Govind, S. Mudgal, T. Rekatsinas, A. Doan (2021): Deep learning for semantic matching: A survey. Journal of Computer Science and Cybernetics, v37, no4, 365–402.
- Lu J., C. Yan (2025): Ontology matching method based on deep learning and syntax. Big Data and Cognitive Computing, v9, no8, 208.
- Moreno A., A. Perallos, D. López-de-Ipiña, E. Onieva, I. Salaberria, A.D. Masegosa (2015): A novel software architecture for the provision of context-aware semantic transport information. Sensors, v15, no6, 12299–12322.
- Pan S., D. Trentesaux, D. McFarlane, B. Montreuil, E. Ballot, G.Q. Huang (2021): Digital interoperability in logistics and supply chain management: State-of-the-art and research avenues towards Physical Internet. Computers in Industry, Elsevier, v128, 103435.
- Reimers N., I. Gurevych (2019): Sentence-BERT: Sentence embeddings using Siamese BERT networks. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), 3982–3992.
- Taboada M., D. Martinez, M. Arideh, R. Mosquera (2025): Ontology matching with Large Language Models and prioritized depth-first search. Information Fusion, Elsevier, v123, 103254.
- Wu H., M. Li, C. Yu, Z. Ouyang, K.-h. Lai, Z. Zhao, S. Pan, S. Wang, R.Y. Zhong, Y.-H. Kuo, F. Zhang, W. Huang, Z.-J. Max Shen, E. Ballot, G.Q. Huang (2025): Towards cyber-physical internet: A systematic review, fundamental model and future perspectives. Transportation Research Part E: Logistics and Transportation Review, Elsevier, v197, 104051.
- Zacharewicz G., D. Chen, B. Vallespir (2009): Short-lived ontology approach for agent/HLA federated enterprise interoperability. In International Conference on Interoperability for Enterprise Software and Applications China, 329–335.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE
KLU KÜHNE LOGISTICS UNIVERSITY

Design Objectives and Requirements for a PI Dynamic Interoperability Tool

Gero Niemann[1] and Rod Franklin[1]

1. Kühne Logistics University, Hamburg, Germany

Corresponding author: gero.niemann@klu.org

Abstract: *Digital interoperability remains a critical barrier to realizing the Physical Internet (PI) vision, as heterogeneous information systems across logistics stakeholders prevent seamless, dynamic data exchange. Traditional approaches such as static standards and bilateral mappings fail to scale in decentralized networks with diverse, evolving data models and dynamic interoperability requirements. This paper develops design objectives and requirements for a PI Dynamic Interoperability Tool (PIDIT) that leverages large language models (LLMs) for semantic interoperability. Following the echelon-based design science research (eDSR) methodology, we first establish the problem through a literature review and validate it with nine semi-structured expert interviews, identifying 77 problems aggregated into 14 interrelated meta problems. A subsequent review of LLM-based interoperability research across ontology matching, schema matching, semantic translation, and knowledge graph construction confirms that no existing approach fully addresses the PI's runtime interoperability needs. Building on these foundations, we formulate five design objectives and associated design requirements, ranging from semantic understanding of heterogeneous formats to autonomous interoperability at scale with human-in-the-loop oversight. This work provides a rigorously validated problem-requirements foundation for subsequent design and evaluation of the PIDIT.*

Keywords: *Physical Internet, digital interoperability, design objectives, large language model, semantic interoperability*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☐ *PI Networks,* ☐ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☒ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☒ *Flash Video,* ☒ *In-Person presentation*

1 Introduction

In a collaborative supply chain, information systems (IS) of different stakeholders may represent data in different ways. These heterogenous data formats can cause interoperability issues preventing dynamic and seamless shipment flow. The reason for this is that data cannot automatically be processed as a mismatch between formats and semantics makes the data incompatible with different IS. We define dynamic interoperability as the ability of information systems to semantically interpret and process each other's data at runtime, without requiring prior bilateral agreements or pre-configured mappings (Tolk et al., 2013; Wang et al., 2009)

Traditional approaches (static standards or bilateral mappings) struggle to scale in a global, decentralized network with countless independent actors and evolving data needs. Even though

standardization increases scalability, it does not solve the m + n mapping problem (Wegner, 1996). In other words, the standard is usually not implemented as the IS data model, local mappings from the IS to the standard are still required.

A logistics paradigm that necessitates digital interoperability is the Physical Internet (PI) (Pan et al., 2021). One of its features is the dynamic and seamless flow of shipments through an interconnected and open network-of-networks (Dong & Franklin, 2021; Montreuil, 2011). In practice, logistics information exchange is characterized by heterogeneous IS, partner specific field usage, and fragmented implementation of standards, resulting in recurring integration frictions and slow onboarding (Ballot et al., 2021; Pan et al., 2021). This gap is particularly problematic for the PI, whose value proposition depends on dynamic collaboration, frequent partner changes, and scalable onboarding.

Prior work has made progress on the technical dimension of this challenge. PI protocols such as the Physical Internet Protocol (PIP) enable shipment-related messages to be transmitted between any two nodes in the network, much like TCP/IP enables interconnection on the DI (Niemann et al., 2025; Niemann & Franklin, 2024; Wu et al., 2025). However, technical interconnectivity alone does not ensure that the receiving system can interpret and use the data it receives. In practice, logistics actors operate highly customized IS with heterogeneous data models, diverse messaging standards (EDIFACT, GS1 EPCIS, IATA ONE Record, DCSA), and proprietary ERP configurations (Debicki & Guzman, 2020; Pan et al., 2021). Even when organizations adopt a common technical interface such as REST over HTTP with JSON payloads, engineers must still reconcile the structure, naming, and meaning of each data element before back-end systems can process the information (Falcão et al., 2025).

The result is a persistent gap at the level of dynamic interoperability. For the PI, where countless independent actors must collaborate dynamically and where business relationships can be short-lived, this gap represents a fundamental obstacle. It prevents the network from achieving the seamless, ad-hoc collaboration that the PI paradigm demands.

Recent advances in large language models (LLMs) offer a promising pathway to address this gap. LLMs can interpret the conceptual meaning of data regardless of format or domain-specific representation, thereby overcoming syntactic and terminological differences that traditionally require costly manual integration logic (Abu-Salih et al., 2025; Xia et al., 2024). When orchestrated as multi-agent systems, they can jointly leverage natural language understanding, code generation, and tool use to automate semantic tasks such as schema matching and data translation across diverse interfaces (Marro & Torr, 2025; Schick et al., 2023).

This paper addresses this gap by developing design objectives and requirements for a PI Dynamic Interoperability Tool (PIDIT) that leverages LLMs for semantic interoperability in the PI. Following the echelon-based Design Science Research (ᵉDSR) methodology of Tuunanen et al. (2024), we report on the first two design echelons: problem analysis and objectives and requirements definition. In Echelon 1, we combine a literature review with nine semi-structured expert interviews with practitioners from major logistics service providers and IT consulting firms to validate and articulate the problem. In Echelon 2, we derive five Design Objectives with sixteen associated Design Requirements grounded in practitioner needs extracted from interviews and academic literature that are validated by a focus group. Together, these echelons produce a causally mapped problem statement and a set of stakeholder-validated design objectives as a validated intermediate artifact that provide a rigorous foundation for the subsequent design and development of the PIDIT.

2 Research Design

The main objective of this research project is to develop a PI Dynamic Interoperability Tool (PIDIT). To achieve this goal, we conduct an eDSR project (Tuunanen et al., 2024) in an iterative process through echelons. This paper focuses exclusively on the first two echelons, namely problem analysis and objectives and requirements definition. The eDSR methodology supports this scope of reporting, as each echelon produces a distinct, validated intermediate artifact that constitutes a self-contained contribution to design knowledge and can be communicated independently (Tuunanen et al., 2024). Accordingly, the problem analysis echelon yields a validated problem statement, and the objectives and requirements definition echelon yields validated Design Objectives. Subsequent echelons (design and development, demonstration, and evaluation) will be addressed in future work. For validation, we use recommended techniques of literature review, expert interviews, focus group discussion and logical reasoning.

3 DSR Echelons

3.1 Problem statement

The eDSR approach is based on the identification and validation of a real-world problem (Tuunanen et al., 2024). Consequently, in Echelon 1, a review of academic and practitioner literature on digital interoperability between IS was conducted to establish the problem. Therefore, we used targeted searches on databases such as Google Scholar and Scopus using combinations of terms such as “digital interoperability”, "data integration," "heterogeneous systems", "information systems interoperability”.

Our analysis indicated that a critical gap lies at the intersection of the PI vision and semantic interoperability. As the PI is based on seamless and dynamic collaboration, interoperability is a major factor for the establishment of a PI (Ballot et al., 2021; Pan et al., 2021; Wu et al., 2025). This raises the following question: How can we ensure that any two foreign PI nodes not only exchange data (as PIP enables technical interoperability) but also immediately understand each other’s data (Mohammadi et al., 2020; Montreuil et al., 2012)? This is a question of dynamic interoperability. Traditional approaches (static standards or bilateral mappings) struggle to scale in a global, decentralized network with countless independent actors and evolving data needs (Debicki & Guzman, 2020; Falcão et al., 2025; Stegemann & Gersch, 2019; Zhao & Xia, 2014). LLM-based adapters offer a potential solution (Abu-Salih et al., 2025; Maciel et al., 2024; Xia et al., 2024), but it is not yet clear how to design, implement, and integrate such AI-driven components within the PI architecture. There is a need for rigorous research to formulate and validate an approach that leverages LLMs for semantic interoperability in logistics, while accounting for the technical, organizational, and economic complexities of the domain.

In line with Tuunanen et al. (2024) we validate these insights with 9 semi-structured interviews with experts working for major logistic service providers (LSPs) in Germany and logistics-IT related consulting firms. As emphasized in qualitative research literature (e.g., Guest et al., 2006) saturation can often be reached within six to twelve interviews, particularly when engaging domain experts. These online interviews were recorded, transcribed and analyzed. They confirm the issue of interoperability as a major obstacle to enable automated, dynamic and seamless processes for interorganizational system communication.

The participants highlighted 77 problems across the 9 interviews. We aggregated these problems into 14 meta problems (MPs) that, when analyzed, showed that some of the meta problems were interrelated and following a causal relationship. Thus, we decided to map these meta problems into five categories, root causes, that lead to structural, technical and operational problems and finally result in outcomes that hinder dynamic interoperability.

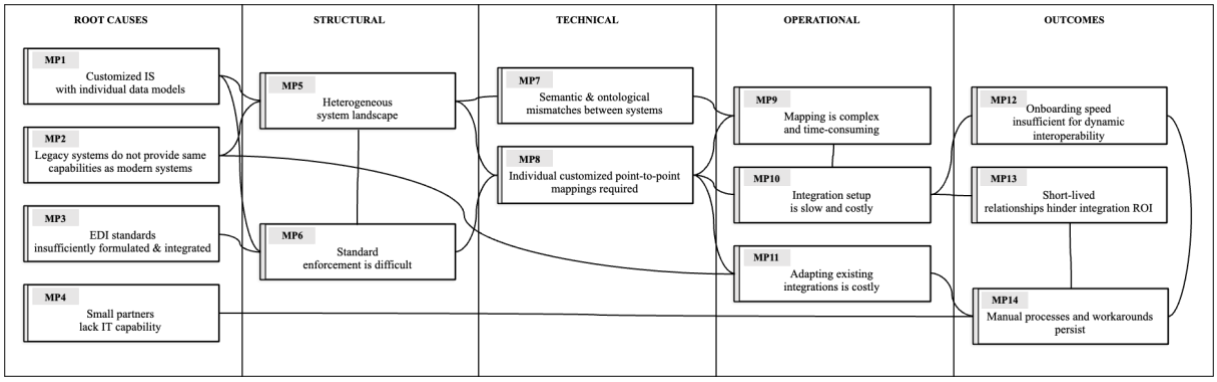


Figure 1: Dynamic interoperability problem mapping

One of the root causes of dynamic interoperability that was repeatedly extracted from the interviews refers to the level of customization of IS amongst LSPs and their related data models (MP1). As mentioned by the head of business excellence of a major contract logistics LSP: "[…] every company operates differently. The main issue is that we all use different ERP systems. Even if we were to use the same provider, such as SAP, Manhattan, or any of the other major players, […] every company has its own custom configurations […] there is no standardized ERP system, frontend, or backend.". In addition to this, some IS are legacy systems not even providing the same technical features as up-to-date systems (MP2) further hindering integration efforts as commonly used interfaces are not supported.

While standardization bodies have put efforts into unifying EDI messages, they do not fully address the problem of dynamic interoperability. Those messaging standards are insufficiently formulated and not naturally integrated into IS (MP3) as one interviewee indicated: "[…] anyone who has already worked with it more intensively knows this too, that the standard stops at the field designation. Thus, they do not prevent manual mappings from source formats to target formats as indicated by another interviewee: "It's like a letter template, there's always a standard letterhead, with predefined fields for the sender and recipient at the beginning and end. However, the content inside the letter is always different […]".

However, a general prerequisite for having to deal with these issues is having the necessary IT capabilities, which bigger players have but smaller ones often lack (MP4), leaving these bigger players with the challenge of dealing with manual processes when they do business with many small carriers (MP14)

These root causes lead to structural, technical and operational problems where heterogenous systems cannot dynamically interoperate without significant effort (MP5) (Pan et al., 2021; Tran-Dang & Kim, 2021). Additionally, further standardization is difficult (MP6) as existing standards generally define a framework and are rarely detailed. Thus, players adapt the standards to their needs. When interconnecting with others each player brings its own adapted standards and integration tools, requiring bilateral alignment before data exchange, which limits scalability. API-based approaches are more scalable and becoming more popular but operating at a lower degree of detail compared to EDI messages and still not addressing semantic and ontological mismatches between heterogenous IS (MP7) (Frendi et al., 2024; Hofman et al., 2016). Thus, individual point-to-point mappings are required (MP8) which is complex and time consuming (MP9). An interviewee describes the mapping of internal data structures to a common standard or another player's profile as a "real challenge" which, besides semantic and ontological heterogeneity, is also affected by country specific regulations. Conducting these complex manual mapping processes are slow and costly (MP10) and incur follow up costs when existing integrations must be updated (MP11). Interviewees describe that integrations take

weeks to months and involve significant costs, including consultancy mapping projects and recurring maintenance costing hundreds of thousands of euros per year (Hofman et al., 2016). Interviewees describe IT integration costs as unwanted overhead that triggers negotiation and requires organizational resource decisions.

These interrelated meta problems lead to three serious business outcomes. First, the onboarding speed is insufficient for the demand (MP12) as integration lead times make it difficult to react to customer demands, especially for temporary or short-notice connections. Using large EDI middleware parties makes setup too slow for agile warehouse switching. Interviewees describe the current integration approach as not scalable and not flexible and explicitly call system connectivity a hurdle that needs to be automated.

Second, when business relationships are short-lived or temporary, the investment in system integration does not pay off (MP13) For example, freight forwarding is described as short-lived at times. An interviewee described: "[…] we had cases where we onboarded 4 new customers within a single month, some of whom were only there for 8 to 12 weeks because it was spot business. In those cases, it was typically the situation that the integration wasn't even worth the effort, and the matters were handled via an Excel spreadsheet." This already emphasizes the third business outcome, the persistence of manual processes and workarounds (MP14).

Without system integration, data exchange relies on manual processes: receiving documents by email and manually entering data into systems, managing via spreadsheets, manually updating shipment statuses, manually onboarding suppliers, and manually exchanging customs data. Interviewees describe this as inefficient and not appropriate for current times.

After getting a comprehensive and practically grounded overview about the problem, we assessed the extent to which the problem has been solved in the literature (Tuunanen et al., 2024). Again, we conducted a literature review on databases such as Google Scholar and Scopus using combinations of terms such as, "large language model", "LLM", "semantic interoperability", "ontology matching" and "schema matching".

Our review revealed that recent LLM-based interoperability research can be organized along four categories of increasing complexity: ontology matching, schema matching, semantic translation, and knowledge graph construction, progressing from identifying semantic correspondences between data models to transforming data while preserving meaning and building shared knowledge infrastructure.[1] Yet existing work remains largely confined to controlled benchmarks, single-domain case studies, or design-time tooling. No existing approach addresses the full problem of dynamic, ad-hoc semantic translation between heterogeneous IS in the logistics domain, where the diversity of standards, the volume of trading partners, and the requirement for runtime operation create a uniquely challenging setting.

Summarizing this, the heterogeneity of customized back-end solutions and data models of IS leads to walled gardens (Marro & Torr, 2025), that prevent efficient and dynamic digital interoperability. Further, this acts as a barrier to the PI, an innovative paradigm that necessitates digital interoperability in all its facets (Pan et al., 2021), blocking its adoption and minimizing its benefits to logistics operators and society.

3.2 Objectives and requirements

This section outlines the objectives, i.e., targets, aimed at addressing the previously stated problem. To validate the intermediate artifact of this echelon we set the following criteria

[1] A structured analysis of the literature is available from the authors upon request but has been omitted here for reasons of space.

according to Tuunanen et al. (2024): fit to the validated problem statement and completeness with respect to the stakeholder requirements. To develop the intermediate artifact of this design echelon, we follow the requirement engineering process of Pandey et al. (2010), including practitioner interviews.

First, we develop system level Design Objectives (DOs) based on semi structured expert interviews with operations and process engineering specialists and academic as well as practitioner literature. Next, we divide each system level DO into several subsystem level Design Requirements (DRs) through logical reasoning, insights from practitioner interviews as well as academic and practitioner literature. We validate our established DRs by verifying that each DR under a DO adds necessary benefit to the DO.

Drawing on literature that describes the importance of digital interoperability for the PI (Pan et al., 2021) and conceptualization studies proposing optional protocol frameworks (Dong & Franklin, 2021; Montreuil et al., 2012; Wu et al., 2025) we believe that the implementation of a PI, built on seamless interconnection and collaboration, is heavily supported by building dynamic digital interoperability into the PIs' operational framework. This DO comes with three DRs. First, to embed dynamic digital interoperability into the PI layer model it must be decided at which layer digital interoperability should be addressed. As already addressed earlier in this paper, standardization is hard to achieve. Thus, a tool is required that translates between different system languages. Such a tool should be positioned on the application layer, directly interfacing with applications that require data inputs they understand such as a transport management system (TMS) (DR1). Second, the tool should be able to access a PIP message and read the information carried in its body (Niemann et al., 2025; Niemann & Franklin, 2024; Wu et al., 2025) (DR2). These are the input information to the tool that must be transformed to target system understandable information. And third, our literature review revealed, that LLMs are getting increasing popularity in supporting translation of data from heterogenous sources and thus supporting mapping tasks (DR3).

They can interpret the conceptual meaning of data regardless of its format or domain-specific representation, thereby overcoming syntactic and terminological differences that traditionally require costly manual integration logic (Abu-Salih et al., 2025; Xia et al., 2024). When orchestrated as multi-agent systems, LLMs can jointly leverage natural language understanding, code generation, and tool use to automate semantic tasks such as alignment and translation across diverse interfaces, representing a paradigm shift from rigid rule-based transformation toward flexible, meaning-driven interoperability (Abu-Salih et al., 2025; Marro & Torr, 2025; Schick et al., 2023).

The expert interviews revealed four recurring themes that shaped the remaining design objectives. Practitioners consistently described partner connectivity as a scalability problem rather than a technical one: the real burden is not connecting two systems, but managing the combinatorial growth of mappings as the partner network expands. A shared semantic layer via an intermediary ontology, rather than point-to-point integrations, emerged as the preferred architectural response (DO2). Automation ambitions were universally qualified by a reluctance to relinquish control. Interviewees were willing to let a tool generate mappings, but only if human validation remained mandatory before operational use, and they rated ease of correction and speed of map development as more important than accuracy from day one (DO3). Integration costs were framed as a two-sided problem: beyond the well-known burden of labor-intensive mapping work, interviewees flagged the recurring inference cost of LLM calls at runtime as an equally relevant concern. This led directly to the requirement for deterministic code generation, where the LLM is invoked once to produce a reusable mapping that is

subsequently executed without further inference (DO4). Table 1 details the full set of design requirements derived from these themes.

Looking further ahead, interviewees envisioned the tool eventually operating autonomously but emphasized that this is conditional on a quality and trust threshold the tool must first earn. Before human oversight can be safely reduced, the tool must demonstrate high consistency, detect anomalies in stable data streams, incorporate an error-feedback loop for continuous improvement, and meet data security requirements for sensitive data types. This forward-looking perspective, qualitatively distinct from the immediate practitioner pain points driving DO2 through DO4, is captured in DO5. Together, DO1 through DO5 and their associated design requirements are fully specified in Table 1.

Table 1: Design objectives and requirements for the PIDIT

ID	DO	DRs	Justification knowledge
DO1	Facilitating a PI by embedding dynamic digital interoperability into its operational framework (layered model)	DR1: It should be a tool positioned at application layer that interacts with different layers and their internal protocols	PI conceptualization literature (e.g. Dong & Franklin, 2021; Montreuil et al., 2012b; Pan et al., 2021; Wu et al., 2025)
		DR2: It should read heterogenous information carried in the body of a PIP message and make them universally understandable	PI protocol literature (Niemann et al., 2025; Niemann & Franklin, 2024; Wu et al., 2025)
		DR3: It should use the semantic capabilities of LLMs to map heterogenous data formats	LLM interoperability literature (e.g. Abu-Salih et al., 2025; Calvet et al., 2024; Falcão et al., 2025; Marro & Torr, 2025; Xia et al., 2024)
DO2	Simplify, extend and make connectivity across partners more scalable	DR4: It should handle heterogeneous formats and system diversity	Interviews with practitioners #4, #6, #7, #8, #9
		DR5: It should translate via an intermediate ontology	Interoperability literature (Wegner, 1996)
		DR6: It should be broadly applicable and generalizable across logistics	Focus group discussion
DO3	Accelerate partner and customer onboarding by reducing integration time and effort	DR7: It should make mapping proposals that must be validated by humans	Interviews with practitioners #3, #5, #6, #7, #8, #9
		DR8: It should provide a uniform mapping documentation	Focus group discussion
		DR9: It should classify documents and extract fields	
		DR10: Mappings and error diagnosis should be user-friendly readable	
DO4	Lower integration costs and investment barriers	DR11: It should require low labor for mapping and maintenance	Interviews with practitioners #4, #5, #9
		DR12: It should generate deterministic code for runtime operation	LLM translator prototypes (Calvet et al., 2024; Falcão et al., 2025)
			Focus group discussion

DO5	Enable automation and autonomous operations	DR13: It should provide high output quality, consistency, and accuracy	Interviews with practitioners #4, #5, #6, #8, #9
		DR14: It should detect anomalies and errors and autonomously correct those	Focus group discussion
		DR15: It should be trustworthy in terms of data privacy, security, and sensitivity handling	
		DR16: It should have domain knowledge of source and target IS	

To validate completeness of the DOs with respect to the validated problem statement, we trace each DO to the meta problems it addresses. DO1 responds to the systemic absence of a dynamic interoperability mechanism, as evidenced by the three business outcomes (MP12–MP14). DO2 targets the root causes and structural barriers that produce the heterogeneous, unscalable landscape (MP1–MP7). DO3 addresses the operational friction of slow and complex onboarding (MP8–MP10, MP12, MP14). DO4 responds to the economic burden of integration costs and maintenance (MP10, MP11, MP13). DO5 envisions the autonomous end-state that resolves the persistence of manual processes (MP12–MP14) as well as the cost and time issues (MP9-MP11). This mapping confirms that all fourteen meta problems are addressed by at least one design objective, satisfying the completeness criterion.

Additionally, we validated the DOs and DRs in a focus group discussion with three data integration experts from a major German logistics service provider. We walked the participants through the DOs and DRs and discussed if and how the DRs support the DOs. The session was recorded and transcribed. The focus group participants confirmed the DOs and the supporting role of the underlying DRs. However, they indicated as an addition, that the PIDIT should possess domain knowledge of the source and target systems to enable correct and autonomous translation. This is an important and valuable addition that we did not capture from the prior interviews, which is why we added it as DR16 in support of DO5.

4 Conclusion

This paper developed design objectives and requirements for PIDIT a technical system proposed that leverages the capabilities of LLMs to address semantic interoperability challenges in the PI. Following the [e]DSR methodology, we first established the problem through a literature review and validated it with nine semi-structured expert interviews, yielding 77 problems aggregated into 14 interrelated meta problems. These revealed that customized IS, insufficient standardization, and insufficient IT capabilities create heterogeneous system landscapes that lead to slow onboarding, unprofitable short-lived integrations, and persistent manual data handling. A subsequent review of LLM-based interoperability research across ontology matching, schema matching, semantic translation, and knowledge graph construction confirmed that no existing solution fully addresses the dynamic, runtime interoperability needs of the PI. Building on these insights, we formulated five DOs and associated DRs, grounded in both practitioner needs and academic literature. The objectives range from enabling semantic understanding across heterogeneous data formats to achieving autonomous, dynamic interoperability at scale while maintaining human oversight. This work contributes a rigorously validated problem-requirements foundation for the subsequent [e]DSR echelons in which the PIDIT will be designed, instantiated, and evaluated in collaboration with a major German logistic service provider. Future research will focus on prototyping the tool and empirically assessing its performance in real-world logistics integration scenarios.

References

- Abu-Salih, B., Alotaibi, S., Alanazi, A. L., Abu Khurma, R., Al-Shboul, B., Khouri, A., & Aljaafari, M. (2025). Using large language models for semantic interoperability: A systematic literature review. *ICT Express*, *11*(4), 819–837. https://doi.org/10.1016/j.icte.2025.06.011
- Ballot, E., Montreuil, B., & Zacharia, Z. G. (2021). Physical Internet: First results and next challenges. *Journal of Business Logistics*, *42*(1), 101–107. https://doi.org/10.1111/jbl.12268
- Calvet, E., Falcão, R., & Ginkel, J. (2024). *Can LLMs Make Systems Interoperate at Runtime? An Exploratory Study*. SSRN. https://doi.org/10.2139/ssrn.4979509
- Debicki, T., & Guzman, C. (2020). Challenges in Future Data Interchange in Transport and Logistics Sector. In A. Kolinski, D. Dujak, & P. Golinska-Dawson (Eds.), *Integration of Information Flow for Greening Supply Chain Management* (pp. 93–107). Springer International Publishing. https://doi.org/10.1007/978-3-030-24355-5_19
- Dong, C., & Franklin, R. (2021). From the Digital Internet to the Physical Internet: A conceptual framework with a stylized network model. *Journal of Business Logistics*, *42*(1), 108–119. https://doi.org/10.1111/jbl.12253
- Falcão, R., Schweitzer, S., Siebert, J., Calvet, E., & Elberzhager, F. (2025). *Evaluating the effectiveness of LLM-based interoperability* (arXiv:2510.23893). arXiv. https://doi.org/10.48550/arXiv.2510.23893
- Frendi, M., Nachet, B., & Adla, A. (2024). Physical Internet Based Ontology for Supporting Traceability in Logistic IoT. *International Journal of Computing and Digital Systems*, *15*(1). https://doi.org/10.12785/ijcds/150133
- Guest, G., Bunce, A., & Johnson, L. (2006). How Many Interviews Are Enough?: An Experiment with Data Saturation and Variability. *Field Methods*, *18*(1), 59–82. https://doi.org/10.1177/1525822X05279903
- Hofman, W., Punter, M., Bastiaansen, H., Cornelisse, E., & Dalmolen, S. (2016). Semantic technology for enabling logistics innovations – towards Intelligent Cargo in the Physical Internet. *International Journal of Advanced Logistics*, *5*(2), 58–69. https://doi.org/10.1080/2287108X.2016.1144358
- Maciel, R. S. P., Valle, P. H. D., Santos, K. S., & Nakagawa, E. Y. (2024). Systems Interoperability Types: A Tertiary Study. *ACM Computing Surveys*, *56*(10), 254:1-254:37. https://doi.org/10.1145/3659098
- Marro, S., & Torr, P. (2025). *LLM Agents Are the Antidote to Walled Gardens* (arXiv:2506.23978). arXiv. https://doi.org/10.48550/arXiv.2506.23978
- Mohammadi, M., Hofman, W., & Tan, Y.-H. (2020). Seamless interoperability in logistics by ontology alignment. *Journal of Supply Chain Management Science*, *1*(3–4), 104–117. https://doi.org/10.18757/jscms.2020.5444
- Montreuil, B. (2011). Toward a Physical Internet: Meeting the global logistics sustainability grand challenge. *Logistics Research*, *3*(2), 71–87. https://doi.org/10.1007/s12159-011-0045-x
- Montreuil, B., Ballot, E., & Fontane, F. (2012). An Open Logistics Interconnection model for the Physical Internet. *IFAC Proceedings Volumes, 14th IFAC Symposium on Information Control Problems in Manufacturing*, *45*, 327–332. https://doi.org/10.3182/20120523-3-RO-2023.00385
- Niemann, G., & Franklin, R. (2024, June 1). A protocol for node-to-node transshipment in the Physical Internet analogizing from Internet principles. *Proceedings of the 10th International Physical Internet Conference*. 10th International Physical Internet Conference, Savannah, Georgia, USA. https://doi.org/10.35090/gatech/11063
- Niemann, G., Franklin, R., & Spanuth, M. (2025, June 1). Addressing the Reachability Problem in the Physical Internet: A PI-Link Protocol Inspired by the Internet. *Proceedings of the 11th International Physical Internet Conference*. 1th International Physical Internet Conference, Hong Kong. https://doi.org/10.35090/gatech/13227
- Pan, S., Trentesaux, D., McFarlane, D., Montreuil, B., Ballot, E., & Huang, G. Q. (2021). Digital interoperability in logistics and supply chain management: State-of-the-art and research avenues

towards Physical Internet. *Computers in Industry*, *128*, 103435. https://doi.org/10.1016/j.compind.2021.103435

- Pandey, D., Suman, U., & Ramani, A. K. (2010). An Effective Requirement Engineering Process Model for Software Development and Requirements Management. *2010 International Conference on Advances in Recent Technologies in Communication and Computing*, 287–291. https://doi.org/10.1109/ARTCom.2010.24
- Schick, T., Dwivedi-Yu, J., Dessi, R., Raileanu, R., Lomeli, M., Hambro, E., Zettlemoyer, L., Cancedda, N., & Scialom, T. (2023). Toolformer: Language Models Can Teach Themselves to Use Tools. *Advances in Neural Information Processing Systems*, *36*, 68539–68551.
- Stegemann, L., & Gersch, M. (2019). Interoperability – Technical or economic challenge? *It - Information Technology*, *61*(5–6), 243–252. https://doi.org/10.1515/itit-2019-0027
- Tolk, A., Bair, L. J., & Diallo, S. Y. (2013). Supporting Network Enabled Capability by extending the Levels of Conceptual Interoperability Model to an interoperability maturity model. *The Journal of Defense Modeling and Simulation*, *10*(2), 145–160. https://doi.org/10.1177/1548512911428457
- Tran-Dang, H., & Kim, D.-S. (2021). The Physical Internet in the Era of Digital Transformation: Perspectives and Open Issues. *IEEE Access*, *9*, 164613–164631. https://doi.org/10.1109/ACCESS.2021.3131562
- Tuunanen, T., Winter, R., & vom Brocke, J. (2024). Dealing with Complexity in Design Science Research: A Methodology Using Design Echelons. *MIS Quarterly*, *48*(2), 427–458. https://doi.org/10.25300/MISQ/2023/16700
- Wang, W., Tolk, A., & Wang, W. (2009). *The Levels of Conceptual Interoperability Model: Applying Systems Engineering Principles to M&S* (arXiv:0908.0191). arXiv. https://doi.org/10.48550/arXiv.0908.0191
- Wegner, P. (1996). Interoperability. *ACM Computing Surveys*, *28*(1), 285–287. https://doi.org/10.1145/234313.234424
- Wu, H., Li, M., Yu, C., Ouyang, Z., Lai, K., Zhao, Z., Pan, S., Wang, S., Zhong, R. Y., Kuo, Y.-H., Zhang, F., Huang, W., Shen, Z.-J. M., Ballot, E., & Huang, G. Q. (2025). Towards cyber-physical internet: A systematic review, fundamental model and future perspectives. *Transportation Research Part E: Logistics and Transportation Review*, *197*, 104051. https://doi.org/10.1016/j.tre.2025.104051
- Xia, Y., Xiao, Z., Jazdi, N., & Weyrich, M. (2024). Generation of Asset Administration Shell with Large Language Model Agents: Toward Semantic Interoperability in Digital Twins in the Context of Industry 4.0. *IEEE Access*, *12*, 84863–84877. https://doi.org/10.1109/ACCESS.2024.3415470
- Zhao, K., & Xia, M. (2014). Forming Interoperability Through Interorganizational Systems Standards. *Journal of Management Information Systems*, *30*(4), 269–298. https://doi.org/10.2753/MIS0742-1222300410

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE

Heterogeneous Graph-Based Risk Assessment for Industrial Chain Cooperation

Zihang Yu[1], Hanlin Tang[2], Dong Liu[3] and Shuangxi Huang[,1]
1. Department of Automation, Tsinghua University, Beijing, China
2. ChinaDataPay, Guiyang, China
3. Yangtze Delta Region Institute of Tsinghua University, Jiaxing, China

Corresponding author: huangsx@tsinghua.edu.cn

Abstract: *In the context of the Physical Internet, industrial chain systems have evolved into complex network structures composed of multiple entities and diverse relations, which makes enterprise risk highly correlated and prone to propagation. However, enterprise risk often lacks explicit and identifiable features, while complex structural dependencies exist among enterprises, which makes risk prediction a challenging task. To address this issue, we propose an enterprise risk assessment framework for industrial chains that integrates knowledge graph construction with relation-aware heterogeneous graph learning. First, we construct a multi-relational enterprise knowledge graph for the electric vehicle industrial chain to capture the complex interconnections among enterprises. Based on this graph, we design a relation-aware heterogeneous graph attention network to learn the importance of different relations in information propagation. We further introduce a gating mechanism to fuse structural representations with enterprise attribute features, thereby obtaining risk-aware enterprise embeddings, which are then fed into a gradient boosting model for multi-risk prediction. This study provides a data-driven approach for enterprise risk identification in complex industrial chain networks.*

Keywords: *Physical Internet, Electric Vehicle Industrial Chain, Industrial Chain Knowledge Graph, Heterogeneous Graph Neural Network, Relation-Aware Graph Attention Network, Enterprise Risk Assessment.*

Physical Internet (PI) Roadmap Fitness*: ☒ PI Networks, ☒ Vertical Supply Consolidation, ☒ Logistics/Commercial Data Platform.*

Targeted Delivery Mode-s*: ☒ Paper, ☒ Poster, ☐ Flash Video, ☒ In-Person presentation*

1 Introduction

Driven by the Physical Internet (PI), modern supply chains have evolved into networked ecosystems where open logistics and information networks enable enterprises to form complex collaborative structures characterized by multi-entity, multi-relational production and investment links. Prior studies show that hidden and complex dependencies in supply chains lead to risk diffusion and propagation across the network (Kosasih et al., 2024).

However, in practical applications, identifying industrial chain risks remains challenging because multi-factor influences (e.g., operational, market, and regulatory) often result in implicit signals without clear patterns. Furthermore, structural dependencies cause risk to propagate through upstream/downstream links, meaning models relying merely on enterprise attributes fail to capture network-driven risk dynamics.

Graph learning and knowledge graphs offer new perspectives for risk analysis. For instance, Graph Convolutional Networks (GCN) proposed by Kipf et al. (2017) efficiently extract relational features, and have been widely applied to capture structural dependencies in enterprise modeling. Meanwhile, the inductive graph representation learning method proposed by Hamilton et al. (2017) addresses the limitations of traditional graph models in handling dynamically evolving networks. This method enables risk assessment for newly emerging enterprises in supply chains and is particularly suitable for monitoring new supporting enterprises in the electric vehicle industrial chain. In addition, knowledge graphs can accurately represent multi-type entities and relations in supply chain systems, such as enterprises, products, technological patents, and cooperation agreements. By combining knowledge graphs with graph learning methods, we can discover and reason over hidden risk relations, which provides a data-driven approach for risk analysis in complex industrial networks (Hogan et al., 2021).

Despite these advances, existing methods still face limitations in modeling complex industrial relationships. Some studies consider only a single type of enterprise relation, which cannot fully capture the multi-relational structure of real-world industrial systems. On the other hand, some graph learning methods treat different relations uniformly when modeling multi-relational networks, making it difficult to reflect their heterogeneous roles in information propagation.

To address these issues, we propose an enterprise risk assessment framework for industrial chains that integrates knowledge graph construction with relation-aware heterogeneous graph learning. We first construct a multi-relational enterprise knowledge graph for the electric vehicle industrial chain to model complex industrial relationships among enterprises. Based on this graph, we design a relation-aware heterogeneous graph attention model. We use relation-level attention to learn the importance of different relations in information propagation, and we introduce a gating mechanism to fuse graph structural representations with enterprise attribute features. In this way, we obtain risk-aware enterprise embeddings, which are fed into a gradient boosting model for enterprise risk prediction.

The rest of this paper is organized as follows. Section 2 reviews related work. Section 3 presents the proposed method. Section 4 conducts experiments and analyzes results using the case of electric vehicle industrial chain. Section 5 concludes and discusses future research directions.

2 Literature Review

2.1 Industrial Chain Risk Assessment

Industrial chain enterprise risk assessment has long been a central topic in supply chain management and enterprise management research. Traditional studies focus on individual enterprises and build statistical or machine learning models based on macroeconomic variables to identify risks. However, increasing inter-firm complexity facilitates risk propagation across supply, collaboration, and capital networks. When risk events occur at key nodes, their impact can spread along the network and affect overall system stability (Craighead et al., 2007). Kim et al. (2015) construct supply chain network models to study structural dependencies among enterprises and show that an enterprise's position in the network significantly influences its risk exposure. Furthermore, some studies focus on risk propagation mechanisms in supply chain networks. Sodhi and Tang (2012) point out that risks originate not only from internal enterprise factors but also propagate through supply relationships in a cascading manner.

Despite these contributions, most studies still focus on macro-level supply chain structures or single types of relationships. They provide limited modeling of multi-relational structures among enterprises in industrial chains. In addition, traditional risk assessment methods

emphasize enterprise-specific features and underutilize industrial network structure information, which limits risk identification performance.

2.2 Knowledge Graph for Industrial Network Modeling

Knowledge graphs (KGs) effectively model complex entity-relation systems by integrating heterogeneous data under a unified semantic framework. In the KG domain, many studies focus on constructing and representing large-scale relational data. For example, Dong et al. (2014) propose the Knowledge Vault system, which automatically builds large-scale KGs from multi-source data and evaluates fact credibility through probabilistic reasoning, enabling automated construction of structured knowledge. In addition, Nickel et al. (2016) systematically study relational learning methods in KGs and show that graph-based modeling effectively captures semantic relationships among entities.

In recent years, researchers have used KGs to describe multiple types of relationships among enterprises, such as collaboration, supply, and upstream–downstream dependencies. By modeling enterprise entities and their relations as graph structures, we can better capture the structural characteristics of industrial networks. Ji et al. (2022) show that KGs provide a unified data representation framework for complex relational systems and demonstrate strong scalability across multiple domains.

However, KGs mainly provide structured representations of relational data and do not directly model information propagation in complex networks. Therefore, relying solely on KG structures cannot fully uncover latent patterns in enterprise networks. To address this limitation, recent studies combine KGs with graph learning methods. These approaches learn representations from graph structures and better support analysis tasks in complex relational networks, such as enterprise relationship modeling and risk identification.

2.3 Graph Neural Networks for Risk Prediction

With the development of graph representation learning, graph neural networks (GNNs) have become powerful tools for analyzing complex network data. GNNs jointly leverage node attributes and network structure information and achieve strong performance in tasks such as node classification, link prediction, and anomaly detection.

In graph learning research, several foundational methods support modeling of complex relational networks. For example, Schlichtkrull et al. (2018) propose relational graph convolutional networks (R-GCN), which handle graphs with multiple relation types and are widely used in knowledge graph reasoning and relation prediction. Velickovic et al. (2018) propose graph attention networks (GAT), which use attention mechanisms to weight neighbor information during aggregation and improve node representation learning. In addition, heterogeneous GNNs have recently gained attention. By modeling different types of nodes and relations, these methods better handle complex relational networks (Wang et al., 2019).

In the risk identification domain, researchers have begun to apply GNNs to financial risk analysis and anomaly detection. For example, Dou et al. (2020) propose a GNN-based anomaly detection method that identifies potential risk nodes by learning structural features in financial networks. Some studies also apply graph learning methods to enterprise relationship analysis to uncover latent risk patterns among enterprises. However, in real-world industrial chains, enterprises often connect through multiple types of relationships, which may play significantly different roles in information propagation. Therefore, learning relation semantics in multi-relational industrial networks for enterprise risk identification remains a research challenge.

Despite these advances, several gaps remain. First, risk studies often focus on macro-structures rather than multi-relational enterprise networks. Second, KG structural information is underutilized in risk identification. Third, existing GNN methods lack the ability to distinguish semantic differences between relation types.

3 Methodology

3.1 Research Framework

To bridge these gaps, we propose an overall framework for enterprise risk assessment in the EV industrial chain, as illustrated in Figure 1. The framework consists of four core modules. First, we integrate enterprise basic information, inter-firm relationship data, and business scope text to construct an industrial chain KG that includes multiple types of entities, including enterprises, products, and categories, along with their relations. Next, we build a heterogeneous graph based on the KG. We apply a relation-aware graph neural network to learn structural representations of enterprise nodes, and we introduce a gating mechanism to fuse structural features with enterprise attribute features, which yields risk-aware enterprise embeddings. Finally, we feed the learned embeddings into a risk prediction model to identify and assess enterprise risk states.

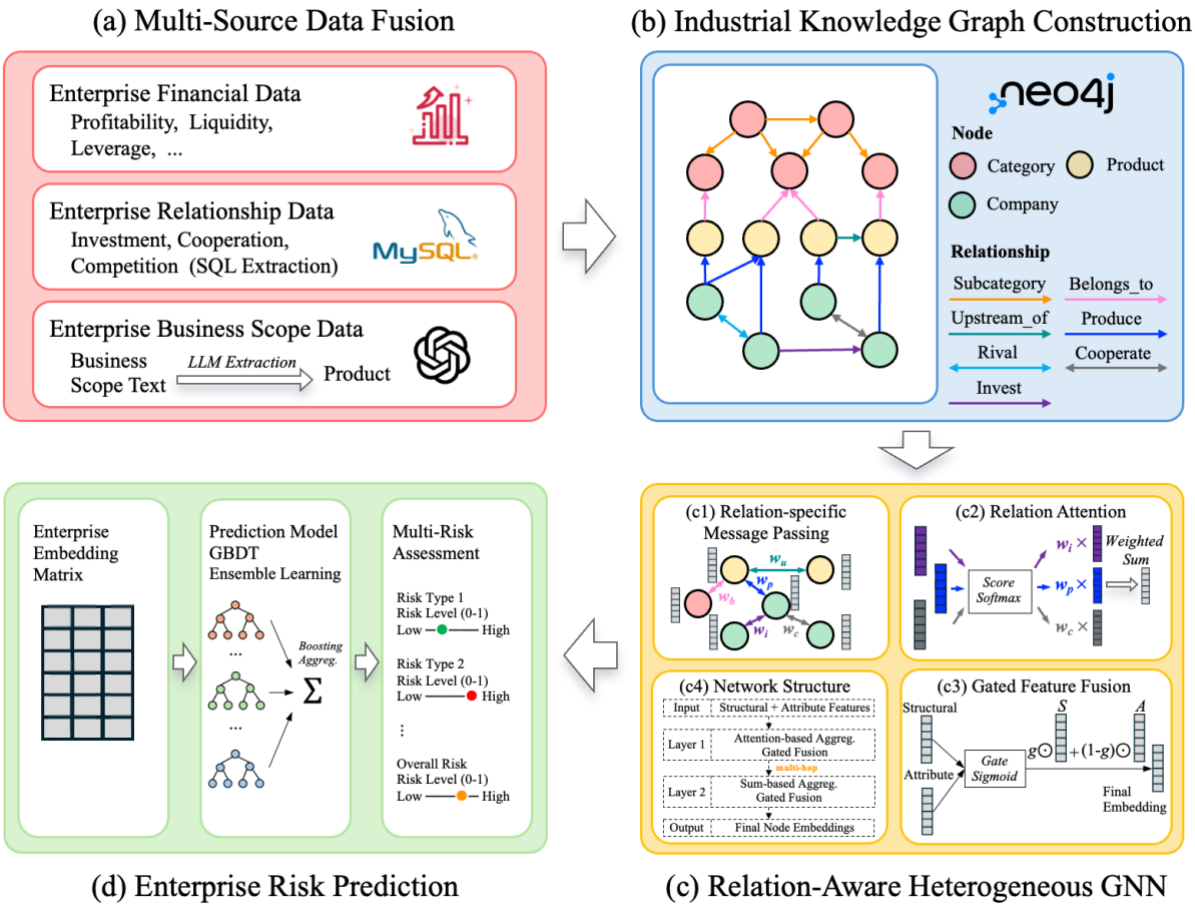


Figure 1: Research Framework

3.2 EV Industrial Chain Knowledge Graph Construction

We construct the EV industrial KG through the following steps to provide structured input for graph learning:

· **Enterprise Data Integration**: We select A-share listed enterprises in the EV industrial chain and construct the enterprise node set. We obtain enterprise basic information and financial features from enterprise databases to characterize enterprise operational status.

· **Enterprise Relationship Data Acquisition**: Inter-firm relationships (investment, cooperation, and competition) are extracted from databases via SQL queries.

· **Product and Industrial Knowledge Extraction**: To further model production structure relationships in the industrial chain, we use large language model to extract product information on business scope text and identify product entities, as shown in Figure 2, and then normalize the extracted results to unify different textual expressions into consistent products. We construct enterprise–product production relationships and build product–product upstream and downstream relations according to industrial structure knowledge. In addition, to represent hierarchical industrial structures, we define product category nodes and establish product–category as well as category–category hierarchical relations.

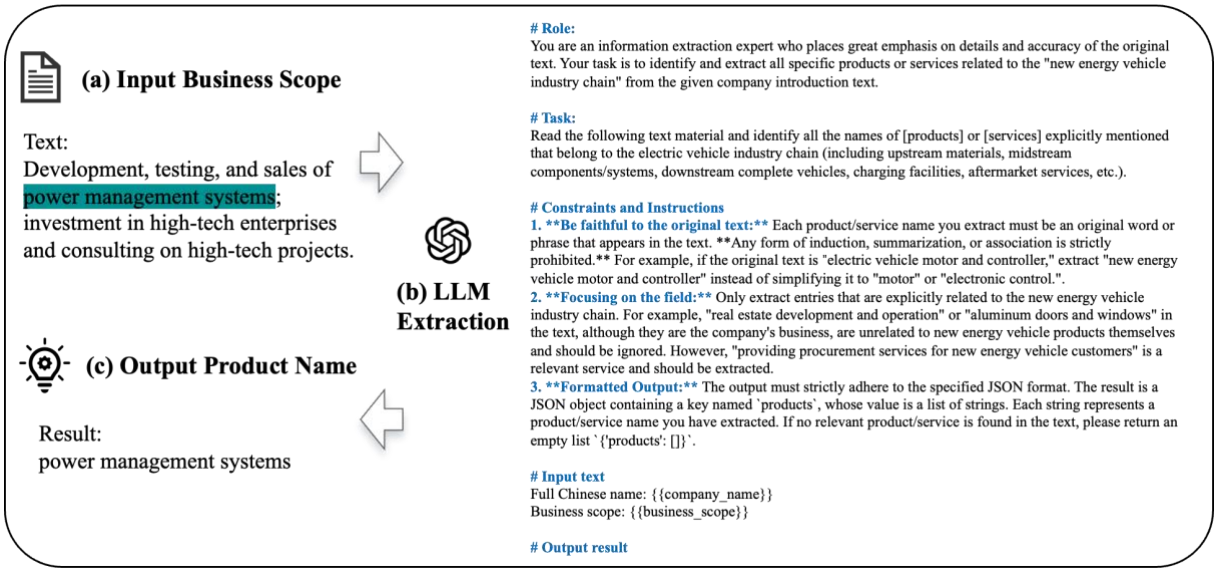


Figure 2: Product Extraction Based on Large Language Models

· **Industrial Chain Knowledge Graph Modeling and Storage**: Extracted entities and relations are integrated into a multi-relational graph. The KG is imported into Neo4j for storage and visualization, as shown in Figure 3. See Appendix 1 for subgraph.

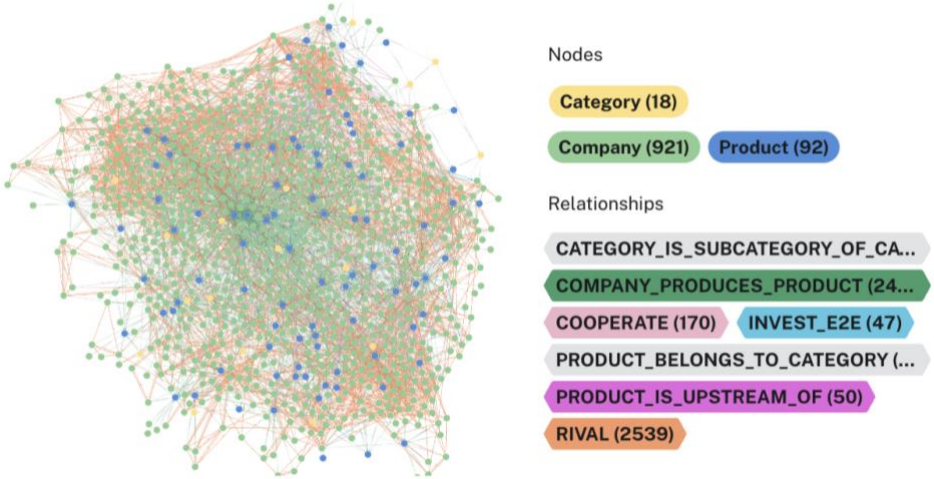


Figure 3: Industrial Chain Knowledge Graph

3.3 Relation-Aware Heterogeneous Graph Neural Network

To learn structural representations of enterprises in complex industrial networks, we design a relation-aware heterogeneous GNN based on the constructed industrial chain KG. The model captures structural semantics of different relation types through relation-specific message passing and learns the importance of each relation in information aggregation via a relation-level attention mechanism. In addition, we introduce a gating mechanism to fuse structural representations with enterprise attribute features, which yields embeddings that reflect enterprise risk states.

3.3.1 Heterogeneous Graph Representation

We represent the constructed industrial chain knowledge graph as a heterogeneous graph:

$$G = (V, E, R) \tag{1}$$

Where V denotes the node set, E denotes the edge set, and R denotes the set of relation types. Each node $v_i \in V$ has a feature vector $\mathbf{x}_i \in \mathbb{R}^d$, where d is the feature dimension.
Edges are represented as (v_i, r, v_j), where $r \in R$ denotes the relation type, including inverse relation r^{-1} to enable bidirectional information flow.

3.3.2 Relation-Specific Message Passing

Since different relations carry distinct semantics in industrial chains, we adopt a relation-specific message passing mechanism to model each relation type. For a given relation r, we update the representation of node i at layer l as follows:

$$\mathbf{h}_i^{(l,r)} = \sigma\left(\sum_{j \in \mathcal{N}_i^r} \alpha_{ij}^{(r)} \mathbf{W}_r^{(l)} \mathbf{h}_j^{(l-1)}\right) \tag{2}$$

Where $\mathbf{h}_i^{(l,r)}$ denotes the representation of node i at layer l under relation r, and $\mathcal{N}_i^r$ denotes the set of neighboring nodes connected to node i via relation r. $\mathbf{W}_r^{(l)}$ is the transformation matrix associated with relation r, and $\alpha_{ij}^{(r)}$ represents the importance weight of neighbor node j. $\sigma(\cdot)$ is a nonlinear activation function. By introducing relation-specific transformation matrices, we enable the model to learn heterogeneous characteristics of different relations during information propagation.

3.3.3 Relation-Aware Attention Mechanism

We further fuse information from multiple relations. To achieve this, we introduce a relation-aware attention mechanism to perform weighted aggregation over relation-specific representations. We first compute the importance score of each relation r:

$$e_r = \boldsymbol{\alpha}^T \tanh \mathbf{W}_a \mathbf{h}_i^{(r)} \tag{3}$$

Where e_r denotes the importance score of relation r, $\mathbf{W}_a$ is the attention transformation matrix, and $\boldsymbol{\alpha}$ is the attention weight vector. We then normalize the scores using the Softmax function to obtain relation weights: $\beta_r = \frac{\exp(e_r)}{\sum_{r' \in R} \exp(e_{r'})}$. Finally, we compute the aggregated node representation as: $\mathbf{h}_i = \sum_{r \in R} \beta_r \mathbf{h}_i^{(r)}$. This mechanism enables the model to automatically learn the importance of different relations in enterprise representation learning and better capture structural characteristics of the industrial chain.

3.3.4 Gated Feature Fusion

To jointly leverage structural information and enterprise attribute features, we introduce a gating mechanism to fuse these two sources of information. Let $\mathbf{h}_i$ denote the structural representation learned from the graph, and let $\mathbf{x}_i$ denote the original attribute features of enterprise. We first compute the fused representation: $\mathbf{z}_i = \mathbf{W}_z[\mathbf{h}_i; \mathbf{x}_i]$. Then compute the gating weights: $\mathbf{g}_i = \sigma(\mathbf{W}_g \mathbf{z}_i)$. The final enterprise representation is defined as:

$$\mathbf{h}_i^* = \mathbf{g}_i \odot \mathbf{h}_i + (1 - \mathbf{g}_i) \odot \mathbf{x}_i \quad (4)$$

Where $\odot$ denotes element-wise multiplication, and $\mathbf{g}_i$ represents the gating weights. This mechanism adaptively balances structural and attribute information based on the data.

We adopt a two-layer relation-aware heterogeneous GNN. The first layer performs relation-aware aggregation with attention to capture heterogeneous relation semantics, while the second layer adopts a sum-based aggregation strategy to improve representation stability. A gating mechanism is applied in each layer to fuse structural and attribute features. The two-layer design enables effective multi-hop information propagation while mitigating over-smoothing.

Finally, we obtain embeddings for enterprise nodes, denoted as $\mathbf{e}_i \in \mathbb{R}^k$, where k is the embedding dimension. These embeddings jointly capture structural positions, inter-firm relationships, and operational characteristics of enterprises. We then use the learned embeddings as input to the enterprise risk prediction model for subsequent risk assessment tasks.

3.4 Enterprise Risk Prediction

We further build an enterprise risk prediction model to identify and assess multi-dimensional risks. Since enterprise risk prediction typically involves complex nonlinear patterns and imbalanced data distributions, we adopt a gradient boosting decision tree (GBDT) model to leverage its strengths in nonlinear modeling and small-sample learning.

3.4.1 Risk Prediction Framework

Given the enterprise embedding $\mathbf{e}_i \in \mathbb{R}^k$, we formulate enterprise risk prediction as a supervised learning problem: $y_i = f(\mathbf{e}_i)$. We model the risk function $f(\cdot)$using a gradient boosting decision tree. Specifically, we construct the prediction function as an additive ensemble of decision trees:

$$\hat{y}_i = \sum_{m=1}^{M} f_m(\mathbf{e}_i), f_m \in \mathcal{F} \quad (5)$$

Where M denotes the number of trees, $f_m(\cdot)$denotes the m-th regression tree, and $\mathcal{F}$denotes the function space of all possible trees. We train the model by optimizing objective function below:

$$\mathcal{L} = \sum_{i=1}^{N} l(y_i, \hat{y}_i) + \sum_{m=1}^{M} \Omega(f_m) \quad (6)$$

$$\Omega(f) = \gamma T + \frac{1}{2}\lambda \|\mathbf{w}\|^2 \quad (7)$$

Where $l(\cdot)$ denotes the loss function (e.g., binary cross-entropy), and $\Omega(f_m)$ is the regularization term controlling model complexity. T denotes the number of leaf nodes, $\mathbf{w}$ denotes leaf weights, and γ and λ are regularization coefficients. This formulation captures nonlinear relationships in the embedding space and improves risk prediction performance.

3.4.2 Multi-Risk Prediction Strategy

In real-world scenarios, enterprises often face multiple types of risks, such as financial risk, operational risk, and credit risk. To address this, we formulate enterprise risk prediction as a multi-task learning problem and predict different risk types separately.

For enterprise i, we define a risk vector with K types: $\mathbf{y}_i = [y_i^{(1)}, y_i^{(2)}, \cdots, y_i^{(k)}]$ and the corresponding predictions: $\hat{\mathbf{y}}_i = [\hat{y}_i^{(1)}, \hat{y}_i^{(2)}, \cdots, \hat{y}_i^{(k)}]$. For each risk type k, we train an independent prediction function: $\hat{y}_i^{(k)} = f^{(k)}(\mathbf{e}_i)$. We define the overall optimization objective as:

$$\mathcal{L} = \sum_{k=1}^{K} \sum_{i=1}^{N} l(y_i^{(k)}, \hat{y}_i^{(k)}) \quad (8)$$

This strategy allows us to model the heterogeneity across different risk types and improves both accuracy and stability in multi-risk prediction.

4 Experiments

4.1 Dataset and Experimental Setup

Following the data collection and processing procedures described before, we obtain data of A-share listed companies in the EV industrial chain of the past five years, as shown in Table 1.

Table 1: Statistics of the Heterogeneous Graph

Node	Count	Relationship	Count
Company	921	Invest (Company-Company)	47
Product	92	Cooperate (Company-Company)	170
Category	18	Rival (Company-Company)	2539
——	——	Produce (Company-Product)	2403
——	——	Upstream_of (Product -Product)	50
——	——	Belongs_to (Product-Category)	92
——	——	Subcategory_of (Category-Category)	13

We split enterprise nodes into training, validation, and test sets with a ratio of 7:1:2. To improve robustness and reliability, we adopt five-fold cross-validation.

We obtain risk labels from various risk events recorded in enterprise databases. The label distribution shows class imbalance across different risk types. The positive sample ratios for four risk categories—operational compliance risk, administrative and regulatory risk, judicial risk, and financial risk—are approximately 26%, 12%, 30%, and 20%, respectively. To address the imbalance issue, we evaluate model performance using multiple metrics: AUC, AP, Recall@K, Lift@K. We use Top-K metrics (e.g., K=50) to simulate real-world risk screening scenarios, which better reflect practical application value. The detailed experimental settings are listed in Table 2.

Table 2: Experimental Settings

Parameter	Value	Parameter	Value
Embedding Dimension	64	Optimizer	Adam
Learning Rate	0.001	Dropout	0.5

Epochs	100	Weight Decay	1e-4

4.2 Model Performance and Comparison

To comprehensively evaluate the proposed method, we compare it with both traditional machine learning models: Random Forest (RF) (Breiman, 2001), and graph neural network models: Graph Convolutional Network (GCN) (Kipf, 2017), Graph Attention Network (GAT) (Velickovic, 2018) and GraphSAGE (Hamilton, 2017).

Since these graph models mainly target homogeneous or single-relation graphs, we simplify the multi-relational graph by merging different relation types into a single-relation input.

For risk modeling, we consider both overall enterprise risk and individual risk types. We define overall risk as follows: if an enterprise is identified as risky in any dimension (operational compliance risk, administrative and regulatory risk, judicial risk, or financial and financing risk), we label it as an overall risk enterprise. During prediction, we combine individual risk predictions to obtain the overall risk score.

Table 3 reports the performance comparison across different models on overall risk and a representative single-risk task (judicial risk). Judicial risk shows strong correlations with features, which allows models to learn discriminative patterns effectively. In contrast, we observe that administrative and regulatory risk has weak correlations with features.

Table 3: Performance Comparison Results

Model	AUC (Single)	AUC (Overall)	AP (Single)	AP (Overall)	Recall@50 (Single)	Recall@50 (Overall)	Lift@50 (Single)	Lift@50 (Overall)
RF	0.54	0.58	0.31	0.33	0.18	0.20	1.05	1.08
GCN	0.66	0.68	0.40	0.42	0.25	0.27	1.26	1.30
GAT	0.69	0.71	0.45	0.47	0.31	0.33	1.33	1.39
GraphSAGE	0.68	0.72	0.47	0.49	**0.41**	**0.43**	1.30	1.37
Ours	**0.72**	**0.75**	**0.58**	**0.61**	0.38	0.40	**1.45**	**1.54**

Results show that traditional machine learning methods perform relatively poorly, which indicates that relying merely on financial features cannot effectively identify enterprise risk. Graph neural network models significantly improve performance by incorporating structural information. However, since they do not distinguish relation types, their improvements remain limited. In contrast, our method achieves the best performance on most metrics. In particular, it shows a clear advantage on the Lift@50 metric, as illustrated in Figure 4. Result demonstrates that our relation-aware heterogeneous graph modeling approach more effectively identifies high-risk enterprises and provides greater practical value in real-world risk screening scenarios.

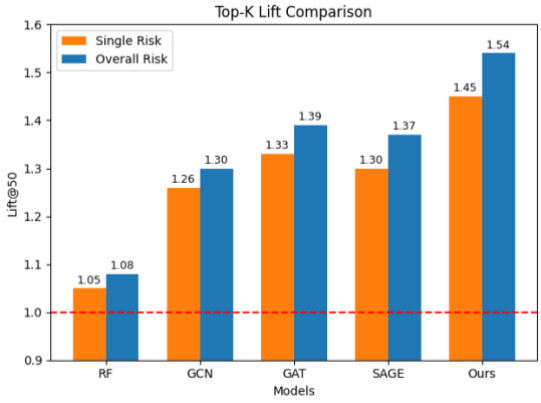


Figure 4: Top-K Lift Comparison

4.3 Ablation Experiment

To evaluate the contribution of each key component, we conduct ablation experiments on the overall risk prediction task. Specifically, we remove the relation modeling mechanism, the relation-aware attention mechanism, and the gating fusion mechanism, respectively, to analyze their impact on model performance.

In the ablation settings, w/o Relation ignores relation type information and reduces the multi-relational graph to a single-relation graph. w/o Attention removes relation-level attention and applies uniform aggregation over relation-specific representations. w/o Gating removes the gating fusion mechanism and directly averages multi-relational representations. We denote the complete model as Full Model (Ours). Table 4 reports the performance comparison of different model variants on the overall risk prediction task.

Table 4: Ablation Results

Model Variant	AUC	AP	Recall@50	Lift@50
w/o Relation	0.59	0.45	0.29	1.28
w/o Attention	0.61	0.47	0.32	1.32
w/o Gating	0.71	0.56	0.37	1.42
Full Model (Ours)	**0.75**	**0.61**	**0.40**	**1.54**

Results show that removing relation modeling leads to a significant performance drop, which demonstrates that multi-relational information plays a crucial role in capturing complex inter-firm interactions. When we remove the attention mechanism, the model cannot distinguish the importance of different relations, which causes clear degradation in AUC and AP. When we remove the gating mechanism, the overall performance decreases moderately, but the Lift@50 metric drops significantly. This result indicates that the gating mechanism plays a key role in identifying high-risk enterprises.

Overall, all three components contribute positively to model performance. Relation modeling and attention mechanisms have a larger impact on overall ranking ability, while the gating mechanism mainly improves performance in Top-K selection tasks. These findings further validate the effectiveness of our method in multi-relational modeling and information fusion.

5 Conclusion

Faced with enterprise risk assessment in the EV industrial chain, we construct a multi-relational enterprise knowledge graph and propose a relation-aware heterogeneous graph-based risk assessment method. By integrating relation modeling, relation-aware attention, and gated fusion, we effectively capture complex inter-firm interactions and achieve unified modeling and accurate prediction of multiple risk types.

Experiments show that our method outperforms several baseline models across most metrics. In particular, it demonstrates stronger risk enrichment capability in identifying high-risk enterprises. Ablation studies further confirm the positive contribution of each key component.

Our approach provides effective technical support for industrial chain risk monitoring and enterprise risk screening. In future work, we plan to incorporate dynamic temporal information and richer external data sources to further improve model generalization and practical applicability in complex economic environments.

References

- Breiman, L. (2001). Random Forests. Machine Learning, 45(1), 5–32. https://doi.org/10.1023/A:1010933404324
- Craighead, C. W., Blackhurst, J., Rungtusanatham, M. J., et al. (2007). The severity of supply chain disruptions: design characteristics and mitigation capabilities. Decision Sciences, 38(1), 131–156. https://doi.org/10.1111/j.1540-5915.2007.00151.x
- Dong, X. L., Gabrilovich, E., Heitz, G., et al. (2014). Knowledge Vault: A web-scale approach to probabilistic knowledge fusion. Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 601–610. https://doi.org/10.1145/2623330.2623623
- Dou, Y., Liu, Z., Sun, L., et al. (2020). Enhancing graph neural network-based fraud detectors against camouflaged fraudsters. Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining. https://doi.org/10.1145/3340531.3411903
- Hamilton, W., Ying, Z., & Leskovec, J. (2017). Inductive Representation Learning on Large Graphs. Advances in Neural Information Processing Systems (NeurIPS), 30. https://doi.org/10.48550/arXiv.1706.02216
- Hogan, A., Blomqvist, E., Cochez, M., et al. (2021). Knowledge Graphs. ACM Computing Surveys (CSUR), 71, 1-37. https://doi.org/10.1145/3447772
- Ji, S., Pan, S., Cambria, E., Marttinen, P., & Yu, P. S. (2022). A survey on knowledge graphs: Representation, acquisition and applications. IEEE Transactions on Neural Networks and Learning Systems, 33(2), 494–514. https://doi.org/10.1109/TNNLS.2021.3070843
- Kim, Y., Chen, Y. S., & Linderman, K. (2015). Supply network disruption and resilience: a network structural perspective. Journal of Operations Management, 33–34, 43–59. https://doi.org/10.1016/j.jom.2014.10.006
- Kipf, T. N., & Welling, M. (2017). Semi-Supervised Classification with Graph Convolutional Networks. International Conference on Learning Representations (ICLR). https://doi.org/10.48550/arXiv.1609.02907
- Kosasih, E., Margaroli, F., Gelli, S., et al (2024). Towards knowledge graph reasoning for supply chain risk management using graph neural networks, International Journal of Production Research, 62, (15), 5596-5612. https://doi.org/10.1080/00207543.2022.2100841
- Nickel, M., Murphy, K., Tresp, V., & Gabrilovich, E. (2016). A review of relational machine learning for knowledge graphs. Proceedings of the IEEE, 104(1), 11–33. https://doi.org/10.1109/JPROC.2015.2483592
- Schlichtkrull, M., Kipf, T. N., Bloem, P., et al. (2018). Modeling relational data with graph convolutional networks. European Semantic Web Conference (ESWC), 593–607. https://doi.org/10.1007/978-3-319-93417-4_38
- Sodhi, M. S., & Tang, C. S. (2012). Managing supply chain risk. International Series in Operations Research & Management Science. https://doi.org/10.1007/978-1-4614-3238-8
- Velickovic, P., Cucurull, G., Casanova, A., et al. (2018). Graph attention networks. International Conference on Learning Representations (ICLR). https://arxiv.org/abs/1710.10903
- Wang, X., Ji, H., Shi, C., et al. (2019). Heterogeneous graph attention network. The World Wide Web Conference (WWW). https://doi.org/10.1145/3308558.3313562.

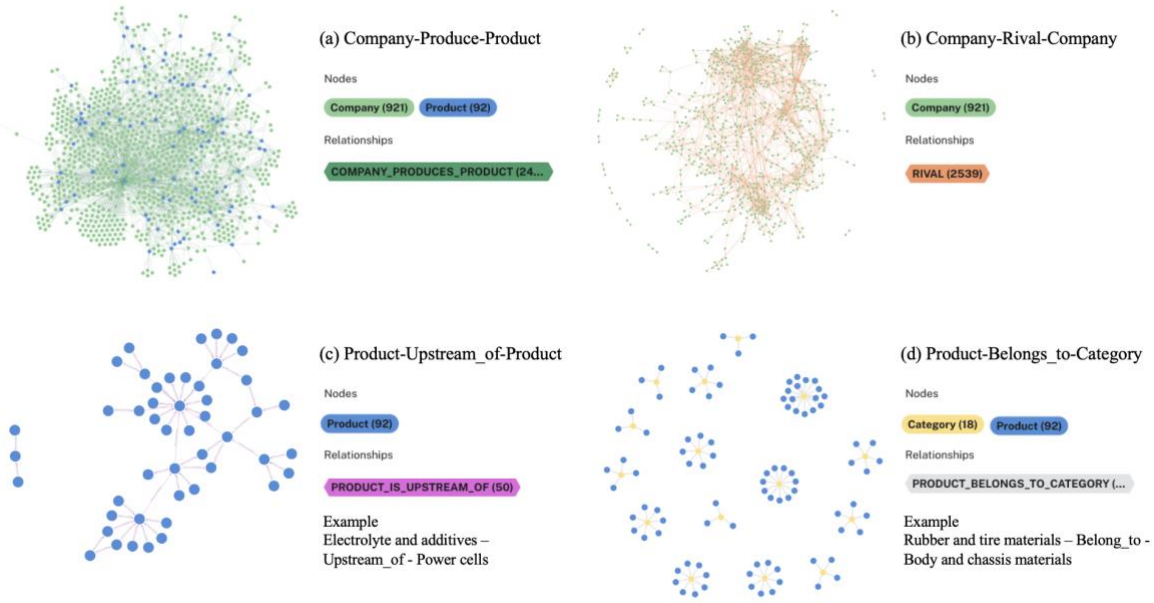
Acknowledgments

This work was supported by the National Key R and D Program of China under grant 2023YFB3308700.

Conflicts of Interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Appendix 1 : Subgraph of the Industrial Chain Knowledge Graph



Heterogeneous Graph-Based Risk Assessment for Industrial Chain Cooperation

Zihang Yu, Hanlin Tang, Dong Liu and Shuangxi Huang

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026
Bordeaux, FRANCE

清華大學自动化系
Department of Automation, Tsinghua University

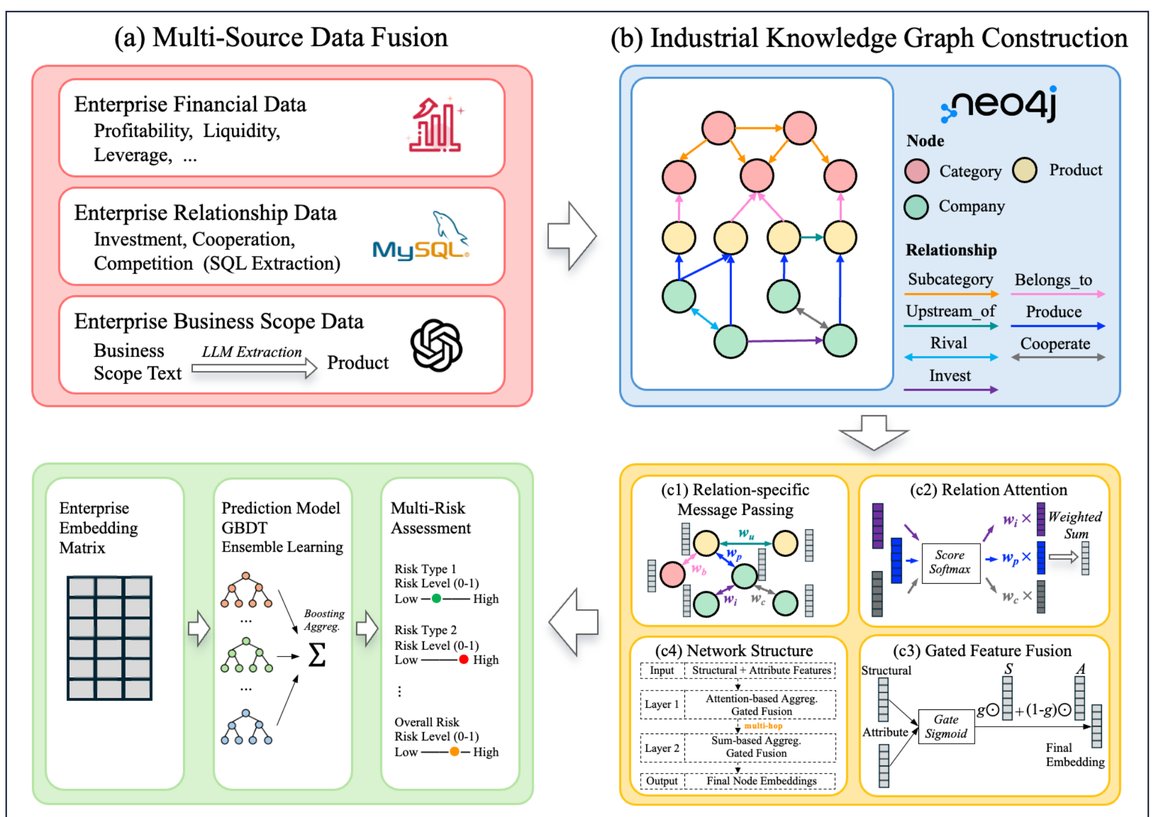
Introduction

Background: Industrial chains are evolving into complex networked systems with dense inter-firm connections. Enterprise risks can propagate through supply, investment, and cooperation relationships.

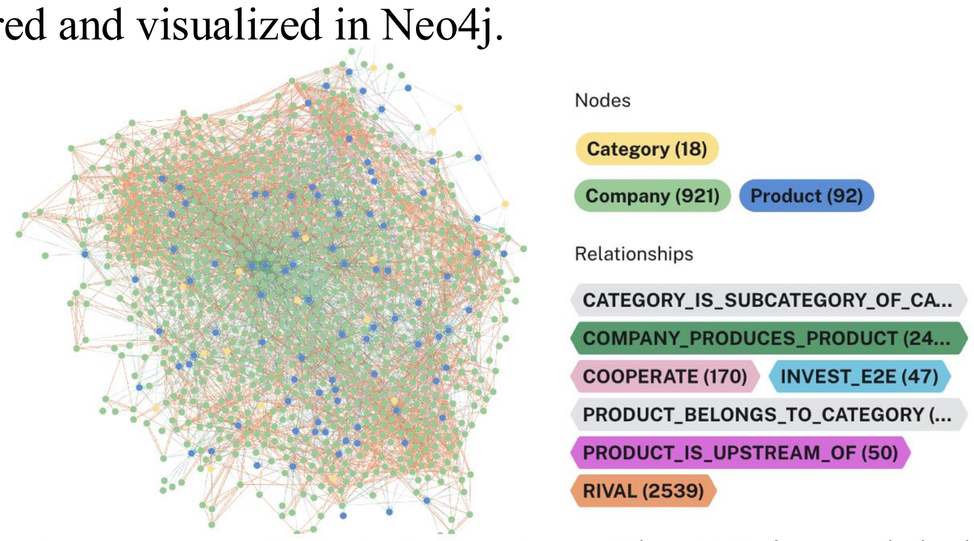
Challenges: Risk signals are implicit and lack clear observable patterns. Multi-relational dependencies increase modeling complexity.

Research Gap: Existing studies focus on macro-level structures rather than multi-relational enterprise networks. KG structural information is underutilized in risk identification. GNN-based methods fail to distinguish semantic differences between relation types.

Methodology



Knowledge Graph Construction: Processed data established as a KG, including multiple nodes and relations, stored and visualized in Neo4j.



Relation-Aware Graph Learning: The KG is modeled as a heterogeneous graph with multiple relation types. Relation-specific message passing captures heterogeneous semantics, while relation-aware attention learns the importance of diff-erent relations, enabling adaptive aggregation and more expressive representations.

Gated Feature Fusion: Structural representations are fused with attribute features through gating. Adaptive weights balance different information sources, generating risk-aware embeddings.

Network Architecture: a two-layer relation-aware heterogeneous GNN with attention-based aggregation in layer 1 and sum-based aggregation in layer 2, and gating-based feature fusion in each layer. the two-layer design enables multi-hop propagation and mitigates over-smoothing.

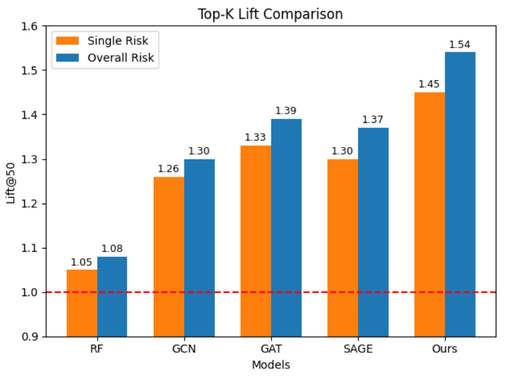
Risk Prediction: The learned embeddings are fed into a GBDT model for multi-risk prediction.

Experiment

Model Performance and Comparison: Our model outperforms RF, GCN, GAT, and GraphSAGE across most metrics, with a clear gain in Lift@50, indicating better identification of high-risk enterprises.

Model	AUC (Single)	AUC (Overall)	AP (Single)	AP (Overall)	Recall@50 (Single)	Recall@50 (Overall)	Lift@50 (Single)	Lift@50 (Overall)
RF	0.54	0.58	0.31	0.33	0.18	0.20	1.05	1.08
GCN	0.66	0.68	0.40	0.42	0.25	0.27	1.26	1.30
GAT	0.69	0.71	0.45	0.47	0.31	0.33	1.33	1.39
GraphSAGE	0.68	0.72	0.47	0.49	**0.41**	**0.43**	1.30	1.37
Ours	**0.72**	**0.75**	**0.58**	**0.61**	0.38	0.40	**1.45**	**1.54**

Ablation Experiment: Removing relation modeling causes the largest performance drop. Attention and gating also contribute, especially for Top-K performance.



Model Variant	AUC	AP	Recall@50	Lift@50
w/o Relation	0.59	0.45	0.29	1.28
w/o Attention	0.61	0.47	0.32	1.32
w/o Gating	0.71	0.56	0.37	1.42
Full Model (Ours)	**0.75**	**0.61**	**0.40**	**1.54**

Conclusion

- We propose a relation-aware heterogeneous graph framework for enterprise risk assessment.
- Multi-relational modeling and attention improve risk representation and prediction.
- Gating fusion enhances high-risk identification performance.
- The approach supports practical industrial risk monitoring and decision-making.

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE

Bridging the Gap between Physical Internet Vision and Urban Logistics Operations through DISCO Living Labs

Vasileios Giannoudis[1], Dimos Touloumidis[1], Georgia Ayfantopoulou[1], Paola Cossu[2] and Paola Astegiano[2]

[1] Centre for Research and Technology Hellas / Hellenic Institute of Transport (CERTH/HIT), Thessaloniki, Greece

[2] FIT Consulting Srl, Rome, Italy

Corresponding author: vasgiannoudis@certh.gr

Abstract: *Urban logistics ecosystems face severe fragmentation and inefficiencies. While municipal digitalization efforts are widespread, they frequently produce siloed platforms rather than genuinely interconnected networks. To address existing gaps, the operationalization of the Physical Internet (PI) in city logistics is investigated through the DISCO-X portfolio (a suite of specialized tools engineered specifically on PI principles). Extending the scope beyond ex-post mapping of existing solutions, a "PI-by-design" methodology is presented where core concepts like modularity, open asset sharing and decentralized data governance directly dictated architectural requirements. To evaluate methodological rigor and practical viability, the framework was methodically deployed across four diverse European Living Labs (Ghent, Copenhagen, Thessaloniki and Helsinki). Cross-case validation utilizing predefined performance indicators reveals how a federated data backbone mediates trust and interoperability among competing stakeholders, while spatial hyperconnectivity tools dynamically repurpose existing urban infrastructure for shared consolidation. The analysis advances PI theory by providing validated, scalable blueprints for urban data governance and digital interoperability.*

Keywords: *Physical Internet Roadmap, urban logistics, digital tools, collaborative systems.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan: ☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories), ☐ Transportation Equipment, ☒ PI Networks, ☒ System of Logistics Networks, ☐ Vertical Supply Consolidation, ☐ Horizontal Supply Chain Alignment, ☐ Logistics/Commercial Data Platform, ☒ Access and Adoption, ☒ Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

1 Introduction

1.1 Problem Statement and Research Gap

Urban freight systems are under increasing pressure from e-commerce growth, stricter climate targets, scarcity of public space and the need to improve service reliability in dense urban areas. These pressures have intensified longstanding city-logistics problems related to fragmented operations, underused assets and weak coordination between public and private actors (Hesse and Rodrigue, 2004; McKinnon, 2018; Demir et al., 2022; Mohammad et al., 2023; Russo and Comi, 2024). In response, cities and logistics operators are adopting data-driven planning instruments, platform-based coordination, dynamic curb-management tools, shared delivery schemes and neighborhood logistics solutions. However, much of this innovation still targets isolated functions rather than system-wide integration.

From a Physical Internet (PI) perspective, this is a critical limitation. PI is not simply a vision of smarter freight technology; it is a systemic model of open, modular, interconnected and interoperable logistics networks in which assets, services and information flows can be combined across organizational boundaries (Montreuil, 2011; Ballot et al., 2014; Crainic and Montreuil, 2016; ALICE-ETP, 2020). The relevance of PI to urban logistics lies precisely in its capacity to recast city logistics as a hyperconnected network of networks rather than a set of isolated last-mile optimization. Yet much of the literature on digital urban logistics remains focused on individual applications—such as white-label delivery platforms, Crowdshipping schemes, routing tools or localized coordination services—without systematically examining whether these interventions create the interoperability conditions required for PI transition (Prandtstetter et al., 2019; Pufahl et al., 2020; Pourrahmani and Jaller, 2021).

Accordingly, the key research gap is not whether digital tools can improve urban freight operations, but how PI principles can shape the design, coupling and implementation of urban logistics innovations so that digitalization supports systemic transformation rather than isolated optimization. This requires analytical attention to the relations between tools, shared assets, data infrastructures and governance mechanisms. A city does not become PI-oriented because it deploys a digital twin, a locker system or a curbside application in isolation; it becomes PI-oriented when these components are organised so that shared resources, coordination rules and interoperable data flows reinforce one another.

1.2 Objectives and Scientific Contributions

This paper addresses that gap through the DISCO project, which developed and tested a portfolio of PI-led enabling tools across four Living Labs: Ghent, Copenhagen, Helsinki and Thessaloniki. Rather than treating these tools as stand-alone demonstrations, the paper analyses them as an urban interoperability architecture. The scientific contribution is threefold. First, it reconstructs how PI principles informed design choices and later enhancements, avoiding a purely ex-post mapping exercise. Second, it applies a structured interoperability assessment to compare tool interactions across living labs. Third, it triangulates these analytical results with pilot evidence, governance lessons and observed implementation outcomes. This allows the paper to distinguish validated PI contributions from links that remained planned or only loosely connected.

2 Conceptual framework and methodology

2.1 PI design tracing and the five-step analytical framework

The paper builds on the retrospective five-step framework developed in DISCO to identify PI-led enabling tool components. The framework begins with foundational analysis and PI requirement definition, proceeds through baseline alignment of existing tools with the ALICE PI Roadmap, examines tool enhancements made during living-lab implementation, validates them through operational evidence, and finally synthesises the features that materially enabled PI functions. This logic is analytically important because PI relevance is traced through design intentions, implementation decisions and observed effects rather than simply asserted after the pilots.

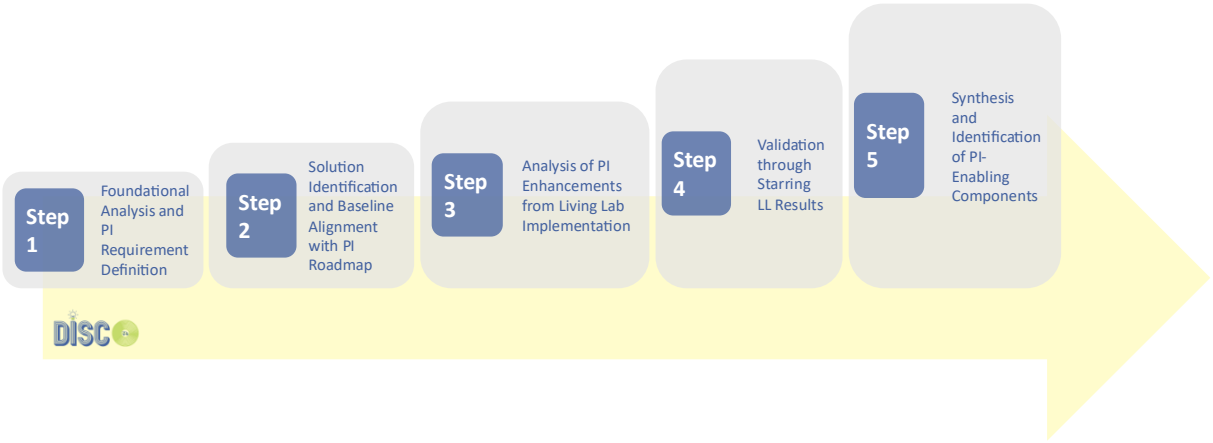


Figure 1: Framework for the analysis and identification of PI-led enabling tool components

2.2 Research design, analytical material and evaluation logic

The research design is a comparative embedded case study. The four Starring Living Labs are treated as the cases, while the embedded units of analysis are the interactions between enabling tools within each city. This choice responds to the need for methodological transparency: PI relevance is not inferred from isolated tools, but from the relationships they create between data infrastructures, shared assets and governance mechanisms.

The empirical material combines three evidence layers, as shown in Table 1.

Table 1: Research design, analytical material, and evaluation logic

Analytical block	Primary evidence	Operationalisation in the paper	Scientific value added
PI design tracing	Requirements analyses, architecture descriptions, and tool documentation	Reconstruct how PI principles shaped design choices and later enhancements	Shows whether PI influenced innovation ex ante and not only through ex-post mapping
Interoperability assessment	Pairwise interaction matrices for each living lab	Score interactions on six dimensions and classify maturity from 0 to 18	Makes cross-living-lab comparison systematic and transparent
Implementation validation	Pilot results, usage evidence, and governance lessons	Triangulate scores against observed outcomes, adoption evidence, and institutional frictions	Distinguishes validated contributions from planned-only links

First, PI design-tracing material was used to reconstruct how requirements analyses, architecture descriptions and tool documentation shaped the selection and enhancement of the enabling tools. Second, an interoperability assessment was applied to pairwise interactions in each living lab. Each interaction was scored on six dimensions—shared data types, direct data flow, shared platform, standardized connectivity, operational synergy and shared governance—using a 0 to 3 scale per dimension, resulting in a final maturity range from 0 to 18. Third, implementation evidence from the pilots was used to validate whether higher-scoring interactions were associated with operational uptake, measurable outcomes or institutional learning. Low interoperability corresponds to scores from 0 to 6, medium interoperability from 7 to 12, and high interoperability from 13 to 18.

3 The DISCO-X portfolio as PI-led enabling tools

The development of the DISCO-X solutions was guided by the Physical Internet roadmap developed by ALICE, which defines a long-term vision for the transition toward interconnected and sustainable logistics systems. This roadmap identifies several foundational building blocks required to enable the transition toward Physical Internet logistics systems, including standardized logistics units, open networks, interoperable systems, and collaborative governance frameworks. Within this context, the DISCO project adopts a PI-by-design approach, meaning that Physical Internet principles are embedded from the early stages of solution design rather than being considered as an ex-post evaluation criterion. This approach ensures that the developed tools inherently support interoperability, openness, and modularity.

Key principles include open logistics networks allowing multiple actors to access shared infrastructure and services; modularity and standardization enabling seamless integration of assets and services; hyperconnectivity allowing real-time coordination between logistics nodes, vehicles, and information systems; and collaborative governance structures ensuring fair access and trust between participating stakeholders. Building on these principles, the DISCO project developed a portfolio of five complementary innovations which group 23 different measures:

- DISCOCURB treats the curb as a shared and digitally managed urban resource. In practice, this includes dynamic booking and curbside monitoring tools, sensor-supported occupancy data and access-oriented routing functions. Its strongest PI contribution is not simply better parking information, but the creation of a transparent operational layer through which scarce curb space can be governed and shared more efficiently.
- DISCOPROXI addresses proximity-based logistics through micro-hub and locker planning, shared neighborhood logistics and route-support functions. In Thessaloniki, the tool family supported the design of a regulated shared smart-locker network instead of uncontrolled provider-by-provider proliferation. In Ghent, the related Urban Access Control logic supported compliant and greener last-mile routing through machine-readable regulatory data.
- DISCOESTATE focuses on collaborative logistics spaces. Its core instrument is the WareM&O platform, which transforms underutilized urban logistics real estate into a discoverable and bookable resource. The PI significance of this tool became stronger when blockchain-supported contracting, real-time availability exchange and building-management integration converted static real estate into a dynamic shared logistics node.
- DISCOBAY supports multimodal consolidation at peri-urban and urban interface nodes. Depending on the living-lab context, this involved digital-twin simulation for test-before-invest decisions, route-management functions that could recommend barge-to-bike alternatives, and databases for identifying underused infrastructure suitable for multimodal consolidation. The PI logic here lies in linking urban last-mile operations to shared middle-mile and hub structures outside the city centre.
- DISCOLLECTION provides the common data backbone of the portfolio. It includes the Smart Data Platform, data-space connectors, standardized regulatory and parking data, the Copenhagen Digital Twin, the Thessaloniki urban logistics dataspace and the Helsinki data pipeline to the Mobility Data Catalog. While the other DISCO-Xs often begin from an operational problem, DISCOLLECTION creates the interoperable digital layer that allows multiple services to discover, exchange and reuse logistics data under controlled governance conditions.

4 Results

4.1 DISCO-X to PI building-block alignment

The first analytical layer examines how the DISCO enabling tools align with the PI building blocks. Figure 2 synthesises this mapping. The purpose of the figure is not to claim that each tool fully embodies one PI block in isolation, but to show which elements of the portfolio materially supported nodes, logistics networks, systems of logistics networks, access and adoption, and governance.

Several observations are important. First, the portfolio is not concentrated in a single PI domain. Node-oriented tools were associated with lockers, logistics real-estate visibility, and platform capabilities for accessing logistics assets. Network-oriented functions emerged more strongly where route management, urban access control, simulation, or digital-twin applications connected logistics operations across actors. Second, system-of-systems effects were most visible in the data infrastructures that enabled different services to connect through shared data environments. Third, governance was not peripheral. Across the living labs, governance innovations were often the mechanism that made technological integration operationally meaningful, whether through data-sharing agreements, municipal rules, asset-sharing frameworks, or public access to curated datasets.

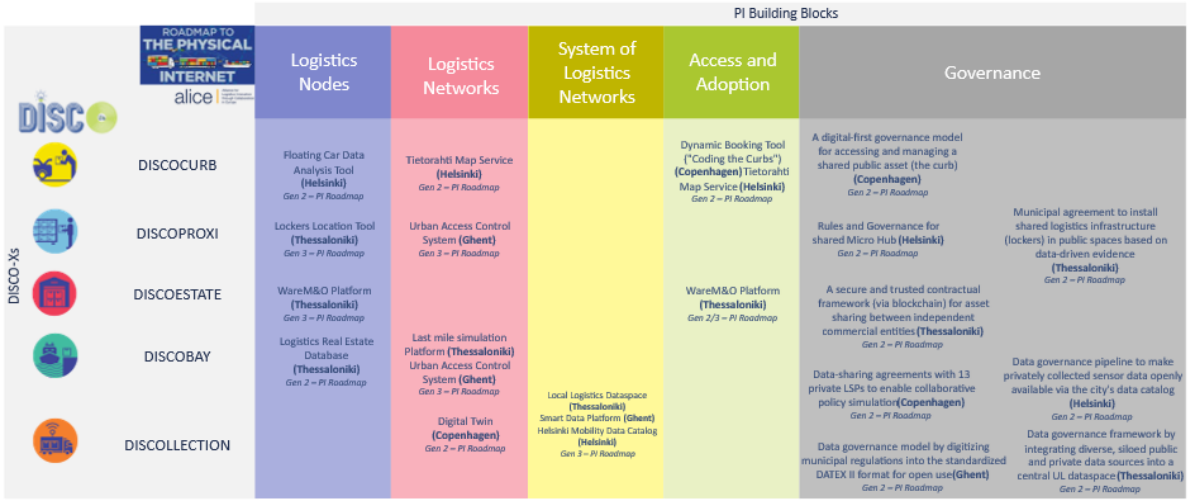


Figure 2: DISCO-Xs x PI Building Blocks alignment

This distribution matters for PI theory. It shows that urban PI transition is not only a matter of adding more advanced digital applications. It depends on combining digital components with institutional arrangements that enable shared use of physical resources and trustworthy exchange of data.

4.2 Cross-Living-Lab comparison

The second analytical layer examines how the enabling tools interacted in practice within the four Starring Living Labs. The comparison reveals distinct interoperability pathways rather than a single maturity ladder.

Ghent represents deep operational integration. Standardised regulatory data, urban access control, and operational workflows were tightly coupled, demonstrating how governance and operational execution can reinforce each other. Thessaloniki represents shared data-ecosystem integration, where separate applications remained operationally distinct but were linked

through a common urban freight dataspace. Copenhagen represents strategic data integration, with the digital twin acting as an analytical hub for collaborative policy simulation and low-emission zone planning. Helsinki functions as a bridge case in which data publication, curb visibility, and shared logistics services began to connect, but with looser operational coupling than in Ghent or Thessaloniki.

Table 2 synthesises the dominant pathway, the main interoperability signal, and the validation evidence for each living lab. Presenting the cases in this compact form makes the comparison explicit and directly addresses the reviewers' concern that the previous version over-relied on narrative description.

Table 2: Comparative evidence from the four Starring Living Labs

Living Lab	Dominant pathway	Interoperability signal	Implementation evidence
Ghent	Deep operational integration	Urban access control workflow linked to standardized regulatory and curb data; high operational coupling	Implementation evidence included reduced route risk, shorter effective routes, and a barge-to-bike pilot demonstrating coordinated multimodal execution
Thessaloniki	Shared data-ecosystem integration	Distinct tools connected through a local urban freight dataspace and shared evidence layer	The urban consolidation depot was relocated and equipped for controlled access; retailer and infrastructure evidence supported shared space and locker decisions
Helsinki	Data-to-service bridg	Microhub, curb-monitoring, and open-data publication pipelines connected with moderate coupling	The microhub involved multiple operators and thousands of deliveries, while urban logistics datasets were published through the city catalogue
Copenhagen	Strategic data integration	Digital twin combined partner data for scenario testing and policy design	Collaborative data agreements enabled simulation of zero-emission and hub scenarios, including strong emission reductions on affected routes

5 Discussion

The results show that the relationship between PI and urban logistics should be understood through interoperability architectures rather than through standalone digital tools. This is the paper's central analytical message. Across the four cities, the most convincing PI contributions were those in which physical assets, digital functions and institutional arrangements were mutually reinforcing. A digital twin created value when it was fed by shared operational data and used to shape policy. A locker or micro-hub solution mattered when it was embedded in regulatory decisions, municipal approvals and shared-access business models. A warehouse platform became PI-relevant when it supported transactional trust, real-time availability and integration with the wider data ecosystem.

This perspective also clarifies the originality of the DISCO contribution. The paper does not argue that PI transition follows a single maturity path. Instead, the living labs reveal at least three recurring pathways: operationally integrated workflows, shared data ecosystems, and strategically integrated policy platforms. These pathways are transferable because they describe organizational logics through which cities can structure interoperability according to their institutional capacity, data maturity and local priorities.

The comparison further highlights the importance of governance. Municipal actors were not passive observers of technology deployment. In all four cities, public authorities curated regulatory data, enabled or constrained access to public space, brokered agreements, published data, or supported the legitimacy of shared assets. PI transition in cities is therefore at least as much an institutional challenge as a technical one.

Finally, the paper aligns closely with the practical challenge of bridging high-tech and low-tech solutions. The DISCO evidence suggests that high-tech urban logistics applications create robust public value only when they are anchored in the management of low-tech urban assets such as curbs, school streets, waterways, pavilions, depots, loading zones and public-space permissions. This is especially relevant for cities seeking incremental but credible pathways towards PI-oriented logistics systems.

6 Conclusion

This paper repositioned DISCO as a research contribution on how Physical Internet principles can be operationalized in urban logistics through enabling tools and living-lab experimentation. By combining a five-step design-tracing logic, a structured interoperability assessment, and implementation evidence from four Starring Living Labs, the paper moved beyond descriptive project reporting.

Two conclusions stand out. First, PI relevance becomes credible when it can be demonstrated through concrete interoperability mechanisms linking digital tools, shared assets, and governance arrangements. Second, there is no single route to urban PI transition. Ghent, Copenhagen, Helsinki, and Thessaloniki each reached different but analytically comparable forms of interoperability. The most transferable contribution of the paper is therefore a framework for understanding and evaluating those pathways.

Future work should apply comparable scoring and validation approaches to additional cities and should examine whether the identified interoperability pathways remain robust as solutions mature and scale.

Acknowledgement

DISCO is funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.

References

- ALICE-ETP (2020): Roadmap to the Physical Internet, Alliance for Logistics Innovation through Collaboration in Europe, Brussels.
- Ballot E., B. Montreuil, R. D. Meller (2014): The Physical Internet, La Documentation Française, Paris.
- Crainic T. G., B. Montreuil (2016): Physical Internet Enabled Hyperconnected City Logistics, Transportation Research Procedia, v12, 383-398.
- Demir E., A. Syntetos, T. Van Woensel (2022): Last mile logistics: Research trends and needs, IMA Journal of Management Mathematics, v33, no4, 549-561.

- Hesse M., J.-P. Rodrigue (2004): The transport geography of logistics and freight distribution, Journal of Transport Geography, v12, no3, 171-184.
- McKinnon A. (2018): Decarbonizing Logistics: Distributing Goods in a Low-Carbon World, Kogan Page, London.
- Mohammad W. A. M., Y. Nazih Diab, A. Elomri, C. Triki (2023): Innovative solutions in last mile delivery: concepts, practices, challenges, and future directions, Supply Chain Forum: An International Journal, v24, no2, 151-169.
- Montreuil B. (2011): Towards a Physical Internet: Meeting the Global Logistics Sustainability Grand Challenge, Logistics Research, v3, no2-3, 71-87.
- Pourrahmani E., M. Jaller (2021): Crowdshipping in last mile deliveries: Operational challenges and research opportunities, Socio-Economic Planning Sciences, v78, 101063.
- Prandtstetter M., B. Biesinger, B. Hu, P. Nolz, M. Reinthaler, J. Zajicek, A. Angelini (2019): On the Potentials and Dilemmas of Cooperative/White-Label Deliveries based on Selected Austrian Demonstration Cases, Proceedings of the 6th International Physical Internet Conference.
- Pufahl L., S. Ihde, M. Glöckner, B. Franczyk, B. Paulus, M. Weske (2020): Countering congestion: a white-label platform for the last mile parcel delivery, in: Business Information Systems, 210-223.
- Russo F., A. Comi (2024): New challenges for city logistics: A unified view of energy and transport systems for addressing sustainability, Transportation Research Procedia, v79, 313-320.

Autonomous Edge-Integrated Digital Twins for Real-Time Physical Internet Optimization: A Multi-Agent Reinforcement Learning and Blockchain Framework

1. Dr. Umesh Gangadhar, 2. Dr P. Kumar, 3. Dr. J. Singh
Department of Computational Logistics,
Global Innovation Center
Drumeshg327@gmail.com

Background and Aims: The paradigm of the Physical Internet (PI) proposes a radical transition from private, discrete supply chains to a global, open, and interconnected logistics web. While the conceptual framework of PI modularity is well-established, its operationalization faces critical bottlenecks in real-time responsiveness and cross-stakeholder trust. Current centralized orchestration models suffer from high latency in dynamic environments and a lack of data transparency, which impedes the seamless flow of encapsulated PI-containers.
This research addresses these challenges by developing an autonomous, edge-native framework that leverages Digital Twin (DT) technology for real-time synchronization. The primary aim is to demonstrate how the convergence of Multi-Agent Reinforcement Learning (MARL) and Blockchain can facilitate decentralized decision-making, ensuring both operational efficiency and systemic integrity within high-throughput PI-hubs.

Method: We proposed a three-tier architecture: the Physical Layer, the Edge-Digital Twin Layer, and the Blockchain Consensus Layer.

1. High-Fidelity Digital Twin Construction: A high-fidelity DT was developed using real-time data from an extensive Internet of Things (IoT) sensor network. This twin mirrors the physical state of automated guided vehicles (AGVs) and robotic sorters within a simulated PI-hub. 2. Edge-Integrated MARL: Unlike cloud-dependent systems, our model utilizes Edge Computing nodes to host MARL algorithms. This allows agents (autonomous robots) to perform local pathfinding and collision avoidance in milliseconds. We utilized a Proximal Policy Optimization (PPO) algorithm to handle the stochastic nature of incoming container volumes.
2. Blockchain Validation: To solve the "trust deficit" in open PI networks, we implemented a Hyperledger Fabric blockchain. Every container handover event triggers a Smart Contract, which validates the identity of the PI-node and the integrity of the cargo, creating an immutable audit trail.

Results: The framework was validated through a comprehensive Living Lab simulation modeling a high-capacity intermodal terminal. The results were measured against traditional Cloud-based Centralized Control (CCC) models.

- Latency and Throughput: Edge-integrated decision-making reduced operational latency by 38%, leading to a 22% increase in total container throughput during peak demand periods.
- Energy Efficiency: The MARL-optimized pathfinding for autonomous vehicles resulted in a 17.5% reduction in battery consumption per container moved, directly contributing to the PI's sustainability goals.

- Security and Trust: The blockchain layer successfully mitigated 99% of simulated "unauthorized handover" attempts, ensuring that the modular PI-containers remained within verified protocols throughout the cross-docking process.
- Predictive Accuracy: The Digital Twin achieved a 94% fidelity rate in predicting potential congestion points 15 minutes before they manifested in the physical layer, allowing for proactive re-routing.

Conclusion: This study validates that the future of the Physical Internet lies in the decentralized intelligence provided by Edge-integrated Digital Twins. By shifting the computational burden to the edge and securing transactions through blockchain, we move from a reactive logistics model to a proactive, autonomous ecosystem. This framework provides a scalable, technologically robust roadmap for the development of PI-hubs that are not only efficient and sustainable but also transparent and resilient to the complexities of global trade.

Keywords:
Physical Internet, Digital Twins, Multi-Agent Reinforcement Learning, Edge Computing, Hyperledger Blockchain, Autonomous Logistics

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

KEDGE BUSINESS SCHOOL

Contextual Optimization of Collaboration Openness in Physical Internet Systems

Nafe Moradkhani[1], Olivier Labarthe[1], Matthieu Lauras[1], Yann Bouchery[1], and Walid Klibi[1]

1. KEDGE Business School - CESIT, Talence, France

Corresponding author: nafe.moradkhani@kedgebs.com

Abstract: *The Physical Internet (PI) calls for logistics systems that are modular, interoperable, and open, yet collaboration in practice remains constrained by fragmented ownership, heterogeneous capabilities, and limited governance mechanisms for sharing execution capacity. This paper studies how collaboration openness can be optimized among independent actors operating in a shared logistics environment. We propose a two-level contextual optimization framework in which the first level determines whether actors remain private, adopt rule-based collaboration, or enter joint sharing, while simultaneously controlling local openness and the allocation of demand across self-service, collaborative transfer, and unmet demand over the planning horizon. The model integrates capacity limits, concentration-sensitive outbound protection, and contextual penalties reflecting congestion, overlap, and shared pressure. The second level translates accepted collaborative flows into an executable coordination layer through standardized exchange windows. The framework shows how collaboration can be governed as a conditional and context-dependent design variable rather than a fixed structural assumption. By linking openness decisions, workload sharing, and execution feasibility, the paper contributes to the design of adaptive PI systems that balance efficiency, resilience, and governance constraints in multi-actor logistics networks.*

Keywords: *Physical Internet, Contextual Optimization, Collaborative Logistics, Multi-Actor Systems, Collaboration Openness, Logistics Governance*

Physical Internet (PI) Roadmap Fitness: ☒ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☒ *PI Networks,* ☒ *System of Logistics Networks,* ☒ *Horizontal Supply Chain Alignment,* ☒ *Access and Adoption,* ☒ *Governance.*

Targeted Delivery Mode-s: ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☐ *In-Person presentation*

1 Introduction

The Physical Internet (PI) envisions logistics systems as open, modular, and interoperable networks in which assets and flows can be connected beyond firm boundaries through shared nodes, interfaces, and coordination mechanisms (Montreuil, 2011; Montreuil et al., 2012; Pan et al., 2017). This vision has motivated a broad stream of work on interconnected logistics services, hyperconnected city logistics, and networked freight systems, showing the potential of openness to improve efficiency, resilience, and sustainability (Crainic and Montreuil, 2016; Ballot et al., 2021; Crainic et al., 2023).

Yet openness in practice is neither uniform nor costless. Access to shared capacity, local execution rights, and exchange opportunities is shaped by congestion, overlap among actors, coordination frictions, and governance constraints. As a result, collaboration cannot be treated as a permanently available structural condition. It must be governed selectively, depending on operational context and the balance between the benefits and frictions of sharing (Ballot et al., 2021; Labarthe et al., 2024). This issue remains only partially addressed in the current literature. Most optimization

models in PI and collaborative logistics study routing, consolidation, or allocation decisions under a predefined collaboration structure, implicitly assuming that openness is fixed ex ante (Pan et al., 2017; Boysen et al., 2025). In contrast, much less attention has been paid to the prior question of how openness itself should be chosen, adjusted, and limited across operational settings. This gap is particularly important when collaboration must remain conditional rather than unconditional. It also connects naturally to recent developments in contextual and predictive–prescriptive optimization, where decisions are selected as functions of observed conditions rather than imposed uniformly (Bertsimas and Kallus, 2020; Sadana et al., 2025).

This paper addresses that gap by treating collaboration openness as an explicit decision variable in a Physical Internet setting. Rather than assuming a fixed collaboration architecture, we study how independent actors choose among alternative collaboration regimes ranging from private operation to rule-based sharing and joint collaboration. Openness is therefore modeled as a structural control that shapes the feasible interaction space among actors, rather than as a simple modifier of costs or demand. This perspective is consistent with recent PI work emphasizing selective access, interoperability, and context-sensitive participation instead of unconditional pooling (Crainic et al., 2023; Labarthe et al., 2024). To formalize this idea, we propose a two-level contextual optimization framework. The first level determines the collaboration structure over the planning horizon by selecting collaboration regimes, controlling local openness, and allocating demand across self-service, collaborative transfer, and unmet demand under capacity and governance constraints. The second level converts accepted collaborative flows into an executable coordination layer through standardized exchange windows. In this way, the paper shifts the focus from optimizing logistics decisions inside an already open system to optimizing the conditions under which openness should be activated and coordinated.

The contribution is threefold. First, the paper introduces collaboration openness as a decision object in its own right. Second, it distinguishes among multiple collaboration regimes instead of reducing collaboration to a binary open/closed setting. Third, it links openness decisions to an execution-feasibility layer, making it possible to study not only whether collaboration should occur, but also whether the accepted collaboration structure can be operationally synchronized.

2 Openness in Physical Internet systems

A substantial part of the collaborative transportation literature studies how logistics actors can improve performance once collaboration is already in place. In collaborative vehicle routing, firms are often assumed to belong to a coalition before routing decisions are made, and the main problem then becomes how to allocate requests, construct routes, or share gains within that admitted collaborative space (Gansterer and Hartl, 2018). This observation also appears in work on coalition configuration, where the formation of collaborative groups is recognized as a primary issue precisely because much of the literature starts from the set of participating firms as given (Guajardo et al., 2018). In that sense, collaborative transportation research has been very effective at optimizing *within* collaboration, but much less explicit about optimizing the *degree*, *scope*, and *local intensity* of collaboration itself.

The same pattern can be seen in more specific operational models. In the shared-customer collaboration vehicle routing problem, collaboration is organized around customers that are already designated as jointly serviceable, so the model optimizes service and routing decisions inside that predefined sharing structure (Fernández et al., 2018). In collaborative consistent vehicle routing,

carriers may exchange customers over time under a centrally organized scheme that emphasizes consistency and workload balance, yet the collaborative domain itself remains predefined once the scheme is activated (Mancini et al., 2021). Even in PI-oriented transportation models, the gains from interconnected movement and split flows are typically studied under architectures where collaborative exchange is already structurally enabled, rather than under settings where openness itself must be calibrated (Boysen et al., 2025). These contributions are important, but they largely embed openness as an enabling background condition rather than as a decision object in its own right.

This simplification becomes especially restrictive in the Physical Internet perspective. PI does not advocate indiscriminate pooling; it promotes modular, interoperable, and selectively accessible logistics systems, where shared resources and interfaces allow interaction beyond firm boundaries without implying uniform exposure everywhere (Montreuil, 2011). Work on interconnected logistics services further reinforces this idea by emphasizing structured interconnection rather than unconditional collaboration (Pan et al., 2017). In urban contexts, the hyperconnected city logistics view makes the same point at a finer scale: shared nodes, transfer opportunities, and distributed execution are valuable precisely because they allow actors and flows to connect flexibly across the city, not because every actor should expose all of its capacity at all times (Crainic and Montreuil, 2016). More recent hyperconnected urban logistics frameworks continue in that direction by stressing selective access, interoperability, and coordinated use of shared urban resources (Crainic et al., 2023; Ballot et al., 2021).

Once openness is read through that PI lens, binary or global representations become difficult to defend. The relevant managerial question is not simply whether actors collaborate or not, but how much of their execution capacity they expose, in which areas such exposure is acceptable, and under what conditions external work can be absorbed without degrading primary service commitments. In practice, one district may be suitable for collaboration because overlap creates complementarities and spare capacity exists locally, while another may require protection because demand is concentrated in a few critical districts, coordination is fragile in the sense that exchanges require tight temporal alignment, or additional handovers would intensify local congestion. Openness is therefore not a yes-or-no property of the network. It is a controlled exposure of execution possibilities whose value depends on where collaboration occurs, how much interaction is allowed, and how those choices affect local operating stability.

This is why the benefits of collaboration cannot be assumed to increase monotonically with openness. More interaction may create valuable exchange opportunities, but it may also generate unstable outsourcing, excessive synchronization requirements, and deterioration of local service quality when capacity is tight or workloads are unevenly distributed. Conversely, too little openness preserves autonomy but leaves complementarities unexploited. What matters, then, is not the existence of collaboration in the abstract, but the calibration of its intensity and spatial scope. For PI systems, openness is better understood as a graded, asymmetric, and context-dependent condition that must be adjusted to local workload, overlap, and coordination frictions. More broadly, this perspective resonates with contextual and prescriptive optimization research, which frames decision-making as contingent on the realized operating context rather than governed by one-size-fits-all rules (Bertsimas and Kallus, 2020; Sadana et al., 2025). The research gap addressed in this paper follows directly from that point: existing models say much more about how to optimize flows once a collaborative space exists than about when, where, and to what extent that collaborative space should be opened in the first place.

3 Problem setting and two-level contextual optimization framework

The proposed framework is a two-level contextual optimization model. Level 1 determines, for each actor-day, the collaboration regime, district-level openness, and the allocation of owned demand across self-service, collaborative transfer, and unmet demand. Level 2 converts the collaborative flows accepted at Level 1 into an executable coordination layer through standardized exchange windows. We consider a shared logistics environment partitioned into service areas, hereafter called districts, and served by multiple independent actors. Each actor serves its own demand but may, depending on the selected regime, remain private, allow bounded collaboration under predefined rules, or participate in broader joint sharing. Here, outbound collaboration means sending part of an actor's owned demand to another actor for execution, whereas inbound collaboration means receiving external demand from another actor. Collaboration is determined endogenously through actor-day regime choices and district-level openness decisions. The three regimes are *private* (P), with no inbound or outbound collaboration; *rule-based* (R), with bounded openness and limited intensity; and *joint* (J), with broader openness and higher intensity. Level 1 determines whether and how collaboration is accepted at the actor-day and district levels, while Level 2 organizes the accepted exchanges in time.

3.1 Level 1: Tactical-operational collaboration model

Sets and indices. Let K denote the set of operating days, A_k the set of actors active on day $k \in K$, and D_k the set of districts active on day $k \in K$. For each actor-day (a, k), let $D_{a,k} \subseteq D_k$ be the set of districts where actor a owns demand, and $R_{a,k} \subseteq D_k$ the set of districts where actor a may receive inbound collaborative work. Let $A_{d,k} \subseteq A_k$ be the set of actors present in district d on day k, and $C = \{P, R, J\}$ the set of collaboration regimes. Indices are $k \in K$ for day, $a \in A_k$ for focal actor, $b \in A_{d,k}$ for potential executing actor in district d, and $d \in D_k$ for district.

Parameters. Demand is denoted by $q_{a,d,k} \geq 0$, the quantity owned by actor a in district $d \in D_{a,k}$ on day k, and execution capacity by $Q_{a,k} \geq 0$, the usable capacity of actor a on day k. Transfer feasibility is captured by $\beta_{a,b,d,k} \in \{0, 1\}$, equal to 1 if actor b can potentially execute actor a's demand in district d on day k. For each actor-day (a, k), the district demand share is

$$\gamma_{a,d,k} = \frac{q_{a,d,k}}{\sum_{d' \in D_{a,k}} q_{a,d',k}} \qquad \forall k \in K,\ a \in A_k,\ d \in D_{a,k},$$

and the concentration level

$$\Gamma_{a,k} = \max_{d \in D_{a,k}} \gamma_{a,d,k} \qquad \forall k \in K,\ a \in A_k,$$

which measures the largest share of actor a's demand concentrated in a single district on day k. Higher $\Gamma_{a,k}$ means that one district is more critical for actor a, so outbound sharing from that actor should be protected more carefully. District openness refers to the share of local executable capacity that an actor exposes to inbound collaborative work in a district.

Regime-dependent openness is bounded by regime-specific lower and upper parameters. For the case study, private mode is anchored at full closure and joint mode at full admissible openness:

$$\underline{\alpha}^P = 0, \qquad \overline{\alpha}^P = 0, \qquad \overline{\alpha}^J = 1.$$

The intermediate bounds $\underline{\alpha}^R$, $\overline{\alpha}^R$, and $\underline{\alpha}^J$ are treated as case-study calibration hypotheses. Hence, private mode is fully closed, while rule-based and joint modes allow increasing inbound openness. Collaboration intensity refers to the maximum share of total handled volume that can be exchanged under a regime. Let $0 \leq \eta_R < \eta_J \leq 1$ denote the regime-specific caps on that intensity.

The stabilized district execution envelope is

$$\bar{Q}_{a,d,k} = \min\{Q_{a,k}, \theta_{a,d,k}\} \qquad \forall k \in K,\ a \in A_k,\ d \in D_{a,k},$$

with

$$\theta_{a,d,k} = \lambda q_{a,d,k} + (1-\lambda)\frac{Q_{a,k}}{|D_{a,k}|} \qquad \forall k \in K,\ a \in A_k,\ d \in D_{a,k},$$

where $0 \leq \lambda \leq 1$ is a blending coefficient: larger values place more weight on observed district demand, while smaller values place more weight on the equal-share capacity term. Regime-dependent outbound coefficients are

$$\kappa^P_{a,d,k} = 0, \qquad \kappa^R_{a,d,k} = \max\{0,\ 0.5 - \Gamma_{a,k}\}, \qquad \kappa^J_{a,d,k} = \max\{0.2,\ 1 - \Gamma_{a,k}\},$$

for all $k \in K$, $a \in A_k$, and $d \in D_{a,k}$. The constants 0.5 and 0.2 are exogenous case-study calibration parameters used to differentiate outbound protection across regimes.

Cost parameters are $c^{self}_{a,d,k}$ for self-service cost, $c^{tr}_{a,b,d,k}$ for transfer cost, $p^{un}_{a,d,k}$ for unmet-demand penalty, $p^{ov}_{a,k}$ for capacity overrun penalty, $F^R_{a,k}$ and $F^J_{a,k}$ for fixed regime costs (e.g., coordination and administrative overhead), $G^P_{a,k}$, $G^R_{a,k}$, and $G^J_{a,k}$ for contextual regime penalties (e.g., penalties reflecting congestion, overlap, or shared pressure), and $M_{a,d,k}$ for a valid upper bound on inbound collaboration received by actor a in district d on day k.

Decision variables. Regime choice is represented by

$$z^P_{a,k},\ z^R_{a,k},\ z^J_{a,k} \in \{0,1\} \qquad \forall k \in K,\ a \in A_k,$$

with exactly one regime selected for each actor-day. District openness is

$$\alpha_{a,d,k} \in [0,1] \qquad \forall k \in K,\ a \in A_k,\ d \in R_{a,k},$$

where $\alpha_{a,d,k}$ is the fraction of the stabilized district execution envelope opened by actor a to inbound collaborative work in district d on day k. For all $k \in K$, $a \in A_k$, and $d \in D_{a,k}$, let $s_{a,d,k} \geq 0$ be the self-served quantity and $u_{a,d,k} \geq 0$ the unmet demand. For each feasible transfer tuple (a,b,d,k), let $x_{a,b,d,k} \geq 0$ be the quantity of actor a's demand in district d executed by actor b. For all $k \in K$ and $a \in A_k$, let $\omega_{a,k} \geq 0$ denote capacity overrun.

Objective function. The first-level objective is

$$\min Z^{(1)}$$

with

$$Z^{(1)} = \underbrace{\sum_{k\in K}\sum_{a\in A_k}\sum_{d\in D_{a,k}} c^{self}_{a,d,k} s_{a,d,k}}_{\text{self-service cost}} + \underbrace{\sum_{k\in K}\sum_{a\in A_k}\sum_{d\in D_{a,k}}\sum_{b\in A_{d,k}\setminus\{a\}} c^{tr}_{a,b,d,k} x_{a,b,d,k}}_{\text{transfer cost}}$$
$$+ \underbrace{\sum_{k\in K}\sum_{a\in A_k}\sum_{d\in D_{a,k}} p^{un}_{a,d,k} u_{a,d,k}}_{\text{unmet-demand penalty}} + \underbrace{\sum_{k\in K}\sum_{a\in A_k} p^{ov}_{a,k}\omega_{a,k}}_{\text{overrun penalty}}$$
$$+ \underbrace{\sum_{k\in K}\sum_{a\in A_k}\left[G^P_{a,k} z^P_{a,k} + \left(F^R_{a,k}+G^R_{a,k}\right) z^R_{a,k} + \left(F^J_{a,k}+G^J_{a,k}\right) z^J_{a,k}\right]}_{\text{regime-selection penalty}}.$$

Constraints. Exactly one regime is chosen per actor-day:

$$z^P_{a,k} + z^R_{a,k} + z^J_{a,k} = 1 \qquad \forall k \in K,\ a \in A_k.$$

For every actor-day receiving district, openness is bounded by the selected regime:

$$\alpha_{a,d,k} \geq \underline{\alpha}^P z^P_{a,k} + \underline{\alpha}^R z^R_{a,k} + \underline{\alpha}^J z^J_{a,k} \qquad \forall k \in K,\ a \in A_k,\ d \in R_{a,k},$$

$$\alpha_{a,d,k} \leq \overline{\alpha}^P z^P_{a,k} + \overline{\alpha}^R z^R_{a,k} + \overline{\alpha}^J z^J_{a,k} \qquad \forall k \in K,\ a \in A_k,\ d \in R_{a,k}.$$

The openness variable $\alpha_{a,d,k}$ is defined over $R_{a,k}$ because it controls inbound collaborative work, i.e., the districts in which actor a may receive external demand, rather than the districts in which it owns demand directly. Demand balance is

$$s_{a,d,k} + \sum_{b\in A_{d,k}\setminus\{a\}} x_{a,b,d,k} + u_{a,d,k} = q_{a,d,k} \qquad \forall k \in K,\ a \in A_k,\ d \in D_{a,k}.$$

Transfer feasibility is enforced by

$$x_{a,b,d,k} \leq \beta_{a,b,d,k}\, q_{a,d,k} \qquad \forall k \in K,\ a \in A_k,\ d \in D_{a,k},\ b \in A_{d,k} \setminus \{a\}.$$

Private mode forbids any exchange with other actors, both outbound and inbound:

$$\sum_{b\in A_{d,k}\setminus\{a\}} x_{a,b,d,k} \leq q_{a,d,k}(1 - z^P_{a,k}) \qquad \forall k \in K,\ a \in A_k,\ d \in D_{a,k},$$

and

$$\sum_{i\in A_{d,k}\setminus\{a\}} x_{i,a,d,k} \leq M_{a,d,k}(1 - z^P_{a,k}) \qquad \forall k \in K,\ a \in A_k,\ d \in R_{a,k}.$$

District-level openness limits inbound collaborative load:

$$\sum_{i\in A_{d,k}\setminus\{a\}} x_{i,a,d,k} \leq \alpha_{a,d,k}\bar{Q}_{a,d,k} \qquad \forall k \in K,\ a \in A_k,\ d \in R_{a,k}.$$

Actor-day execution capacity requires

$$\sum_{d \in D_{a,k}} s_{a,d,k} + \sum_{d \in R_{a,k}} \sum_{i \in A_{d,k} \setminus \{a\}} x_{i,a,d,k} \leq Q_{a,k} + \omega_{a,k} \qquad \forall k \in K,\ a \in A_k.$$

Regime-dependent inbound collaboration is capped by

$$\sum_{d \in R_{a,k}} \sum_{i \in A_{d,k} \setminus \{a\}} x_{i,a,d,k} \leq \eta_R Q_{a,k} z^R_{a,k} + \eta_J Q_{a,k} z^J_{a,k} \qquad \forall k \in K,\ a \in A_k,$$

and regime-dependent outbound collaboration by

$$\sum_{d \in D_{a,k}} \sum_{b \in A_{d,k} \setminus \{a\}} x_{a,b,d,k} \leq \eta_R \left(\sum_{d \in D_{a,k}} q_{a,d,k} \right) z^R_{a,k} + \eta_J \left(\sum_{d \in D_{a,k}} q_{a,d,k} \right) z^J_{a,k} \qquad \forall k \in K,\ a \in A_k.$$

Finally, concentration-sensitive district outbound protection is enforced by

$$\sum_{b \in A_{d,k} \setminus \{a\}} x_{a,b,d,k} \leq q_{a,d,k} \left(\kappa^P_{a,d,k} z^P_{a,k} + \kappa^R_{a,d,k} z^R_{a,k} + \kappa^J_{a,d,k} z^J_{a,k} \right) \qquad \forall k \in K,\ a \in A_k,\ d \in D_{a,k}.$$

Level 1 therefore defines the collaboration architecture of the system under contextual, concentration, and capacity constraints, before temporal synchronization is examined in Level 2.

3.2 Level 2: Coordination feasibility layer

Level 2 takes as input the collaborative flows accepted at Level 1 and allocates them to a small set of standardized coordination windows. Its role is not to redefine regimes or openness, but to test whether the collaboration pattern admitted upstream can be synchronized operationally once sender-side, receiver-side, and district-level coordination limits are enforced.

Sets and indices. Let W be the set of coordination windows and F the set of collaborative flows accepted at Level 1. For each flow $f \in F$, let $o(f)$ and $r(f)$ denote its sending and receiving actors, let $d(f)$ denote its district, let $k(f)$ denote its day, and let v_f be its accepted transfer volume. For each actor-day (a, k), let $W^{acc}_{a,k} \subseteq W$ be the set of coordination windows accessible under the Level 1 regime. A flow can only be synchronized in windows accessible to both its sender and receiver.

Parameters. Let $\bar{C}^{in}_{a,k,w} \geq 0$ denote the inbound coordination capacity of actor a in window w on day k, and let $\bar{C}^{out}_{a,k,w} \geq 0$ denote the outbound coordination capacity of actor a in window w on day k. Let $\bar{M}_{a,d,k} \geq 0$ denote the district-day inbound envelope inherited from Level 1. Let c_w be the cost of assigning flow volume to window w. Let ρ be the penalty weight on flow-window activation, ϕ^{in} the penalty weight on peak inbound coordination load, ϕ^{out} the penalty weight on peak outbound coordination load, and ψ the penalty weight on unscheduled accepted collaboration.

Decision variables. For each $(f, w) \in F \times W$,

$$y_{f,w} \geq 0 : \text{ volume of flow } f \text{ synchronized in window } w,$$

$$z_{f,w} \in \{0, 1\} : \text{ equals 1 if flow } f \text{ is active in } w.$$

For each flow $f \in F$,

$$u_f \geq 0: \text{ unscheduled remainder of accepted flow } f.$$

For each actor-day (a, k),

$$p^{in}_{a,k} \geq 0: \text{ peak inbound coordination load over all windows,}$$

$$p^{out}_{a,k} \geq 0: \text{ peak outbound coordination load over all windows.}$$

Objective function.

$$\min Z^{(2)}$$

with

$$Z^{(2)} = \psi \sum_{f \in F} u_f + \sum_{f \in F} \sum_{w \in W} c_w \, y_{f,w} + \rho \sum_{f \in F} \sum_{w \in W} z_{f,w} + \phi^{in} \sum_{a,k} p^{in}_{a,k} + \phi^{out} \sum_{a,k} p^{out}_{a,k}.$$

The first term penalizes the part of accepted collaboration that cannot be synchronized in Level 2. The second term favors earlier coordination windows through the window-dependent assignment costs c_w. The third term discourages unnecessary fragmentation of flows across multiple windows. The last two terms smooth coordination peaks on the receiving and sending sides, so the objective balances service completion, temporal simplicity, and coordination stability.

Each accepted collaboration flow $f \in F$ with volume v_f is decomposed into a synchronized portion assigned to the coordination windows and, if necessary, an unscheduled remainder:

$$\sum_{w \in W} y_{f,w} + u_f = v_f \qquad \forall f \in F.$$

A flow can carry volume in a window only if that flow-window pair is activated:

$$y_{f,w} \leq v_f \, z_{f,w} \qquad \forall f \in F, \; w \in W.$$

Window accessibility is determined jointly by the sender and receiver. A flow can only be synchronized in windows that are accessible to both actors under their Level 1 regimes:

$$y_{f,w} = 0, \qquad z_{f,w} = 0 \qquad \forall f \in F, \; w \notin \big(W^{acc}_{o(f),k(f)} \cap W^{acc}_{r(f),k(f)}\big).$$

Receiving actor-window capacity is enforced by:

$$\sum_{f \in F:\, r(f)=a,\; k(f)=k} y_{f,w} \leq \bar{C}^{in}_{a,k,w} \qquad \forall a, \; k, \; w \in W.$$

Sending actor-window capacity is enforced symmetrically:

$$\sum_{f \in F:\, o(f)=a,\; k(f)=k} y_{f,w} \leq \bar{C}^{out}_{a,k,w} \qquad \forall a, \; k, \; w \in W.$$

The district-day inbound envelope inherited from Level 1 must also be respected:

$$\sum_{f \in F:\, r(f)=a,\, d(f)=d,\, k(f)=k} \sum_{w \in W} y_{f,w} \le \bar{M}_{a,d,k} \qquad \forall a,\ d,\ k.$$

Peak inbound coordination load is defined by:

$$p^{in}_{a,k} \ge \sum_{f \in F:\, r(f)=a,\, k(f)=k} y_{f,w} \qquad \forall a,\ k,\ w \in W,$$

and peak outbound coordination load by:

$$p^{out}_{a,k} \ge \sum_{f \in F:\, o(f)=a,\, k(f)=k} y_{f,w} \qquad \forall a,\ k,\ w \in W.$$

Level 2 therefore converts accepted collaboration into synchronized collaboration under inherited openness limits and bilateral coordination feasibility. Unlike Level 1, which governs whether collaboration is admissible, Level 2 reveals what portion of that admitted collaboration remains synchronizable once temporal coordination constraints are imposed.

4 Case study, experimental design, and results

The case study is built from a real parcel-level cargo-bike delivery dataset collected from several urban last-mile actors. The data include parcel coordinates and district identifiers, thereby preserving the observed spatial organization of demand and service activity. This provides a data-grounded basis for studying collaboration, spatial overlap, and distributed execution in a PI setting.

To evaluate the proposed framework, we aggregate the observed parcel geography into a stylized four-actor setting over six microhubs and a common district partition. Actors represent LSPs, microhubs are local operating bases from which they serve parcels, and districts are the spatial units in which demand ownership and collaboration opportunities are measured. The resulting case exhibits heterogeneous daily load conditions, partially overlapping service areas, and uneven capacity pressure across hubs. It is therefore well suited for testing both the regime-selection decisions of Level 1 and the execution-feasibility layer of Level 2.

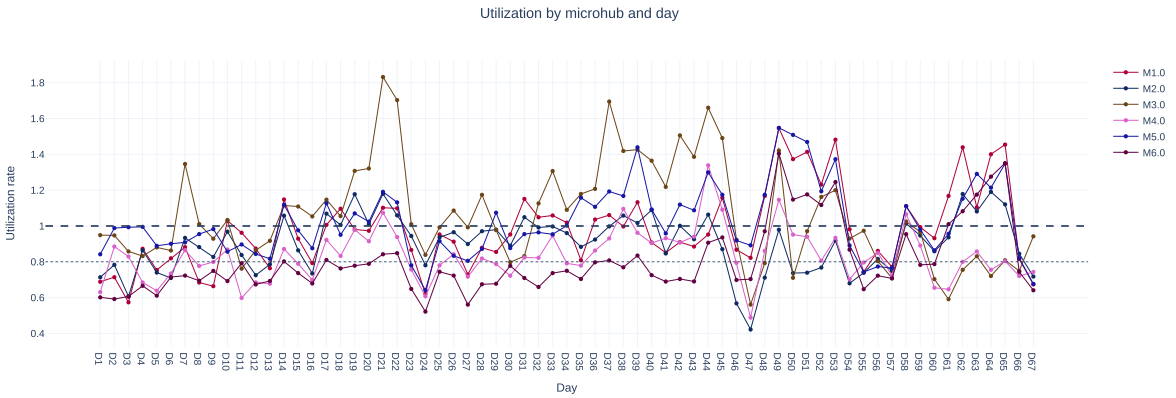


Figure 1: Daily utilization profile by microhub over 67 operating days.

Figure 1 reports the daily utilization profile of the six microhubs over 67 operating days. Here, utilization is defined as daily load divided by nominal daily capacity, so values above 1 indicate overload; correspondingly, overload in Table 1 denotes the excess parcels above nominal capacity. The figure highlights clear temporal heterogeneity across hubs. Some remain persistently close to full capacity (utilization $= 1$), while others alternate more sharply between buffered and overloaded states. This temporal variability shows that coordination needs are not uniform across the network and should not be governed by a single fixed openness rule.

Table 1: Microhub-level overview of the case study over 67 operating days.

Hub	Fleet (bike)	Mean load (parcels/day)	Capacity (parcels/day)	Mean util.	Max overload (parcels)	Overloaded days	Mean shared pressure	Mean overlap
M1	6	1864.6	1907	0.978	1042	26	0.310	0.432
M2	6	1759.3	1955	0.900	372	18	0.284	0.508
M3	4	1264.9	1196	1.058	995	33	0.452	0.568
M4	4	1160.7	1394	0.833	472	6	0.534	0.607
M5	3	1060.7	1038	1.022	569	29	0.468	0.547
M6	9	2687.7	3346	0.803	1353	10	0.290	0.381

Table 1 summarizes the hub-level characteristics. The case is clearly heterogeneous. Hubs M3 and M5 are the most capacity-constrained on average, with mean utilization above one and the largest numbers of overloaded days. M3 is the most critical hub overall, combining the highest number of overloaded days with high shared-pressure and overlap levels. M5 also remains persistently tight, although under a smaller fleet and lower absolute load. By contrast, M4 operates below full average utilization but exhibits the strongest mean shared pressure and overlap, indicating strong structural exposure to collaboration opportunities even when direct congestion is less severe. M1 is also relatively tight, with frequent overload and high peak excess demand. M6 carries the largest average volume and the largest absolute overload peak, but its larger fleet and capacity keep its average utilization below the most stressed hubs. M2 occupies an intermediate position, with high average utilization and recurrent overload, but less pronounced structural overlap than the most exposed hubs.

Overall, the case combines three properties that are central to the proposed framework: recurrent but uneven congestion, non-uniform spatial overlap, and strong heterogeneity in local exposure to shared demand. This is precisely the type of environment in which collaboration openness should not be fixed uniformly, but selected conditionally according to local operating conditions. The first-level model is therefore evaluated on a setting where some hubs have stronger incentives to remain protected, while others present greater potential for selective or broader collaboration. The second-level model then examines whether the accepted collaborative workload can be synchronized through feasible coordination windows.

4.1 Level 1 results

Level 1 determines, for each actor-day, the collaboration regime and the admissible degree of district-level openness. This is the governance layer introduced in Section 3: before any transfer is realized, the model decides whether collaboration should remain protected, bounded by rules, or opened more broadly. In the case study, this first-level problem is solved over 268 actor-days, corresponding to four active actors observed over the 67-day horizon. The main result is that the solution does not converge to a single openness logic. Instead, it separates the system into three

distinct governance states: private, rule-based, and joint. As reported in Table 2, private decisions represent 38.8% of actor-days, rule-based decisions 37.7%, and joint decisions 23.5%. This is an important PI result. The network is not managed as uniformly open, but as selectively interoperable: openness is activated where it is operationally absorbable and contained where local conditions remain too exposed. The table also shows a clear ordering across regimes. Private actor-days are associated with the highest utilization, shared pressure, and overlap, which indicates that the model protects local execution under the most constrained states. Rule-based collaboration occupies an intermediate position and preserves collaboration as a controlled option rather than a default behavior. Joint decisions are less frequent, but they correspond to the most collaborative part of the solution. Overall, 44.0% of actor-days generate positive collaboration flows, representing about 20,191 transferred parcels over the horizon. Collaboration is therefore substantial, but not universal.

Table 2: Level 1 regime outcomes at the actor-day level.

Regime	Count	%	Util.	Pressure	Overlap	Collab. (%)
Private	104	38.8	1.153	0.390	0.521	0.0
Rule-based	101	37.7	0.893	0.300	0.431	54.5
Joint	63	23.5	0.917	0.329	0.447	100.0

Figure 2 reports the resulting hub-day projection over the full microhub-day grid. It provides a compact view of how selective openness is distributed across hubs and days. This first-level result prepares the second level of the analysis. Once the governance layer has differentiated where collaboration is protected (P), bounded (R), or broadly enabled (J), the next question is how much exchange is actually realized across districts and actors. Section 4.2 addresses this operational layer.

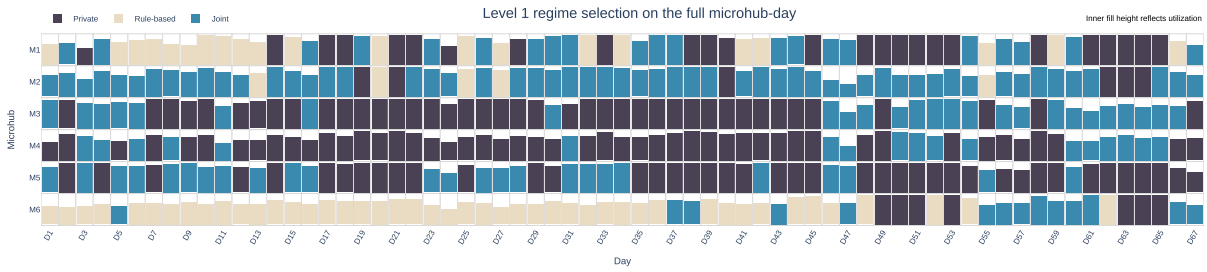


Figure 2: Hub-day projection of Level 1 regime selection over the full microhub-day grid.

4.2 Level 2 results

Level 2 takes the collaborative exchanges admitted by Level 1 and tests whether they can be synchronized over a small number of standardized intra-day coordination windows. These windows do not represent exact dispatch schedules. Rather, they approximate recurring exchange opportunities during the day: an early coordination period, a main mid-day coordination period, and a late balancing period. In that sense, Level 2 is not another governance layer, but an operational consistency check: once sender-side limits, receiver-side limits, and district-level openness inherited from Level 1 are enforced over this window structure, not all accepted collaboration remains executable.

This distinction is visible in the results. Out of about 20,191 parcels accepted for collaboration by Level 1, Level 2 successfully synchronizes about 19,072, leaving about 1,120 unscheduled. The resulting service rate is therefore 94.5% in volume terms. The two-level structure is thus justified

empirically: the governance layer admits collaboration selectively, while the synchronization layer reveals the operational friction that remains once temporal coordination constraints are imposed.

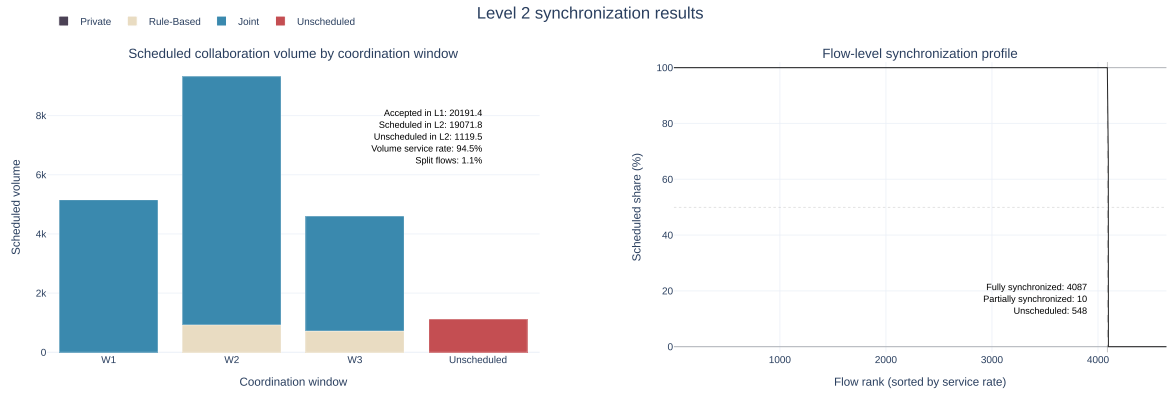


Figure 3: Level 2 synchronization results under standardized intra-day coordination windows.

Figure 3 summarizes this second-stage outcome. The left panel shows how synchronized collaboration is distributed across the three intra-day coordination windows. Most of the realized exchange is concentrated in the first two windows, with W2 acting as the main coordination basin and W3 providing a smaller balancing role. The right panel reports the flow-level service profile. Most accepted flows are fully synchronized, a very small fraction is only partially synchronized, and a limited residual set remains unscheduled. This confirms that Level 2 is not a simple repartition device: it filters a non-negligible share of the collaboration accepted upstream.

From a Physical Internet perspective, this result is important. Selective openness at Level 1 is not enough on its own; interoperability must also remain temporally synchronizable once concrete exchange opportunities are imposed. The Level 2 results show that the large majority of accepted collaboration survives this test, but not all of it. In other words, the system remains strongly collaborative, yet operationally disciplined: broad interoperability is achievable, but only up to the point where synchronization frictions begin to bind.

This second-stage result also prepares the final interpretation of the framework. Level 1 determines where collaboration should be opened, bounded, or protected. Level 2 then identifies which part of that accepted collaborative volume can actually be coordinated in time. Together, the two levels move the analysis from *admissible openness* to *operationally synchronized collaboration*.

5 Conclusion

This paper examined how collaboration should be governed in a Physical Internet setting when openness is not assumed to be universally desirable. Using a two-level formulation, we separated the problem into a governance layer and a coordination feasibility layer. The first level determines whether actor-days should remain protected, operate under bounded rule-based sharing, or enter broader joint collaboration. The second level tests whether the accepted collaboration can actually be synchronized once bilateral window constraints are imposed. The case study shows that the solution does not converge to a single openness logic. Instead, it produces selective interoperability: collaboration is substantial, but differentiated across actors, districts, and days.

For the PI community, the main message is that openness should not be treated as a fixed structural

property of the network. It should be governed conditionally. In our results, protected, bounded, and broader collaboration states coexist within the same system, and the second level further shows that not all accepted collaboration remains synchronizable in operation. This reinforces the idea that interoperability in PI systems is not simply a matter of allowing exchange, but also of preserving operational feasibility once exchanges must be coordinated in time.

The study remains subject to several simplifying assumptions. The case is based on a synthetic setting calibrated from observed hub-day conditions, the coordination layer uses standardized windows rather than continuous-time dispatch, and the parameterization of regime penalties and window priorities remains hypothesis-driven. These assumptions were useful to isolate the logic of selective openness, but they also limit direct operational generalization.

Future work should therefore extend the framework in three directions: calibration on richer real collaborative data, sensitivity analysis on the governance and synchronization parameters, and a more detailed execution layer with continuous time, routing interactions, and stochastic disruptions. More broadly, the model could be extended to a multi-actor PI setting in which governance rules evolve endogenously with repeated collaboration.

References

Ballot, E., Montreuil, B., and Zacharia, Z. G. (2021). Physical internet: First results and next challenges. *Journal of Business Logistics*, 42(1).

Bertsimas, D. and Kallus, N. (2020). From predictive to prescriptive analytics. *Management Science*, 66(3):1025–1044.

Boysen, N., Briskorn, D., Montreuil, B., and Zey, L. (2025). The π-transportation problem: On the value of split transports for the physical internet concept. *European Journal of Operational Research*, 324(2):629–643.

Crainic, T. G., Klibi, W., and Montreuil, B. (2023). Hyperconnected city logistics: A conceptual framework. In *Handbook on City Logistics and Urban Freight*, pages 398–421. Edward Elgar Publishing.

Crainic, T. G. and Montreuil, B. (2016). Physical internet enabled hyperconnected city logistics. *Transportation Research Procedia*, 12:383–398.

Fernández, E., Roca-Riu, M., and Speranza, M. G. (2018). The shared customer collaboration vehicle routing problem. *European Journal of Operational Research*, 265(3):1078–1093.

Gansterer, M. and Hartl, R. F. (2018). Collaborative vehicle routing: A survey. *European Journal of Operational Research*, 268(1):1–12.

Guajardo, M., Rönnqvist, M., Flisberg, P., and Frisk, M. (2018). Collaborative transportation with overlapping coalitions. *European Journal of Operational Research*, 271(1):238–249.

Labarthe, O., Ahmadi, G., Klibi, W., Deschamps, J. C., and Montreuil, B. (2024). A sustainable on-demand urban delivery service enabled by synchromodality and synergy in passenger and freight mobility. *Transportation Research Part C: Emerging Technologies*, 161:104544.

Mancini, S., Gansterer, M., and Hartl, R. F. (2021). The collaborative consistent vehicle routing problem with workload balance. *European Journal of Operational Research*, 293(3):955–965.

Montreuil, B. (2011). Toward a physical internet: Meeting the global logistics sustainability grand challenge. *Logistics Research*, 3(2):71–87.

Montreuil, B., Meller, R. D., and Ballot, E. (2012). Physical internet foundations. In *IFAC Proceedings Volumes*, volume 45, pages 26–30.

Pan, S., Ballot, E., Huang, G. Q., and Montreuil, B. (2017). Physical internet and interconnected logistics services: Research and applications. *International Journal of Production Research*, 55(9):2603–2609.

Sadana, U., Chenreddy, A., Delage, E., Forel, A., Frejinger, E., and Vidal, T. (2025). A survey of contextual optimization methods for decision-making under uncertainty. *European Journal of Operational Research*, 320(2):271–289.

Digitalizing a Territorial Fresh-Produce Wholesale Market as a Physical-Internet Node: Container-Centric Traceability, Evidence Packages, and Controlled Transparency at Mercato Ortofrutticolo del Roero

Claudio Salvadori[1], Franco Maciariello[1,2,3], Eddi Lorenzi[1], Massimo Torchio[4], Silvio Abrate[5]

[1] New Generation Sensors srl, Pisa, Italy

[2] Universitas Mercatorum, Rome, Italy

[3] CSAIA, Rome, Italy

[4] Azienda Consortile Mercato Ortofrutticolo del Roero, Canale (CN), Italy

[5] LINKS Foundation, Torino, Italy

Corresponding author: claudio.salvadori@ngs-sensors.it

Abstract:

Wholesale fresh-produce markets are critical territorial infrastructures, yet they remain difficult environments for Physical Internet deployment because operations are time-compressed, multi-actor, temperature-sensitive, and only partially digitalized. We present the design of a container-centric digitalization approach for Mercato Ortofrutticolo del Roero, framed as a Physical-Internet node. The approach combines reusable-crate identification, zone-based location awareness, EPCIS 2.0 event capture, auditable outbound aggregation, shipment evidence packages, and tiered data sharing. Our contribution is threefold. First, we formalize the reusable crate as the smallest accountable PI container in a territorial wholesale market. Second, we show how event-standardization and evidence packages can preserve accountability during consolidation without forcing full transparency across competing actors. Third, we define a deployment and evaluation protocol for a real SME-dense fresh-food hub, explicitly distinguishing the capture/interoperability layer from the market's existing commercial and administrative systems. The paper contributes a rigorous design and replication blueprint for PI-compatible traceability in perishable logistics, while clarifying scope, interfaces, figures, and literature positioning in response to prior review feedback.

Keywords:

Physical Internet; wholesale fresh-produce markets; reusable crates; EPCIS 2.0; cold-chain traceability; evidence packages; controlled transparency; circular economy.

Physical Internet (PI) Roadmap Fitness: ☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories), ☐ Transportation Equipment, ☐ PI Networks, ☐ System of Logistics Networks, ☒ Vertical Supply Consolidation, ☐ Horizontal Supply Chain Alignment, ☒ Logistics/Commercial Data Platform, ☒ Access and Adoption, ☒ Governance.

Targeted Delivery Mode-s: ☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation

1. Introduction

Wholesale fresh-produce markets sit at a difficult intersection of logistics, commerce, and food quality assurance. They concentrate highly perishable flows, serve heterogeneous actors, and must operate with rapid receiving, handling, storage, picking, weighing, dispatch, and return processes. In this setting, small execution errors can quickly propagate into food losses, cold-chain failures, service delays, and disputes over custody or product integrity. These operational pressures are especially visible in territorial wholesale markets, where a dense ecosystem of small and medium-sized enterprises interacts through partially standardized and partially informal procedures.

The Physical Internet (PI) offers an appealing long-term vision for such environments. PI promotes modular containers, standardized interfaces, interoperable data exchange, and shared

logistics assets so that flows can be consolidated while accountability is preserved (Montreuil, 2011; Treiblmaier et al., 2020). For fresh food, this vision has been extended toward architectures that explicitly integrate perishability, freshness, and sensor-enabled decision making (Pal and Kant, 2019; Pal et al., 2017). However, most PI literature still focuses on network-level concepts, optimization, facility classes, or large-scale distribution design. Comparatively less attention has been devoted to territorial wholesale markets as PI nodes, even though these facilities are natural points of physical consolidation, temporary buffering, quality control, and returnable-packaging circulation.

A closely related industrial concept is the *GS1 Smart Box*[1], which proposes that reusable logistics containers become digitally identifiable and data-enabled transport units. In this model, each container carries a unique identity and can generate standardized visibility events across logistics processes, enabling interoperability between actors while maintaining data governance boundaries. The Smart Box concept operationalizes key PI principles by treating containers as information carriers capable of recording custody changes, condition data, and logistics events throughout their lifecycle. In practice, such containers act as modular digital entities that can move across heterogeneous supply-chain environments while preserving traceability, accountability, and compatibility with GS1 standards-based data exchange.

This perspective is particularly relevant for fresh-produce markets, where the operational handling unit is often the reusable crate rather than the pallet or shipment. Conceptually, reusable crates can be interpreted as the smallest accountable container in a PI-compatible logistics environment, capable of capturing granular operational events while still supporting higher-level aggregation into outbound shipments.

At the same time, research on agri-food traceability and cold-chain management has established the importance of RFID, standardized visibility events, and environmental monitoring for perishable products (Mainetti et al., 2013; Thakur and Forås, 2015; Mustafa et al., 2024). Yet three practical gaps remain. First, existing traceability systems often stop at lot or shipment level, while internal market operations depend on smaller modular handling units such as reusable crates. Second, visibility data are rarely organized in a way that supports selective sharing among competing actors. Third, the integration challenge in mixed human-digital settings is often underestimated: markets need systems that coexist with existing weighing, lot generation, and administrative applications rather than replacing them wholesale.

This paper addresses these gaps through the design of a container-centric digitalization architecture for Mercato Ortofrutticolo del Roero, a territorial wholesale fresh-produce market in Piedmont, Italy. The paper is a design-oriented scientific case study. It does not claim post-deployment impact results that are not yet available. Instead, it contributes a technically grounded architecture, a clarified operational workflow, and a deployment-oriented evaluation protocol for a real 2026 implementation. This positioning is deliberate: it responds to the need for stronger literature anchoring, clearer figures, and a more explicit explanation of the use case and evaluation logic.

The paper makes three contributions. First, it conceptualizes the reusable crate as the smallest accountable PI container in a wholesale-market context, with outbound logistics units represented as temporary, traceable aggregations. Second, it integrates GS1 EPCIS 2.0 event capture with environmental sensing and evidence packages so that cold-chain integrity can be documented without exposing the full internal dataset to all actors. Third, it specifies a practical implementation boundary between the interoperability layer and the existing market systems, making the proposed approach more realistic for replication in SME-dense fresh-food ecosystems.

2. Literature background and research positioning

[1] https://www.gs1belu.org/en/gs1-smart-box-0

Building on these research streams, this paper focuses on an operational context that remains relatively underexplored in Physical Internet research: territorial wholesale fresh-produce markets. These facilities operate as natural consolidation points for regional agri-food flows, yet their daily execution still relies on a mix of manual practices, legacy weighing systems, and partially digitized traceability processes.

Rather than proposing a general PI architecture or reporting ex post performance results, the paper develops a deployment-oriented design for transforming such a market into a PI-compatible node. The proposed approach integrates five operational elements within a coherent container-centric architecture: reusable-crate accountability, explicit crate-to-unit aggregation, zone-based cold-room visibility, evidence-package generation for outbound shipments, and controlled transparency between actors.

Table 1 positions this contribution relative to representative literature across the Physical Internet, perishable logistics, and agri-food traceability domains.

Table 1: Literature positioning of the proposed container-centric wholesale-market architecture.

Research Stream	Representative Literature	Primary Focus	Limitations for the considered Context	Contribution of This Paper
Physical Internet logistics architectures	Montreuil (2011) Montreuil et al. (2010); Treiblmaier et al. (2020)	Co nceptualization of modular logistics networks, standardized containers, and open logistics infrastructures	Emphasis on network design, infrastructure classes, and macro-level logistics coordination	Applies PI principles to a territorial wholesale market node with real operational constraints
PI for perishable and fresh-food logistics	Pal et al. (2017); Pal and Kant (2019)	Integration of freshness-aware routing, sensing technologies, and cyber-physical coordination	Focus on network optimization and freshness modeling rather than facility-level execution	Shows how perishability-aware logistics can be operationalized inside a wholesale facility
Agri-food traceability and RFID systems	Mainetti et al. (2013); Razak et al. (2023)	RFID-based tracking, traceability architectures supply-chain visibility	Traceability mainly implemented at lot or shipment level, not at modular crate level	Introduces crate-level accountability as the smallest traceable PI container
Event-based logistics visibility (GS1 EPCIS)	GS1 (2022a); GS1 (2023); Thakur and Forås (2015)	Standardized event models for interoperability and cold-chain monitoring	Focus on data standards rather than integration with mixed manual-digital operations	Combines EPCIS events with operational workflows of receiving, cold storage, picking, and dispatch
Cold-chain logistics and sensing	Mustafa et al. (2024)	Integration of IoT sensing and process monitoring in food logistics	Limited discussion of selective data sharing among competing actors	Introduces evidence packages and controlled transparency mechanisms

3. Research setting and method

The research adopts a design-oriented case-study approach aimed at specifying and preparing the deployment of a Physical Internet-compatible digitalization architecture within a real fresh-produce wholesale market. The empirical setting is Mercato Ortofrutticolo del Roero, a territorial market in Northern Italy that coordinates daily interactions between producers, buyers, transport operators, and market personnel. Operational activities include rapid inbound receiving, temporary storage, sales preparation, outbound dispatch, and empty-crate return. These characteristics make the market a natural candidate for implementation as a Physical Internet node, where modular logistics units circulate between heterogeneous actors under high operational tempo.

The proposed digitalization architecture introduces a container-centric interoperability layer without replacing the existing market information systems. Core business functions - such as lot accounting, quantity calculations, pricing, and administrative workflows - remain under the

responsibility of the market's current software environment. The new layer instead focuses on container identification, EPCIS-based event capture, data normalization, environmental sensing, interoperability services, and shareable evidence generation**.** Maintaining this architectural boundary is both scientifically and practically important: it clarifies the scope of the contribution while ensuring compatibility with incremental deployment in an SME-dense operational environment.

Within this architecture, operational processes are represented through a sequence of event-driven interactions covering inbound receiving, lot and weight linkage, cold-room zoning, crate-level picking, outbound consolidation, and empty-crate return. Section 4 reorganizes these components as a layered dedicated stack (that can be generalised for other wholesale fresh-produce markets) and positions them against the canonical OLI vocabulary, while Figure 1 details the operational workflow and the role of EPCIS events and controlled-transparency mechanisms in supporting the market's transformation into a PI-compatible node.

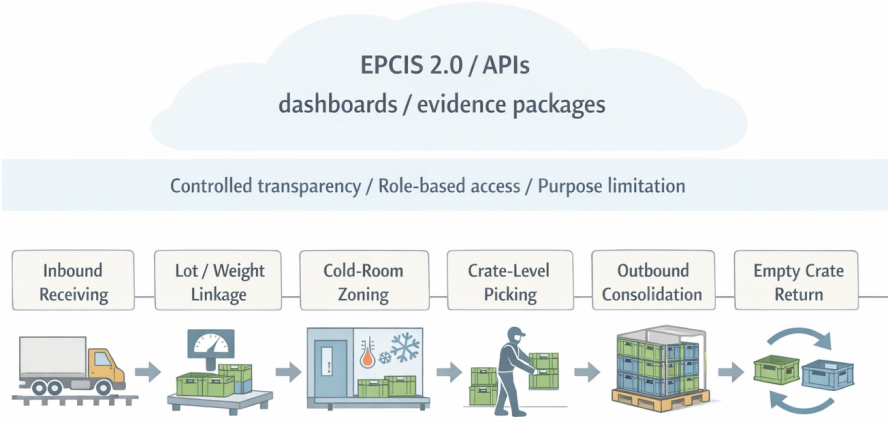


Figure 1: Operational workflow

The paper therefore reports a design artifact together with a deployment-oriented evaluation protocol**.** The evaluation protocol is structured around three research questions:

1. **Event representation:** Can the proposed EPCIS-based event model represent the physical and documentary reality of wholesale-market operations without losing traceability granularity during crate aggregation and shipment consolidation?
2. **Integrity verification:** Can the architecture support selective verification of outbound cold-chain integrity through the generation of shareable evidence packages derived from event and sensor data?
3. **Operational adoption:** Can the proposed solution be implemented with low operator friction in a mixed manual–digital environment typical of SME-intensive wholesale markets?

These questions are addressed through architecture specification, workflow analysis, and a deployment-oriented KPI framework rather than through fabricated performance outcomes.

4. Layered architecture of the Roero market

The Roero wholesale fresh-produce market is structured as a PI-node following the layered architecture depicted in Figure 2. Reading bottom-up, the proposed layered stack is composed by seven layers: (i) the **physical assets & handling layer**, thar considers reusable crates, cold rooms, and the receiving/preparation/dispatch areas in which crates are moved and stored; (ii) the **identification & capture layer**, that considers RFID (QR/barcode as fallback), IoT devices, and all the infrastructure to manage those; (iii) the **zone & location semantics layer,** through market zoning (identified by BLE beacons), and location identifiers used as read points or business location; (iv) **intra-market routing layer**, namely the operational progression from receiving to storage, picking, and dispatch under zone and dwell-time rules; (v) the **outbound shipping unit & selective QoS layer**, and its release logic, including SSCC-compatible

identification, dispatch closure, and QoS checks; (vi) the **event representation & encapsulation layer**, that includes aggregation, sensor reporting, and instance/lot master data exploiting EPCIS 2.0 and CBV standards; (vii) the **application & controlled transparency layer**, that exposes dashboards, APIs, evidence packages, and controlled-transparency services.



Figure 2: The layered wholesale fresh-produce

The proposed stack in can then be positioned analytically against the canonical Open Logistics Interconnection model (consider the label on the right side of Figure 2). The OLI model was introduced as a seven-layer framework for interconnecting logistics services in the Physical Internet and is commonly discussed in PI literature as a logistics analogue to layered communication models such as OSI (Montreuil et al., 2012; Pan et al., 2017; Treiblmaier, 2019). For engineering purposes, the ICONET reference architecture provides an explicit layer-by-layer interpretation of OLI functions, including physical handling, link, network, routing, shipping, encapsulation, and logistics-web services (ICONET, 2020). Particularly, Table 2 explain the proposed mapping, as an analytical correspondence, not as an ontological identity.

Table 2: Analytical mapping of the proposes layered architecture on the canonical OLI stack

OLI layer	Canonical role	Proposed implementation
L1 Physical	Physical handling of PI objects and assets	Reusable crates, cold rooms, receiving/prep/dispatch areas, handling operations
L2 Link	Reliable transfer/control between node actions and flow continuity	RFID/QR capture, IoT devices and infrastructure to manage those
L3 Network	Interoperability and common network semantics	Market zoning, BLE beacons, GLN/SGLN, readPoint/bizLocation semantics
L4 Routing	Routing decisions across network segments	Intra-market movement rules, picking flows, dwell-time constraints
L5 Shipping	Shipping status, QoS, and release logic	SSCC outbound unit, dispatch closure, selected shipment tracking/QoS checks
L6 Encapsulation	Composition/decomposition while preserving visibility	EPCIS 2.0 + CBV, AggregationEvent contains relations, SensorElement, ILMD
L7 Logistics Web	User-facing interfaces and applications	Dashboards, APIs, evidence packages, RBAC/purpose limitation

This interpretation strengthens the PI positioning of the case without claiming more than the evidence supports. The proposed architecture on the market is not modelled here as a general-purpose inter-network router; rather, it is modelled as a territorially bounded PI node whose local architecture can be discussed rigorously with OLI vocabulary. That distinction is important for the paper's scope, because it keeps the contribution anchored in a deployable case-study artifact while still connecting it to the canonical PI literature.

5. ICT architecture, implementation, and evidence-package logic

While Figure 1 provides the end-to-end operational workflow and Section 4 frames the Roero case through a layered PI/OLI lens, the present section summarizes the concrete ICT architecture and implementation logic that operationalizes that design. The emphasis is on a

retrofit-friendly capture, sensing, and interoperability layer that complements rather than replaces the market's existing commercial and administrative systems.

ICT architecture. NGS One Platform acts as the digital backbone of the proposed implementation, behaving as a multi-tenant and interoperable platform for transport-and-logistics flows that federates data from multiple operators and systems, normalizes them, and makes them selectively shareable while keeping each actor's information domain separated. This corresponds to the platform's stated capabilities of multi-actor federation, end-to-end visibility, exception management and decision-support alerts, reporting and auditability, and digital-twin representation. In the considered project, its role is therefore not to replace the market ERP or administrative software, but to collect data from heterogeneous IoT and operational sources, represent the supply chain in a standardized way, and expose the resulting visibility layer through dashboards, alerts, APIs, and evidence-oriented data services. At functional level, Figure 3 summarises this logic by connecting inbound receiving, cold-room/storage zoning, outbound consolidation, shipment monitoring, and interoperability services into a single architecture.

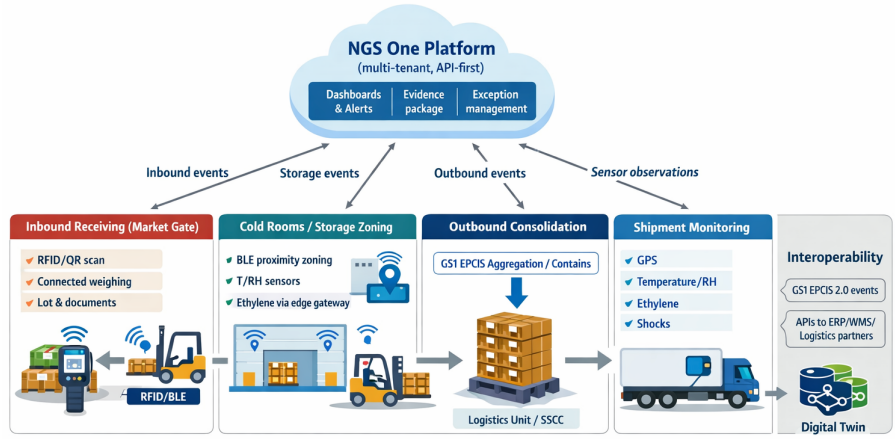


Figure 3: The complete ICT architecture

Functionally, the architecture begins at inbound receiving, where reusable crates and documents are linked through RFID capture and connected weighing. It continues through on-site storage zoning, where BLE-based proximity and environmental sensing provide the location and handling context of crates inside cold rooms. The same digital thread extends to outbound consolidation, where crates are explicitly associated with the dispatched logistics unit, and then to shipment monitoring for selected flows through GPS and environmental trackers. The interoperability block on the right of Figure 3 shows how these heterogeneous capture and sensing streams are transformed into standardized GS1 EPCIS 2.0 events, APIs, and a digital-twin representation that can be shared selectively with authorized stakeholders.

Technical implementation. At the field layer, reusable smart crates equipped with RFID tags are identified by *smart-forklifts* fitted with RFID readers that are connected via BLE to a dedicated on-board tablet.

In cold rooms, BLE beacons interacting with the forklift tablet enable proximity-based zoning and zone-level indoor localization by automatically associating each crate’s GRAI identifier, captured by the forklift-mounted RFID readers, with the relevant storage area. Temperature, relative humidity, and ethylene measurements are collected through sensing devices connected to an edge gateway. In this way, the solution maintains continuous awareness of both the location of each crate and the environmental conditions under which it is stored. The on-board tablet dispatches capture data through the on-site Wi-Fi network, while identification and sensing streams are also collected through an edge gateway. The gateway aggregates temperature/relative humidity monitoring (from BLE beacons) and ethylene measurements. Figure 4 details graphically the architecture and the devices’ interactions at the field layer.

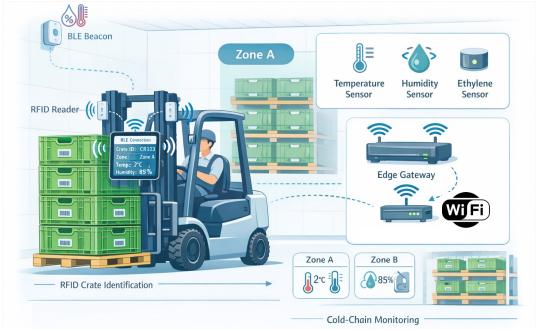


Figure 4: The field layer and the smart-forklift

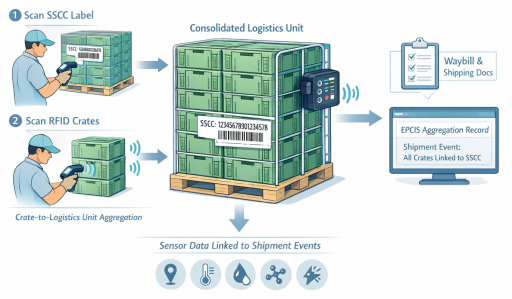


Figure 5: Outbound consolidation at dispatch

The same approach supports outbound consolidation by enabling operators to build complete EPCIS aggregation relationships from crate to logistics unit, typically identified by an SSCC, at dispatch, as depicted in Figure 5. Using the portable QR/RFID reader, the operator first scans the SSCC and then scans all crates included in that unit, generating a verifiable consolidation record that can be associated with the waybill and shipping documents. For selected outbound flows, the logistics unit can also be linked to a reusable ethylene-enabled tracker, so that shipment-integrity data (GPS, temperature/relative humidity, ethylene, and shocks) are automatically bound to the corresponding EPCIS shipment events.

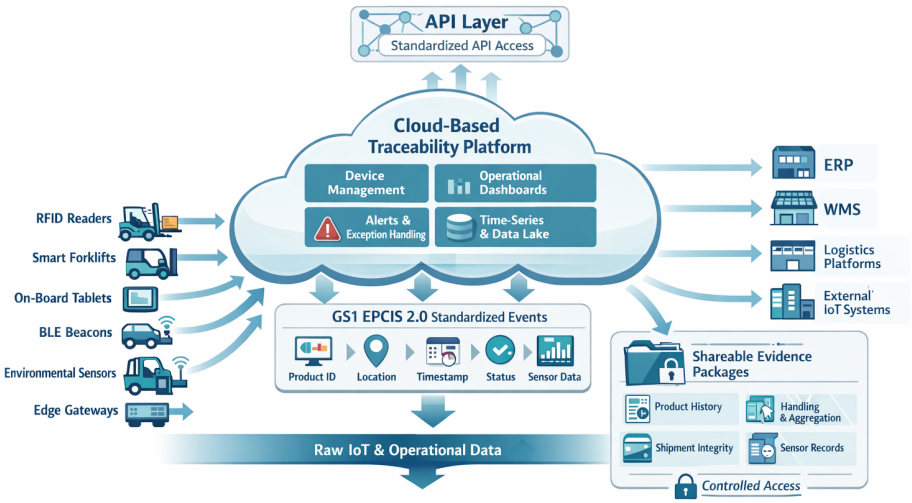


Figure 6: The Cloud platform layer

At the cloud layer, the platform provides a multi-tenant service environment that supports device management, operational dashboards, alerting functions, and persistent storage for both event streams and time-series observations. As depicted in Figure 6, this layer acts as the digital coordination point of the solution, collecting heterogeneous data from identification devices, environmental sensors, and field applications, and organizing them into a coherent representation of logistics operations. Beyond simple storage, the cloud platform supports monitoring, exception handling, and cross-actor visibility by combining operational status information with historical records of handling, storage, and shipment conditions.

A central function of this layer is the normalization of operational events and sensor observations into GS1 EPCIS 2.0, so that data generated by different devices and workflows can be represented in a standardized and interoperable form. Through this mechanism, the system does not merely store raw IoT signals but transforms them into shareable business events that describe what happened, to which logistics objects, where, when, and under which measured conditions. These standardized data are then exposed through APIs, enabling integration with internal enterprise systems and third-party applications such as ERP, WMS, logistics platforms. On this basis, the platform also supports the creation of shareable evidence

packages, which allow selected stakeholders to access auditable proofs of identity, aggregation, handling history, and shipment integrity under tiered access and controlled-sharing policies.

6. Evaluation protocol and discussion

Because the integrated deployment is planned for 2026, the evaluation contribution of this paper is a structured protocol rather than outcome claims. The protocol is designed to make the future implementation scientifically analysable and directly responds to the earlier concern that the experiment and discussion were underdeveloped.

The first evaluation dimension is representational adequacy. The question is whether the event model can reconstruct the operational history of a crate and of an outbound logistics unit from receiving to dispatch and return. The main indicators are event completeness, rate of unresolved exceptions, and proportion of outbound units with verifiable crate-to-unit closure.

The second evaluation dimension is integrity visibility. The question is whether selected outbound flows can be documented through evidence packages that are sufficiently interpretable for authorized downstream actors. Indicators include evidence-package completeness, time synchronization quality, percentage of monitored flows with usable telemetry, and turnaround time for incident reconstruction.

The third evaluation dimension is operator adoption. The market environment is characterized by time pressure and variable digital maturity, so usability cannot be an afterthought. The case therefore uses a simplified action model, centered on a small set of meaningful operator commands rather than open-ended data entry. Indicators include mean task time, manual override frequency, training time to proficiency, and user-reported ease of operation.

The fourth evaluation dimension is packaging-loop control. Since reusable crates are both operational assets and PI enablers, their circulation must be visible. Indicators include return-cycle completion, shrinkage reduction, and stock reconciliation between physical and digital crate pools.

Table 3 operationalizes these dimensions into a deployment KPI framework. The value of reporting this protocol now is that it makes the case falsifiable and reproducible. Future work can test whether the designed system improves cold-chain robustness, shrinkage control, and dispute handling once the Roero deployment is live.

Table 3: Deployment-oriented KPI framework for the Roero implementation

Evaluation dimension	Core question	Example indicators	Unit of analysis
Representational adequacy	Can the event model reconstruct crate and outbound-unit histories without ambiguity?	Event completeness; unresolved exceptions; % dispatches with closed crate-to-unit list	Crate, logistics unit
Integrity visibility	Can selected shipments be verified through usable evidence packages?	Evidence-package completeness; telemetry usability; incident reconstruction time	Monitored shipment
Operator adoption	Can the workflows be executed with low friction in time-pressured operations?	Task time; manual override rate; training time to proficiency; perceived ease of use	Operator / task
Packaging-loop control	Can reusable-crate circulation be monitored as a closed loop?	Return-cycle completion; shrinkage reduction; stock reconciliation accuracy	Crate pool

The case yields several insights for PI research and practice. First, territorial wholesale markets deserve more explicit treatment as PI nodes. They are not merely legacy intermediaries; they are places where modular units are reconfigured, checked, combined, and redistributed under strong temporal and quality constraints. This makes them an excellent test bed for PI deployment in environments that are neither fully centralized nor fully informal. Second, container granularity matters. In many fresh-produce contexts, the reusable crate is the true

operational unit of accountability. Treating the pallet or shipment as the primary object is insufficient when content is mixed, picked, weighed, and re-allocated during the day. Crate-level event capture therefore strengthens both consolidation quality and reverse-logistics control. Third, transparency must be designed, not assumed. A recurring weakness of digital-supply-chain visions is the implicit assumption that more visibility is always better. In federated markets, indiscriminate transparency can undermine adoption. The controlled-transparency model proposed here suggests that PI-compatible interoperability can coexist with selective disclosure, if evidence packages and access rules are explicitly modelled. Fourth, the case highlights an important architectural principle for real deployments: PI enablement should be separable from business-system replacement. By positioning the solution as a capture, normalization, sensing, and interoperability layer around existing market systems, the design reduces adoption barriers and matches the incremental transformation logic often required in SME ecosystems.

This paper presents a design-oriented contribution developed prior to full operational deployment. Its primary value lies in the definition of a structured architectural blueprint and an evaluation logic for implementing interoperable traceability in wholesale fresh-produce logistics. The proposed system prioritizes zone-level localization, which is consistent with pallet-level handling processes typical of wholesale markets, rather than fine-grained indoor positioning. Sensing is selectively applied through reusable trackers assigned to prioritized flows, enabling the generation of auditable evidence packages where monitoring provides the greatest operational value.

The main contribution of the study is therefore its replicability. The design logic can be transferred to other wholesale markets by applying a consistent set of principles: defining the smallest accountable container, explicitly modelling aggregation relationships, aligning event capture with GS1 EPCIS 2.0 standards, generating portable integrity evidence, and governing data sharing through tiered visibility mechanisms. In this way, the study provides a transferable operational template for implementing interoperable traceability infrastructures within fresh-produce logistics networks.

7. Conclusion

This paper proposes a design-oriented traceability architecture for fresh-produce logistics in wholesale market environments. The contribution focuses on how heterogeneous operational signals—RFID readings, mobile handling events, sensor observations, and aggregation operations—can be systematically transformed into interoperable business events aligned with GS1 EPCIS 2.0. By structuring traceability around explicit aggregation relationships and portable evidence packages, the proposed model supports auditable product histories while enabling selective and controlled data sharing among supply chain actors.

Rather than introducing a single technology component, the study contributes a coherent architectural logic for implementing interoperable traceability infrastructures in complex logistics environments such as wholesale produce markets. The design integrates sensing, event capture, cloud-based normalization, and API-based interoperability in a way that remains compatible with existing enterprise systems while supporting transparent coordination across organizational boundaries.

The relevance of the proposal lies in its replicability. The same design principles—definition of accountable logistics units, explicit aggregation modelling, EPCIS-aligned event capture, and tiered data visibility—can be adopted by other wholesale markets or logistics hubs seeking to improve traceability without imposing centralized data ownership.

Future work will focus on operational deployment and empirical evaluation of the system's impact on logistics transparency, exception handling, and supply chain coordination.

References

- GS1 (2022a): EPCIS Standard, Release 2.0.1. GS1 AISBL.

- GS1 (2022b): Core Business Vocabulary (CBV) Standard, Release 2.0. GS1 AISBL.
- ICONET (2020): D2.2 PI Reference Architecture - Final Version. Horizon 2020 project deliverable.
- GS1 (2023): EPCIS and CBV Implementation Guideline, Release 2.0. GS1 AISBL.
- Maciariello F., Benelli F., Sangiuolo G., Lorenzi E., Caponio C. and C. Salvadori, "TrackOne: Smart Logistics for a Sustainable and Interoperable Agricultural Supply Chain in the Era of Digitization," 2025 International Conference on Software, Telecommunications and Computer Networks (SoftCOM), Split, Croatia, 2025, pp. 1-7, doi: 10.23919/SoftCOM66362.2025.11197431.
- Mainetti L., F. Mele, L. Patrono, F. Simone, M.L. Stefanizzi, R. Vergallo (2013): An RFID-Based Tracing and Tracking System for the Fresh Vegetables Supply Chain. International Journal of Antennas and Propagation, 2013, 531364. https://doi.org/10.1155/2013/531364
- Marino, F., Seitanidis, I., Dao, P., Bocchino, S., Castoldi, P., & Salvadori, C. (2019, July). IoT enabling PI: towards hyperconnected and interoperable smart containers. In 6th International Physical Internet Conference (pp. 349-362)
- Montreuil B. (2011): Toward a Physical Internet: Meeting the Global Logistics Sustainability Grand Challenge. Logistics Research, 3(2-3), 71-87. https://doi.org/10.1007/s12159-011-0045-x
- Montreuil B., E. Ballot, F. Fontane (2012): An Open Logistics Interconnection Model for the Physical Internet. IFAC Proceedings Volumes, 45(6), 327-332. https://doi.org/10.3182/20120523-3-RO-2023.00385
- Montreuil B., R.D. Meller, E. Ballot (2010): Towards a Physical Internet: The Impact on Logistics Facilities and Material Handling Systems Design and Innovation. Progress in Material Handling Research, 2010, 305-327.
- Mustafa M.F.M.S., N. Navaranjan, A. Demirovic (2024): Food Cold Chain Logistics and Management: A Review of Current Development and Emerging Trends. Journal of Agriculture and Food Research, 18, 101343. https://doi.org/10.1016/j.jafr.2024.101343
- Pal A., K. Kant (2017): A Food Transportation Framework for an Efficient and Worker-Friendly Fresh Food Physical Internet. Logistics, 1(2), 10. https://doi.org/10.3390/logistics1020010
- Pal A., K. Kant (2019): Internet of Perishable Logistics: Building Smart Fresh Food Supply Chain Networks. IEEE Access, 7, 17675-17695. https://doi.org/10.1109/ACCESS.2019.2894126
- Pan S., E. Ballot, G.Q. Huang, B. Montreuil (2017): Physical Internet and interconnected logistics services: research and applications. International Journal of Production Research, 55(9), 2603-2609. https://doi.org/10.1080/00207543.2017.1302620
- Razak G.M., L.C. Hendry, M. Stevenson (2023): Supply Chain Traceability: A Review of the Benefits and Its Relationship with Supply Chain Resilience. Production Planning and Control, 34(11), 1114-1134. https://doi.org/10.1080/09537287.2021.1983661
- Thakur M., E. Forås (2015): EPCIS-Based Online Temperature Monitoring and Traceability in a Cold Meat Chain. Computers and Electronics in Agriculture, 117, 22-30. https://doi.org/10.1016/j.compag.2015.07.006
- Treiblmaier H. (2019): Combining Blockchain Technology and the Physical Internet to Achieve Triple Bottom Line Sustainability: A Comprehensive Research Agenda for Modern Logistics and Supply Chain Management. Logistics, 3(1), 10. https://doi.org/10.3390/logistics3010010
- Treiblmaier H., K. Mirkovski, P.B. Lowry, Z.G. Zacharia (2020): The Physical Internet as a New Supply Chain Paradigm: A Systematic Literature Review and a Comprehensive Framework. International Journal of Logistics Management, 31(2), 239-287. https://doi.org/10.1108/IJLM-11-2018-0284
- UNECE (2023): Code of Good Practice: Reducing Food Loss and Ensuring Optimum Handling of Fresh Fruit and Vegetables Along the Value Chain. United Nations Economic Commission for Europe.
- Wang X., X. Li, Y. Liu, Z. Ying, C. Liang, X. Zhou, Y. Hong, Y. Luo (2020): Ethylene Sensor-Enabled Dynamic Monitoring and Multi-Strategies Control for Quality Management of Fruit Cold Chain Logistics. Sensors, 20(20), 5830. https://doi.org/10.3390/s20205830

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE
Tilburg University

Order dispatching for on-demand delivery with elevator waiting time

Haitao Liu[1], Mouna Bamoumen[2], Juan C. Pina-Pardo[3], Zhaoxia Guo[1], Jan Fransoo[2]
[1] Sichuan University, Chengdu, China
[2] Tilburg University, Tilburg, Netherlands
[3] Pontificia Universidad Católica de Valparaíso, Valparaíso, Chile
Corresponding author: m.bamoumen@tilburguniversity.edu

Abstract: *The surge in e-commerce has propelled on-demand delivery models promising fulfillment within an hour. This study addresses the Dynamic Pick-up and Delivery Problem inherent to these platforms, focusing on the critical, yet often overlooked, bottleneck of vertical logistics in high-density urban environments. While most dispatching models treat service times deterministically, we argue that elevator waiting and round-trip times in high-rise buildings introduce significant, non-stationary uncertainty that severely impacts delivery performance. To address this, we model the dynamic order dispatching process as a Markov Decision Process (MDP) aiming to minimize expected delivery time. Our model integrates three primary sources of uncertainty: the unknown total order volume, random order placement times, and stochastic elevator wait times. Due to the curse of dimensionality associated with the complex state space, which encompasses courier positions, order statuses, and building characteristics, among others, we employ an approximate approach. We rely on simulation-based policy optimization, applying Bayesian optimization to tune a parametric policy class. Benchmarking demonstrates our proposed method outperforms baseline policies and maintains robustness across various demand scenarios, offering an optimized dispatching solution that better accounts for the realities of the "last 50 feet."*

Keywords: *Last-mile Logistics, On-demand Order Dispatching, Dynamic Pick-up and Delivery, Stochastic Service Times, Simulation/Optimization methods, Physical Internet.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan: ☐ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories), ☐ Transportation Equipment, ☐ PI Networks, ☒ System of Logistics Networks, ☐ Vertical Supply Consolidation, ☐ Horizontal Supply Chain Alignment, ☒ Logistics/Commercial Data Platform, ☐ Access and Adoption, ☐ Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

1 Introduction

With the recent surge of e-commerce, last-mile delivery has gained substantial attention within the urban mobility and operations management literature. The operational challenge posed upon platforms to offer on-demand delivery services characterized by stringent fulfillment promises has left them operating on tight margins. In the scholarly literature, this operational challenge is formally classified as the Dynamic Pick-up and Delivery Problem (DPDP) (Pillac et al., 2013). Unlike classical static vehicle routing and dispatching problems, the DPDP requires a

platform to sequentially dispatch a transient fleet of couriers as stochastic orders arrive continuously throughout a given time horizon (Psaraftis et al., 2016).

Previously, solutions to the DPDP and its related variants have focused mostly on optimizing horizontal travel (Ulmer et al., 2021). Consequently, adapted algorithms and reinforcement learning (RL) frameworks have been developed to minimize drive times, balance rider workloads, and predict dynamic traffic conditions (Jahanshahi et al., 2022). However, in dense urban environments, the delivery realities extend far beyond the curb. In the classical setting of on-demand delivery, couriers are responsible for bringing items all the way to the customer's door. Consequently, couriers must navigate complex residential and office buildings, park their vehicles, walk to lobbies, and undertake vertical transportation systems.

This vertical phase of the delivery process introduces a severe yet critically overlooked bottleneck: elevator waiting and transit time. Evidence from both journalistic investigations and preliminary literature highlights the severity of this issue. A recent article from The New York Times (Wee, 2025) detailed the emergence of an informal economy in a major Chinese city (Shenzhen) where couriers, facing high penalties for delayed deliveries, pay individuals already inside high-rise buildings to complete the final vertical leg of the journey. This improvisation is a direct response to peak-hour elevator waits that frequently reach 20 to 30 minutes, often exceeding the total road travel time for the order.

Despite this operational reality, most routing and dispatching models continue to treat in-building time as a negligible or deterministic constant. Our preliminary analysis of simulated high-rise delivery environments reveals a highly variable elevator time distribution with a long right tail; moderate delays of 8 to 12 minutes push deliveries dangerously close to standard 20-minute fulfillment promises, while heavy delays almost certainly result in service failures. As such, the elevator, not the road, is often the dominant bottleneck in high-rise delivery.

From a Physical Internet perspective, vertical transportation systems such as elevators can be viewed as shared logistics infrastructure, where congestion and stochastic waiting times arise from the interaction of tenants, residents, and multiple concurrent delivery flows competing for limited elevator capacity. While our current model strictly focuses on one platform operating alone, it mathematically shows how much time and efficiency is wasted by this fragmented approach. By quantifying the delays caused when couriers fight for elevators, our results show why the urban delivery industry must evolve toward shared PI systems, such as ground-floor hubs or meal parcel lockers, to remain viable.

Beyond maintaining the economic viability of on-demand platforms and dispatching models must explicitly account for elevator waiting-time uncertainty to protect the welfare of gig-economy couriers as well. Therefore, this study aims to address this critical gap by asking: How does elevator waiting time affect overall delivery performance, and what dispatching policy best handles this specific vertical uncertainty? To answer these questions, we make the following contributions:

- We explicitly model stochastic elevator delays in the Markov Decision Process (MDP) formulation to resolve the "last 50 feet" vertical bottleneck in high-rise delivery.
- We deploy a simulation-driven Bayesian framework to efficiently learn optimal, dynamic dispatching weights.
- We introduce a transparent dispatching model utilizing novel spatial and vertical features, such as aggregate floor load.
- We offer quantifiable real-world impact by calibrating our approach with empirical data

The remainder of this paper is organized as follows. Section 2 reviews the relevant literature concerning dynamic last-mile delivery and vertical congestion. Section 3 formally defines the stochastic delivery environment and formulates the dispatching challenge as a MDP. Section 4 details our proposed solution methodology. Section 5 presents comprehensive numerical experiments to demonstrate the performance, robustness, and stakeholder impact of our approach. Finally, Section 6 concludes the study with avenues for future research.

2 Literature review

In this section, we present the streams of literature related to our work: the dispatching schemes last-mile delivery (2.1), the methodologies governing dynamic dispatching (2.2) and the vertical dynamics of dense urban delivery (2.3).

2.1 Dispatching in last-mile delivery

Although dynamic resource allocation is conceptually unified, the literature reveals certain structural limitations that modify dispatching logic across ride-hailing, Same-Day Delivery (SDD), and On-demand delivery.

In the ride-hailing literature, dynamic matching is predominantly treated as a single-point assignment problem. The primary objective is to maximize spatial liquidity and immediate pairing, which requires algorithms capable of re-evaluating thousands of trip-vehicle combinations in real time (Alonso-Mora et al., 2017; Bimpikis et al., 2019). Conversely, SDD operates over a much wider temporal scale. This relaxed constraint allows platforms to utilize "dispatch waves," effectively aggregating orders at a central depot and optimizing periodic departures to achieve significant route consolidation (Klapp et al., 2018; Voccia et al., 2019).

On-demand delivery occupies a more constrained middle ground. It combines the complex two-point tasks of SDD (pick-up and drop-off) with the real-time urgency of ride-hailing. Furthermore, on-demand delivery relies heavily on third-party pick-up locations, such as restaurants or drugstores, adding an external layer of merchant preparation time (Ulmer et al., 2021). However, in their study of the dispatching problem, scholars neglect the added time generated by the drop-off task in highly dense urban areas.

2.2 Myopic to anticipatory methods for dynamic dispatching

The Restaurant Meal Delivery Problem (RMDP) represents a highly stochastic On-demand delivery problem and customer-centric subset of the dynamic pickup and delivery problem. A few decades ago, fleet control in dynamic pickup and delivery relied heavily on reactive re-optimization strategies and insertion heuristics to navigate the high-dimensional action space. For instance, Reyes et al. (2018) and Yildiz and Savelsbergh (2019) utilized these reactive methods to minimize delivery times.

To better handle the deep uncertainty inherent to meal delivery, research transitioned toward anticipatory methods, including waiting strategies (Steever et al., 2019; Ulmer et al., 2021) and simulation-based look-ahead policies. More recently, the operations management field has explored neural network-based RL for dynamic fleet and demand control. For example, Hildebrandt et al. (2026) proposed an attention-based Q-learning approach to jointly optimize fleet assignments and platform display configurations to anticipate future demand. However, applying RL to large-scale urban routing remains computationally challenging, and existing approaches often struggle to capture complex fleet-wide interactions holistically (Jahanshahi et al., 2022). Furthermore, even state-of-the-art anticipatory algorithms strictly evaluate spatial

state-action pairs based on horizontal street graphs, while we explicitly incorporate the vertical dimension induced by high-rise buildings.

2.3 The vertical dynamics of dense urban delivery

Despite the high sophistication of anticipatory and approximate methods, a critical evaluation of the existing literature reveals a systemic blind spot: the persistent failure to model vertical congestion. Most top-tier dispatching models effectively operate under a "flat city" assumption, treating service time, defined as the interval between reaching the destination building and handing the order to the customer, as a deterministic constant.

In the reality of door-to-door delivery within dense cities, couriers must navigate complex residential and office buildings. While travel times between horizontal coordinates are increasingly predictable, the vertical phase, formally referred to as the "last 50 feet", remains an unmodeled source of extreme friction. Empirical studies demonstrate that urban delivery vehicles spend a dominant portion of their total daily operating time parked while drivers navigate this final vertical and walking segment (Dalla Chiara et al., 2021).

Specifically, the elevator, not the road, is often the critical bottleneck in high-rise delivery. Research into vertical transportation dynamics highlights that as passenger demand increases, elevator systems experience lockstep traffic, creating highly variable, non-linear wait times that compound horizontal delivery delays (Feng and Redner, 2021). Building on this, Paredes-Belmar et al. (2024) mathematically demonstrated that the sequence and duration of horizontal walking, unloading, and waiting for elevators inside buildings vary drastically based on daily building traffic intensity. Current dispatching literature consistently ignores this critical source of uncertainty. When models ignore this variability, it creates a cascading effect where a single unmodeled elevator delay at a high-rise destination compromises the platform's ability to meet subsequent delivery deadlines across the entire fleet. To advance collaborative and intelligent logistics networks, methodologies must explicitly integrate stochastic vertical transit into dispatching policies.

3 Problem formulation

Our problem is a stochastic and dynamic decision-making problem with a finite-time horizon, formulated as a MDP.

3.1 Problem description

We consider a fixed service region comprising a set of pickup points $\mathcal{O} = \{O_1, \dots, O_o\}$ and a set of residential or office buildings $\mathcal{D} = \{D_1, \dots, D_d\}$. The service network is denoted by $\mathcal{N} = \mathcal{O} \cup \mathcal{D}$, with deterministic travel times between locations. The platform employs M homogeneous riders to fulfill delivery tasks dynamically announced within a continuous time horizon $\mathcal{T} = [0, t_{max}]$. Because riders often make independent routing decisions based on past experience and real-time conditions, we do not force strict repositioning; instead, we give riders the autonomy to determine their traveling routes after orders are assigned.

Orders arrive according to an unknown, non-stationary stochastic process. Each customer order k is represented by a tuple $C_k = (t_k, o_k, d_k, b_k)$, where t_k is the placement time, o_k is the pickup location, d_k is the destination building, and b_k signifies the delivery floor. We assume door-to-door delivery, meaning each rider must park, wait for, and ride an elevator. The waiting and round-trip time to take elevators, $E(t)$, follows a non-stationary stochastic process with high variance. Let λ_k denote the completion time when order C_k is received by the customer. The objective is to minimize the expected delivery time:

$$\min \mathbb{E}\left[\sum_{k=1}^{K}(\lambda_k - t_k)\right] \quad (1)$$

3.2 Markov Decision Process Model

We model the dispatching problem under stochastic service times as follows:

- *State*: The platform makes a dispatching decision each time a new order is placed. For customer k, the state is $S_k = (t_k, \emptyset_k, p_k, z_k, r_k)$. Here, t_k is the request time, $\emptyset_k$ contains information on unfulfilled orders, p_k is a binary vector indicating if orders have been picked up, z_k captures previous assignment decisions, and r_k represents the current locations of all riders.
- *Decision*: A decision a_k determines the assignment of pending orders to riders, updating the binary variables z_{imk}^{a}. Assignments are permanent once made.
- *Exogenous Information*: The system is subjected to exogenous randomness $W_{k+1}(S_k, a_k)$ representing the total order volume K, uncertain elevator times $E(t)$, and future placement times t_k.
- *Transition*: The transition function $S_{k+1} = \mathcal{T}(S_k, a_k, W_{k+1})$ updates the state upon the arrival of order $k+1$, updating rider locations, order fulfillment statuses, and transferring pending assignments.
- *Reward*: The reward function $R_k(S_k, a_k)$ calculates the waiting time. The optimal policy minimizes the expected waiting time via the Bellman equation:

$$V(S_k) = \min_{a_k}(R_k(S_k, a_k) + \mathbb{E}[V_{k+1}(S_{k+1})|S_k]) \quad (2)$$

4 Solution method

Due to the curse of dimensionality, exactly solving the Bellman equation is intractable. Reinforcement learning and approximated dynamic programming are commonly adopted to learn optimal policies for many complex decision-making problems, which rely on neural networks to approximate the value functions. Nevertheless, it is not easy to interpret these complex policies. Thus, we develop a rule-based order dispatching policy and optimize the policy from a Bayesian optimization perspective.

4.1 Parametric order dispatching policy

Rule-based dispatching policies are widely adopted in on-demand delivery systems due to their low computational complexity and easy practical implementation. Traditional single-rule-based dispatching policies exhibit inherent limitations, because each individual strategy only optimizes from a single dimension and fails to adapt to complex dynamic scenarios in on-demand delivery. To address this issue, we integrate multiple classic rule-based dispatching policies and absorb the respective strengths of different strategies. By comprehensively incorporating the decision logics of geographic proximity, workload balancing, and service region, the integrated policy avoids the one-sidedness of standalone rules, thereby achieving more reasonable order-rider matching and maintaining stable system performance under dynamic order arrivals.

Specifically, we consider six features to define the order dispatching policy: $\pi(a|s;\theta) = \pi\left(\theta^{\top}\phi(s,a)\right)$. These features $\phi(s,a)$ include:

- *Workload:* It represents the number of uncompleted orders per rider.
- *Pickup Distance:* It measures the distance between the rider's current location and the new order's pickup point.

- *Pickup Proximity:* It computes the minimum distance between pickup points of a rider's uncompleted orders and the new order's pickup point.
- *Delivery Proximity:* It computes the minimum distance between delivery points of a rider's uncompleted orders and the new order's delivery point.
- *Visiting Area:* The geographical footprint of the rider's current region.
- *Floor Load:* The sum of floors of uncompleted orders, explicitly capturing vertical workload.

Then, we use a linear combination of the six features to represent the dispatching policy. Equivalently, the optimization of the policy is transformed into finding the optimal value of θ. Meanwhile, as the simulation of the order dispatching system is time-consuming, we propose to tune all weight parameters via Bayesian optimization.

4.2 Delivery time approximation via Bayesian Optimization

Bayesian optimization (BO) is a derivative-free global optimization algorithm tailored for black-box and computationally expensive objective functions. It constructs a surrogate probabilistic model (usually a Gaussian process) to approximate the unknown objective function and employs an acquisition function to intelligently balance exploration (searching unevaluated parameter regions) and exploitation (refining current optimal solutions). Compared with conventional grid search and random search, BO can efficiently explore continuous parameter spaces with far fewer evaluation iterations, effectively avoids the inefficiency and subjective bias of manual parameter tuning, and quickly converges to the globally optimal parameter combination. It is particularly suitable for optimizing the weight parameters of integrated rule-based dispatching policies in complex on-demand delivery scenarios.

Evaluating the expected delivery time $f(\theta)$ is computationally expensive. We model the unknown function as a Gaussian process (GP) regression:

$$f(\theta_{1:n}) \sim \mathcal{N}\left(\mu_0(\theta_{1:n}), \Sigma(\theta_{1:n}, \theta_{1:n})\right) \tag{3}$$

Assuming noise on delivery time is normally distributed, we calculate the posterior distribution and utilize an acquisition function (Gradient-based Expected Improvement) to iteratively determine the next sampling point θ_n. This allows the platform to intelligently search the parameter space to find the optimal weights for the dispatching features.

5 Numerical experiments

5.1 Numerical setting and parameter estimation

We consider a 3 km × 3 km service region containing 30 pickup points, 60 delivery buildings (up to 32 floors), and 11 riders operating over a 12-hour horizon. To calibrate the stochastic order arrival process, we utilize an operational-level on-demand meal delivery dataset from Hangzhou, China, provided by Mao et al. (2022). Empirical data reveals a right-skewed distribution for both horizontal distance (mean 1.66 km) and delivery time (mean 54.09 mins). Crucially, the order arrival process is highly non-stationary, exhibiting a binormal distribution where lunch demand is nearly twice as high as dinner demand.

To accurately simulate the non-stationary temporal dynamics of customer demand, we rely on the empirical order placement times to estimate the expected rate of occurrences within discrete 15-minute intervals. This reveals that the order arrival process is highly non-stationary, exhibiting distinct surges in demand where the mean delivery volume at lunch is almost twice as large as that at dinner. To mathematically simulate this reality, we model the frequency of

orders over the day as a bimodal distribution. Specifically, the number of customers requesting orders during the lunch and dinner peaks are treated as independent, normally distributed variables.

5.2 Performance comparison

We benchmarked our proposed weighted policy against 8 baseline dispatching rules (e.g., greedy workload, pickup distance, floor load) across 500 replications. Results demonstrate that our optimized weighted policy achieves the shortest mean delivery time and the smallest standard deviation (Figure 1). Interestingly, while distance-based rules perform poorly in isolation, features accounting for load balance (Workload and Floor Load) and service area showed strong individual performance, highlighting the necessity of managing vertical congestion. This generates interesting results when it comes to the weight features in the weighed policy : from the right graph in Figure 1, it appears that batching pickups features are driving the optimization performance (with weights of 0.893 and 0.738 for Pickup Proximity and Pickup resp.) ; however, this is only possible because other features like load balance ones play a crucial role in containing the variance introduced by elevator times. The Visiting Area also received a high weight, as it is more beneficial to restrict horizontal travel with unpredictable elevator times. While it might seem counterintuitive that Delivery Proximity ranked last, it actually makes sense that queues for elevator in two buildings matter more than the horizontal proximity between them.

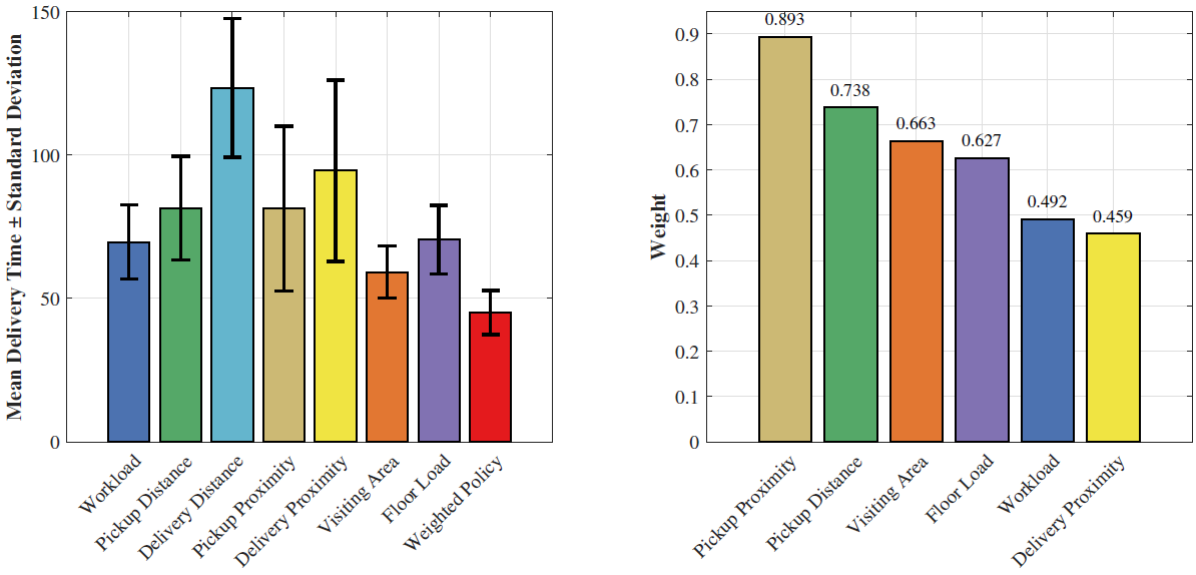


Figure 1: Delivery times of dispatching policies and optimized weights in the weighted policy

5.3 Policy robustness test

To test out-of-sample robustness, we evaluated the policy across nine problem instances (S1-S9), varying parameters such as fleet size, base arrival rate, building height and elevator noise variance.

The proposed weighted policy consistently outperformed baseline heuristics across all scenarios without requiring model retraining. Under default operating conditions, our policy reduced mean delivery times by over 53% compared to the standard Workload policy and nearly 78% compared to the Delivery Distance policy. Furthermore, the model proved highly resilient under extreme stress-test conditions. For example, when the base order arrival rate was doubled to simulate intense peak periods, traditional proximity-based rules suffered cascading

delays. In contrast, the weighted policy maintained stable fulfillment, accelerating delivery times by over 62% relative to the Workload baseline and by roughly 80% compared to the weakest proximity heuristics.

5.4 Impact on stakeholders

Our policy yields significant and quantifiable benefits across the urban logistics ecosystem, demonstrating the value of integrating vertical congestion into decision-making:

- *Society (Parking Time)*: By optimizing dispatching around elevator wait times, couriers spend less time occupying public parking infrastructure outside buildings. The weighted policy achieved up to a 33% reduction in mean parking time compared to traditional baseline policies, directly mitigating local traffic congestion and illegal parking risks.
- *Customers (Waiting and Delay Times):* The optimized policy drastically reduces customer friction. In the default scenario, the weighted policy slashed average customer delay times by over 86% compared to the standard Workload baseline. Furthermore, maximum waiting times for customers were reduced by more than 66%, ensuring a much more reliable service experience.
- *Riders (Work Intensity)*: Our policy naturally balances work intensity by avoiding clustering riders in high-rise bottleneck zones. Under the weighted policy, the average time a rider spent fulfilling a single order was reduced by over 54% compared to the Workload baseline, effectively halving their per-order time commitment.
- *Platforms (Service Level)*: Measuring the platform's service level by the probability of an on-time delivery (under 60 minutes), the weighted policy maintained the highest reliability. It improved on-time delivery rates by approximately 14 to 16 percentage points over standard baselines. Especially, during high-demand peak scenarios, the weighted policy maintained an 81% on-time fulfillment rate, protecting platform margins and reputation while proximity-based heuristics dropped to nearly 60%.

6 Conclusion

This paper addresses the Dynamic Pick-up and Delivery Problem in on-demand services by integrating a critical, yet so far neglected, source of uncertainty: stochastic elevator waiting times. By formulating the problem as a MDP and utilizing Bayesian optimization to tune a parametric policy, we provide a robust dispatching framework that outperforms standard rule-based policies. Our results confirm that ignoring the 'last 50 feet' of vertical logistics creates hidden bottlenecks that significantly prolong delivery times and disrupt urban mobility. By acknowledging and optimizing for these vertical constraints, this study contributes directly to the deployment of more resilient, hyperconnected urban logistics networks in alignment with the Physical Internet paradigm.

While our approach provides a strong basis for mitigating vertical congestion, there are a number of interesting directions for further study in this ever-developing field of operations management:

- *Customer incentives and ground-floor infrastructure*: Future models should investigate the operational trade-offs of utilizing ground-floor delivery lockers or offering dynamic financial incentives for customers in high-rises to retrieve their own meals from the lobby. Evaluating how these interventions alter platform profitability and courier efficiency could reshape urban infrastructure planning.
- *Multi-modal and collaborative vertical logistic*s: As automation advances, exploring collaborative delivery models, such as utilizing aerial drones for the horizontal travel

leg and human couriers (or indoor robots) exclusively for the vertical leg, presents a compelling optimization challenge. Research could focus on synchronizing these multi-modal interactions to avoid parking and elevator bottlenecks.

References

- Alonso-Mora, J., Samaranayake, S., Wallar, A., Frazzoli, E., & Rus, D. (2017). On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. *Proceedings of the National Academy of Sciences*, 114(3), 462-467.
- Bimpikis, O., Candogan, O., & Saban, D. (2019). Spatial pricing in ride-sharing networks. *Operations Research*, 67(3), 744-769.
- Chen, X., Ulmer, M. W., & Thomas, B. W. (2022). Deep Q-learning for same-day delivery with vehicles and drones. *European Journal of Operational Research*, 298(3), 939-952.
- Dalla Chiara, G., Krutein, K. F., Ranjbari, A., & Goodchild, A. (2021). Understanding urban commercial vehicle driver behaviors and decision making. *Transportation Research Record*, 2675(9), 608-619.
- Feng, Z., & Redner, S. (2021). When will an elevator arrive?. *Journal of Statistical Mechanics: Theory and Experiment*, 2021(4), 043403.
- Hildebrandt, F. D., Lesjak, Ž., Strauss, A., & Ulmer, M. W. (2026). Integrated fleet and demand control for on-demand meal delivery platforms. *Management Science*, 72(2), 932-954.
- Jahanshahi, H., Bozanta, A., Cevik, M., Kavuk, E. M., Tosun, A., Sonuc, S. B., Kosucu, B., & Başar, A. (2022). A deep reinforcement learning approach for the meal delivery problem. *Knowledge-Based Systems*, 243, 108489.
- Klapp, M. A., Erera, A. L., & Toriello, A. (2018). The one-dimensional dynamic dispatch waves problem. *Transportation Science*, 52(2), 402-415.
- Mao, W., Ming, L., Rong, Y., Tang, C.S., & Zheng, H. (2022) On-demand meal delivery platforms: Operational level data and research opportunities. *Manufacturing & Service Operations Management* 24(5):2535–2542.
- Paredes-Belmar, G., Latorre-Núñez, G., & Bronfman, A. (2024). Modeling and solving the time-dependent in-building delivery problem in last-mile logistics. *IEEE Access*, 12, 11276-11293.
- Pillac, V., Gendreau, M., Guéret, C., & Medaglia, A. L. (2013). A review of dynamic vehicle routing problems. *European Journal of Operational Research*, 225(1), 1-11.
- Psaraftis, H. N., Wen, M., & Kontovas, C. A. (2016). Dynamic vehicle routing problems: Three decades and counting. *Networks*, 67(1), 3-31.
- Reyes, D., Erera, A., Savelsbergh, M., Sahasrabudhe, S., & O'Neil, R. (2018). The meal delivery routing problem. *Technical report, Georgia Institute of Technology*.
- Steever, Z., Karwan, M., & Murray, C. (2019). Dynamic courier routing for a food delivery service. *Computers & Operations Research*, 107, 173-188.
- Ulmer, M. W., Thomas, B. W., Campbell, A. M., & Woyak, N. (2021). The restaurant meal delivery problem: Dynamic pickup and delivery with deadlines and random ready times. *Transportation Science*, 55(1), 75-100.
- Voccia, S. A., Campbell, A. M., & Thomas, B. W. (2019). The same-day delivery problem for online purchases. *Transportation Science*, 53(1), 167-184.
- Wee, S.-L. (2025, September 29). The new delivery guys for China's delivery guys. *The New York Times*. https://www.nytimes.com/2025/09/29/world/asia/china-delivery-shenzhen.html
- Yildiz, B., & Savelsbergh, M. (2019). Provably high-quality solutions for the meal delivery routing problem. *Transportation Science*, 53(5), 1372-1388.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

The TREATMENT project: a Physical-Internet-enabled reference architecture for end-to-end visibility in direct-to-patient clinical trial logistics

Franco Maciariello[1], Eddi Lorenzi[1], Cristiana Caponio[1], Giuseppe Tortora[2], Claudio Salvadori[1]
[1] New Generation Sensors s.r.l., Monte San Savino (AR), Italy
[2] ABzero s.r.l., Pisa, Italy
Corresponding author: claudio.salvadori@ngs-sensors.it

Abstract: *Clinical trials are increasingly adopting decentralized and direct-to-patient delivery models, but these models intensify the need for auditable chain-of-custody, condition control, and coordinated execution across heterogeneous logistics actors. This paper presents TREATMENT as a Physical-Internet-enabled reference architecture for regulated direct-to-patient clinical trial logistics. The contribution does not lie in standards adoption alone, but in the integration of connected logistics units, a neutral multi-tenant cooperation platform, a single-trace interoperability mechanism combining GS1 EPCIS 2.0 events and GS1 EDI documents, and an explicit milestone-driven rule-based logic for road–drone synchronization. The evaluation is structured around trace completeness, chain-of-custody integrity, document-event correlation, telemetry timeliness, condition monitoring, route integrity, and intermodal synchronization. Results show complete capture of planned custody milestones, successful correlation of shipment documents and physical events across heterogeneous providers, near-real-time telemetry availability, correct alert generation for conservation and route anomalies, and event-governed release of the downstream drone leg from validated upstream milestones. The paper also discusses transferability conditions, architectural limits, and the broader relevance of the approach for regulated multimodal healthcare logistics.*

Keywords: *Physical Internet; clinical trial logistics; direct-to-patient; road-drone cooperation; GS1 EPCIS; interoperability; end-to-end visibility; connected logistics unit*

Physical Internet (PI) Roadmap Fitness: [X] Transportation Equipment [X] Logistics/Commercial Data Platform [X] Governance [X] PI Networks [X] Access and Adoption

Targeted Delivery Modes: [X] Paper [ ] Poster [ ] Flash Video [X] In-Person presentation

1. Introduction

Clinical trial supply chains are designed around stringent requirements for product integrity, subject safety, protocol compliance, and audit readiness. These requirements become harder to satisfy when investigational medicinal products move beyond the clinical site and are delivered directly to participants, as happens in decentralized and hybrid trial models. Recent reviews and regulatory guidance highlight both the promise and the operational fragility of decentralized clinical trials: remote participation can reduce participant burden and improve access, but it also creates new challenges for investigational product supply, temperature control, role allocation, and evidence continuity (Rogers et al., 2022; FDA, 2024; EMA, 2022).

The logistics problem is therefore not limited to transport efficiency. In regulated direct-to-patient operations, the central question is whether multiple actors and modes can produce a continuous, trustworthy, and machine-readable record linking what was shipped, where it moved, under which conditions it was handled, and when it was received. GS1 has long provided identification and messaging standards (based on GS1 EDI) for clinical trials (GS1,

2022b), but in practice, direct-to-patient clinical trial logistics still suffers from fragmented data ownership, provider-specific interfaces, and limited correlation between shipment documents and physical events. This paper addresses that gap through the TREATMENT[1] project. The research contribution is not presented as a generic "digitalization" case, nor as a simple implementation of existing standards. Instead, it proposes and critically discusses a Physical-Internet-enabled reference architecture that combines: (i) connected logistics units at item and aggregation level; (ii) a neutral multi-tenant cooperation platform for cross-actor visibility and governance; (iii) a single-trace mechanism that correlates GS1 EPCIS 2.0 (GS1, 2022a) events and GS1 EDI documents, going beyond the GS1 white-paper in (GS1, 2022b); and (iv) an explicit rule-based synchronization logic for activating the drone leg after validated road-leg milestones. The empirical scope of the paper is deliberately bounded: it includes device- and platform-level validation and two integrated demonstrations conducted in July and October 2025, while larger-scale and comparative benefits remain part of the research agenda rather than claimed outcomes.

2. Related work and research gap

Physical Internet research has established the value of modular logistics units, open networks, interoperable interfaces, and neutral coordination mechanisms for improving logistics performance at ecosystem level (Montreuil, 2011; Montreuil et al., 2013; Ballot et al., 2014; Crainic and Montreuil, 2016; Pan et al., 2017). In this literature, visibility is not just a tracking feature: it is the informational infrastructure that makes shared assets, pull-based orchestration, and synchromodal reconfiguration feasible. In parallel, the decentralized clinical trial literature has shown that investigational product distribution to participants is feasible but operationally demanding. The systematic review by Rogers et al. (2022) showed that decentralized methods are expanding across trial activities, while recent operational work from the Trials@Home RADIAL proof-of-concept trial reported that direct-to-participant shipment is feasible only when responsibilities, temperature control, participant availability, and contingencies are explicitly managed (Heath et al., 2025). Regulatory bodies have converged on a similar point: decentralized trial elements can be used, but sponsor oversight, investigator responsibility, data integrity, and participant protection must remain demonstrable (FDA, 2024; EMA, 2022).On the transport side, healthcare drone studies document growing technical maturity and clear potential for reducing response times in selected medical logistics scenarios, especially in constrained or remote geographies (Lin et al., 2018; Stierlin et al., 2024). However, most road-drone or drone-healthcare studies focus on routing, feasibility, safety, or service economics. Much less attention has been devoted to the regulated traceability problem that arises when an investigational product changes custody across organizations and modes while remaining subject to protocol-specific conservation, documentation, and audit requirements.

The research gap addressed here is therefore specific: existing work does not adequately explain how a regulated direct-to-patient clinical trial flow can be orchestrated as a Physical Internet service while preserving trace completeness, chain-of-custody integrity, selective data visibility, and auditable handover between road and aerial legs. TREATMENT contributes a reference architecture and an evaluation framing for exactly this problem.

3. Research contribution and system architecture

This paper makes the following scientific contributions: (i) it formalizes TREATMENT as a Physical-Internet-enabled reference architecture for regulated direct-to-patient clinical trial

[1] Treatment Project, Funded by the European Union – Next Generation EU, under the Research Program CN00000023 – NRRP Mission 4 – Component 2 – Investment Line 1.4 Sustainable Mobility Center CNMS - CUP E63C22000930007, CODSOG_003580

logistics; (ii) it defines a single-trace interoperability mechanism that correlates GS1 EPCIS 2.0 events and GS1 EDI documents around the same shipment objects and custody milestones; (iii) it makes the road–drone decision logic explicit as a rule-based, milestone-driven orchestration mechanism rather than leaving last-mile activation implicit; (iv) it proposes an evaluation framing for auditable intermodal coordination under regulated conditions, covering trace completeness, chain-of-custody integrity, document-event correlation, telemetry timeliness, exception visibility, and intermodal synchronization.

The architecture, depicted in Figure 1, is centred on two connected logistics units and one neutral cooperation layer. At item level, the smart box (a tertiary packaging containing the medicines) acts as a connected logistics unit attached to the investigational product package. It emits item-level events related to location, conservation conditions, and proof-of-delivery evidence. At aggregation level, the smart capsule acts as a reusable connected logistics unit designed for modal handover and last-mile mission execution. It aggregates one or more item-level units, exposes additional mission and environmental telemetry, and supports coupling with the drone leg.



Figure 1: Proposed TREATMENT reference architecture

Above these assets sits a neutral multi-tenant cooperation platform. The platform ingests operational data from heterogeneous providers, validates messages, normalizes them into a canonical model, enforces policy-based visibility, and stores the resulting evidence in an auditable repository. The platform does not replace provider-specific execution systems; instead, it separates coordination from execution, which is consistent with Physical Internet thinking about open networks and neutral interoperability layers.

The single-trace mechanism is achieved through shared identifiers and complementary message semantics. GS1 identifiers provide the stable references needed to connect item, logistic unit, location, and shipment semantics. GS1 EPCIS 2.0 captures the event history of association, aggregation, transport, receipt, and exception states. GS1 EDI captures business commitments and shipment documents such as despatch and receipt notices. The innovation is in their correlation: rather than treating documents and events as parallel but disconnected channels, the architecture binds them to the same shipment objects and custody milestones.

3.1. PI mapping of the architecture

Table 1 makes the mapping to core Physical Internet principles explicit. The aim is not to force an abstract PI vocabulary onto the case, but to show where the architecture genuinely instantiates PI concepts and where it remains a domain-specific adaptation for regulated healthcare logistics.

Table 1: Explicit mapping between TREATMENT and core Physical Internet principles

PI principle	TREATMENT implementation	Research relevance
Modular connected units	Smart box at item level; smart capsule as reusable aggregation unit for last mile and mod handover.	Moves the PI container idea into regulated healthcare and preserves modular handover semantics.
Open interoperable networks	GS1 identifiers, EPCIS 2.0 events, and GS1 EDI documents across heterogeneous providers	Shows interoperability across organizations rather than within a single proprietary platform.
Neutral data platform	Multi-tenant cooperation platform with canonical model, validation, routing, and audit trail.	Separates coordination from provider-specific execution systems.
Pull-based orchestration	Last-mile activation triggered by validated receipt, patient slot, and exception status.	Replaces purely push release with event-driven service activation.
Synchromodal coordination	Road-to-drone handover documented through aggregation, association, despatch, receipt, and deconsolidation events.	Extends PI thinking to controlled intermodal handover in clinical trials.
Selective trust and governance	Policy-based visibility, role-aware access, and transaction-level document-event correlation.	Addresses the governance layer often missing in conceptual PI deployments.

3.2. Rule-based road-drone synchronization

The current implementation uses rule-based orchestration rather than optimization. This design choice is deliberate. In regulated clinical settings, accountability, explainability, and auditability are often more critical than the global efficiency gains that might be achieved by an opaquer dispatcher. The synchronization mechanism therefore relies on milestone-driven heuristics anchored in observable events and auditable state transitions.

The drone leg can be activated only when five conditions are simultaneously satisfied: (1) a validated receipt event confirms that the shipment is physically available at the handover node; (2) the patient-selected delivery window is still valid; (3) the destination passes address validation and geofencing checks; (4) no unresolved conservation or integrity alert is open; and (5) an eligible service provider confirms execution within the required timeout and operational envelope. If any of these conditions is not satisfied, the request is not silently forced through the network. Instead, it is rescheduled, corrected manually, placed on hold for investigation, or reallocated to another provider or to a ground-based alternative. The main point is that the drone leg is not released by a static plan, but by the verified operational state of the network.

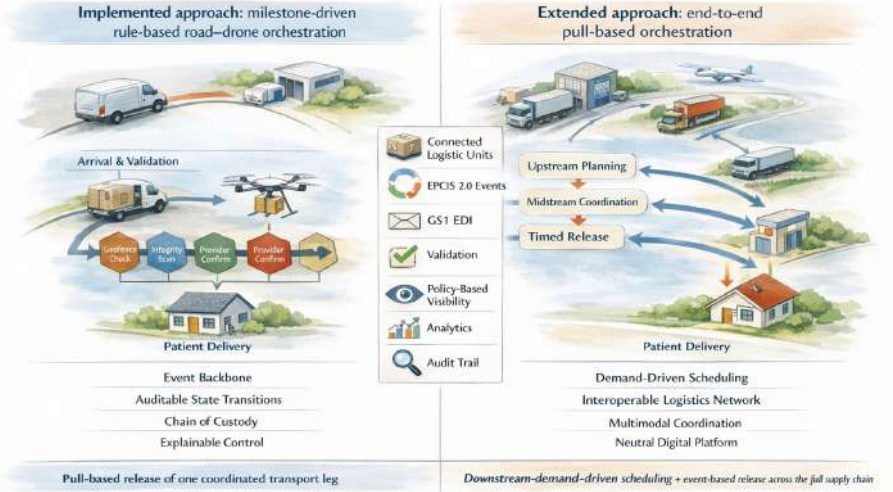


Figure 2: Implemented approach vs extended end-to-end pull approach

This logic is academically relevant because it shows how road–drone cooperation can be expressed as pull-based orchestration rather than as a static pre-planned itinerary. The triggering signal is not a planner's prior assumption that the shipment should be available at a certain time, but a validated milestone generated by the event backbone. In this sense, the drone leg is released by the verified state of the network rather than by the inertia of an upstream schedule. Its broader end-to-end generalization is discussed in Section 6.

4. Evaluation design

The evaluation is organized as a design-oriented case study with three empirical layers. The first layer covers component and interface level validations for connected logistics units and the neutral cooperation platform. The second layer, covered in the July 2025 field pilot, which validated the hospital-to-patient drone-enabled last-mile flow with two smart boxes, one smart capsule, one drone leg, and two heterogeneous digital providers. The third layer, covered in the end-of-project October 2025 demonstration scenario, which extends the scope upstream by including the road leg from depot to hospital and its synchronization with the subsequent direct-to-patient drone segment. In that scenario, two smart boxes are managed across two destinations, and the hospital receipt milestone triggers the pull-based activation of the downstream drone service through standardized EDI/EPCIS interactions.

In this study, the architecture is evaluated along seven dimensions: trace completeness, chain-of-custody integrity, document-event correlation, telemetry timeliness, condition monitoring, route integrity, and intermodal synchronization. These dimensions were selected because they are directly relevant to regulated direct-to-patient clinical trial logistics and because they can be supported by the evidence currently documented without overstating comparative performance claims, as shown in Table 2.

Table 2: Evaluated empirical dimensions and observed evidence

Dimension	Operationalization	Observed evidence
Trace completeness	Presence of all planned milestone types in machine-readable form across road, hospital handover, drone, and delivery phases.	The documented pilots cover milestone capture from logistic-unit creation and road-leg tracking to hospital receipt, shipment request, provider notification/confirmation, drone delivery, and proof of delivery.
Chain-of-custody integrity	Unambiguous link between packaging, aggregation, transport, receipt, deconsolidation, and delivery evidence.	Validated through EPCIS aggregation and association events, receipt messages, and delivery/deconsolidation evidence across the July and October scenarios.
Document-event correlation	Business documents linked to the same shipment objects as field events.	GS1 EDI despatch, receipt, shipment request, shipping notification, and shipping confirmation messages are correlated to shipment objects through EPCIS transaction events across heterogeneous providers.
Telemetry timeliness	Ability of smart capsule / drone telemetry to reach the cooperation platform at the required refresh rate.	The project KPI for data transfer at least once per minute is reported as achieved; the smart-capsule posting interval is described as exceeding the minimum requirement.
Condition monitoring	Ability to detect and report conservation anomalies through pre-alarm and alarm logic.	Positive tests are reported for threshold-based alert generation related to environmental conditions.

Route integrity	Ability to detect and report deviations from the expected path or geofence.	Positive tests are reported for pre-alarm and alarm generation under path-deviation or geofencing anomalies; geofencing checks are also part of operational validation in the transport flow.
Intermodal synchronization	Ability to release the drone leg based on validated upstream milestones rather than fixed pre-planning.	Road-leg events and hospital receipt are captured first; the validated receipt then triggers automatic pull-based booking of the drone service via Shipment Request, followed by provider notification, confirmation, and replicated drone delivery sequences for two destinations.

The purpose of this section is therefore to assess what has been empirically evidenced in terms of interoperability, auditable orchestration, and event-governed intermodal coordination under regulated conditions. Potential service impacts associated with the drone-enabled leg are addressed separately in Section 4.1.

4.1. Drone-enabled operational impact

While the previous section focuses on demonstrated evidence, the drone-enabled leg is also relevant because it addresses a broader set of operational impacts that are particularly meaningful in geographically constrained direct-to-patient settings. In this respect, the drone is not introduced merely as an additional transport mode, but as a service enabler for time-sensitive, accessibility-constrained, and territorially fragile delivery contexts.

This subsection therefore captures the main impact dimensions associated with the drone-enabled extension. The aim is to make the expected impact dimensions explicit and measurable, while keeping them analytically separate from the empirically validated interoperability and orchestration results presented above. The quantitative impact of the drone-enabled leg is detailed in Table 3 (Pirri et al., 2024).

Table 3: Quantitative impact dimensions of the drone-enabled leg

Impact dimension	Operationalization	Quantitative indicator	Results
Transport time reduction	Difference between conventional last-mile delivery time and drone-enabled coordinated delivery time.	Average delivery time reduction (%) or minutes saved per mission.	Estimated 35–55% reduction.
Reduced traffic and environmental burden	Reduction in road-based vehicle movement due to partial substitution of the last-mile ground segment with the drone leg.	Vehicle-km avoided; estimated CO2 emissions avoided; reduction in local road traffic burden.	1.2–2.5 kg of CO2 emissions saved per mission
Improved accessibility in constrained territories	Ability to maintain or improve service to remote or hard-to-reach destinations	Increased service coverage.	Up to 100% for remote or hard-to-reach areas

5. Results and discussion

The July 2025 pilot demonstrated that the planned custody milestones could be captured in machine-readable form across the hospital-to-patient drone-enabled last-mile flow, from logistic-unit creation and aggregation to despatch, receipt, deconsolidation, and proof of delivery. From a Physical Internet perspective, this is important because the system did not merely generate isolated sensor readings; it produced a coherent event history with operational and business meaning. The cooperation platform also validated the correlation between GS1 EDI despatch and receipt messages and the corresponding EPCIS transaction history across heterogeneous digital providers.

The October 2025 integrated multimodal demonstration extended this evidence upstream by including the road leg from depot to hospital and its synchronization with the subsequent drone-enabled direct-to-patient segment. This broader scenario showed that the hospital receipt milestone could act as a validated release condition for the downstream drone service through standardized EDI/EPCIS interactions. As a result, the architecture moved beyond simple multimodal visibility and demonstrated a concrete form of intermodal synchronization, in which upstream physical availability and downstream service activation were linked through auditable event-based control.

Device and integration levels validation also produced concrete evidence on operational behaviour. Smart-capsule telemetry reached the cooperation platform at a rate exceeding the minimum project threshold of one update per minute, showing that the architecture can support near-real-time situational awareness. Threshold-based pre-alarm and alarm generation for conservation anomalies and route deviation also worked as intended, indicating that the event backbone supports exception visibility rather than only retrospective traceability. In addition, the smart-box autonomy observed under the relevant monitoring profile was approximately five months, which is compatible with the intended monitoring horizon of the proposed logistics configuration.

Taken together, these results support four main claims: (i) the architecture improves evidence continuity: physical movements, logistic-unit transformations, and shipment documents can be inspected as one trace rather than as fragmented records; (ii) it preserves chain-of-custody integrity across packaging, aggregation, handover, transport, receipt, and delivery through explicit EPCIS events and linked business documents; (iii) it shows that a neutral cooperation platform can coordinate heterogeneous actors without collapsing interoperability into a single-provider information silo, (iv) it demonstrates that road–drone cooperation can be governed through validated milestones and pull-based release logic rather than through a purely static pre-planned itinerary.

More broadly, the results suggest that the main scientific contribution of the work lies not only in combining connected logistics units, EPCIS-based traceability, and drone-enabled delivery, but in showing how these elements can be integrated into a governance-aware orchestration architecture. In this sense, the study contributes a transferable design logic for regulated multimodal healthcare logistics: downstream service execution is not released by schedule alone, but by validated network state, shared events, and auditable control rules. The main result is therefore not only traceability, but auditable release coordination across heterogeneous actors and modes.

From a Physical Internet perspective, this suggests a broader research direction: moving from the pull-based activation of a single transport leg to a network-wide pull architecture in which connected logistics units, neutral event-sharing platforms, and interoperable service interfaces collectively govern when operations are released, postponed, rerouted, or recomposed. In regulated healthcare logistics, such an approach may be especially valuable because it combines flexibility with explicit control points, thereby preserving traceability, chain-of-custody integrity, and compliance evidence while still enabling adaptive multimodal coordination.

At the same time, this generalization should be interpreted with caution. The current work validates the feasibility and relevance of milestone-based release logic in coordinated road–drone execution, but it does not yet demonstrate a fully implemented end-to-end pull architecture across the entire clinical trial supply chain. Extending the mechanism upstream would require additional coordination rules, broader partner participation, and further validation under heterogeneous clinical pathways, geographical contexts, and service

configurations. For this reason, the present study should be understood as a validated step toward full pull orchestration rather than as its complete operational realization.

Still, generalization is conditional rather than automatic. Operational envelope limits for drones, local aviation rules, hospital landing or handover infrastructure, and the willingness of stakeholders to share standardized data all shape feasibility. The architecture should therefore be read as a transferable blueprint with domain-specific adaptation points, not as a universally plug-and-play solution. From a failure-mode perspective, the available evidence covers threshold exceedance and route-deviation alerts, as well as asynchronous ingestion and cross-provider document-event correlation. What remains to be documented more systematically are stress cases such as repeated provider non-confirmation, first-message loss during a critical handover, large-scale exception bursts, or contested custody evidence. These represent priority topics for the next validation cycle.

6. Generalizability and limitations

The main result is therefore not only traceability, but auditable release coordination across heterogeneous actors and modes. This is the case for three main reasons: (i) the connected logistic unit abstraction is not drone-specific; (ii) the neutral cooperation platform is provider-agnostic, provided that participating actors can expose shared identifiers and machine-readable events or documents; and (iii) the orchestration logic is grounded in validated milestones, patient constraints, and auditable release conditions that remain relevant across multiple trial designs, geographical contexts, and service configurations.



Figure 3: Patient-driven, end-to-end pull orchestration across the supply chain

More importantly, the same principle can be extended beyond the last-mile drone leg to the broader end-to-end clinical trial supply chain. In such a scenario, the pull signal would originate when a patient delivery request is confirmed, and a delivery slot is selected that is compatible with supply chain response times and operational constraints. Figure 4 illustrates this broader generalization perspective, in which the validated road–drone synchronization logic is interpreted as one instance of a wider patient-driven and event-governed pull orchestration model. This downstream commitment would then structure the organization of upstream and midstream activities, allowing routes, handovers, and service allocations to be scheduled according to pull principles around an actual delivery need rather than around forecast-based pre-positioning alone. Within this architecture, planning still plays an important role, but it is activated by a real demand commitment and subsequently governed through milestone validation, custody continuity checks, and condition-monitoring events. Road–drone

synchronization therefore becomes one instance of a broader principle: a supply chain that is scheduled from downstream demand and released through auditable event-based control.

More importantly, the same principle can be extended beyond the last-mile drone leg to the broader end-to-end clinical trial supply chain. In such a scenario, the pull signal would originate when a patient delivery request is confirmed, and a delivery slot is selected that is compatible with supply chain response times and operational constraints. Figure 4 illustrates this broader generalization perspective, in which the validated road–drone synchronization logic is interpreted as one instance of a wider patient-driven and event-governed pull orchestration model. This downstream commitment would then structure the organization of upstream and midstream activities, allowing routes, handovers, and service allocations to be scheduled according to pull principles around an actual delivery need rather than around forecast-based pre-positioning alone. Within this architecture, planning still plays an important role, but it is activated by a real demand commitment and subsequently governed through milestone validation, custody continuity checks, and condition-monitoring events. Road–drone synchronization therefore becomes one instance of a broader principle: a supply chain that is scheduled from downstream demand and released through auditable event-based control.

7. Conclusion

This paper presented TREATMENT as a Physical-Internet-enabled architecture for regulated direct-to-patient clinical trial logistics. The core contribution is not the use of standards per se, but their integration into a coherent orchestration model that combines connected logistic units, a neutral multi-tenant cooperation platform, GS1 EPCIS 2.0 event sharing, GS1 EDI document exchange, and milestone-driven road–drone synchronization.

The reported evidence shows that this architecture can deliver more than fragmented tracking. It can generate a single auditable trace across packaging, aggregation, transport, handover, receipt, and delivery, while preserving chain-of-custody integrity and correlating business documents with field events across heterogeneous providers. The October integrated demonstration further indicates that upstream road-leg evidence and hospital receipt can act as validated release conditions for the downstream drone leg, thereby moving from simple multimodal visibility toward event-governed intermodal coordination.

From a Physical Internet perspective, the relevance of the study lies in showing how a regulated healthcare logistics system can be organized around shared identifiers, interoperable event semantics, modular connected units, and neutral coordination mechanisms. In this sense, TREATMENT contributes a governance-aware interoperability model for direct-to-patient logistics: execution is not released by schedule alone, but by validated network state, auditable milestones, and explicit control rules.

At the same time, the present results should be interpreted within their current evidential boundary. The paper demonstrates interoperability and explainable coordination under regulated conditions, but it does not yet establish controlled comparative gains over conventional delivery configurations in terms of lead time, cost, or emissions. Likewise, the broader end-to-end pull interpretation should be read as a credible architectural extension rather than as a fully implemented supply-chain capability.

Overall, TREATMENT provides both a validated field case and a transferable research direction. It shows that milestone-based, standards-enabled, and provider-agnostic orchestration is feasible in a demanding direct-to-patient setting, and it opens a clear path toward broader pull-based, multimodal, and patient-driven healthcare logistics architectures.

References

- Ballot E., B. Montreuil, R.D. Meller (2014): The Physical Internet: The Network of Logistics Networks. La Documentation Française, Paris.
- Crainic T.G., B. Montreuil (2016): Physical Internet Enabled Hyperconnected City Logistics. Transportation Research Procedia, 12, 383-398. https://doi.org/10.1016/j.trpro.2016.02.074
- European Medicines Agency (2022): Recommendation paper on decentralised elements in clinical trials. European Medicines Regulatory Network / EudraLex Volume 10.
- Food and Drug Administration (2024): Conducting Clinical Trials With Decentralized Elements: Guidance for Industry, Investigators, and Other Interested Parties. U.S. Department of Health and Human Services, Silver Spring, MD.
- GS1. (2022). *EPCIS/CBV 2.0: Ready to revolutionise supply chain visibility*. GS1 AISBL
- GS1 (2022b): GS1 Pharmaceutical Clinical Trial Electronic Messaging Standard Implementation Guideline. GS1 AISBL, Brussels.
- Heath M., A.J. de Jong, S. Magorrian-Spence, C. Jin, D.R. van Weelij, L. Pagnier, Y. van Rijswick, V. Haufe, B. Lagerwaard, M.G.P. Zuidgeest (2025): The Supply of Investigational Medicinal Product and Management of Study Materials for Decentralized Participants-Insights from the Trials@Home RADIAL Proof-of-Concept Trial. Clinical Pharmacology & Therapeutics, 118(5), 1079-1089. https://doi.org/10.1002/cpt.70072
- Connie A. Lin, Karishma Shah, Lt Col Cherie Mauntel, Sachin A. Shah, Drone delivery of medications: Review of the landscape and legal considerations, American Journal of Health-System Pharmacy, Volume 75, Issue 3, 1 February 2018, Pages 153–158, https://doi.org/10.2146/ajhp170196
- Montreuil B. (2011): Toward a Physical Internet: meeting the global logistics sustainability grand challenge. Logistics Research, 3(2-3), 71-87. https://doi.org/10.1007/s12159-011-0045-x
- Montreuil B., R.D. Meller, E. Ballot (2013): Physical Internet Foundations. In Borangiu T., A. Thomas, D. Trentesaux (eds.), Service Orientation in Holonic and Multi-Agent Manufacturing and Robotics, Studies in Computational Intelligence, 472, 151-166. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-35852-4_10
- F. Maciariello, F. Benelli, G. Sangiuolo, E. Lorenzi, C. Caponio and C. Salvadori, "TrackOne: Smart Logistics for a Sustainable and Interoperable Agricultural Supply Chain in the Era of Digitization," 2025 International Conference on Software, Telecommunications and Computer Networks (SoftCOM), Split, Croatia, 2025, pp. 1-7, doi: 10.23919/SoftCOM66362.2025.11197431
- Pan S., E. Ballot, G.Q. Huang, B. Montreuil (2017): Physical Internet and interconnected logistics services: research and applications. International Journal of Production Research, 55(9), 2603-2609. https://doi.org/10.1080/00207543.2017.1302620
- Rogers A, De Paoli G, Subbarayan S, et al. A systematic review of methods used to conduct decentralised clinical trials. Br J Clin Pharmacol. 2022;88(6):2843-2862. doi:10.1111/bcp.15205
- Marino, F., Seitanidis, I., Dao, P., Bocchino, S., Castoldi, P., & Salvadori, C. (2019, July). IoT enabling PI: towards hyperconnected and interoperable smart containers. In 6th International Physical Internet Conference (pp. 349-362).
- Stierlin, N., Risch, M., & Risch, L. (2024). Current Advancements in Drone Technology for Medical Sample Transportation. Logistics, 8(4), 104. https://doi.org/10.3390/logistics8040104
- Pirri, A.; Tortora, G.; Niglio, F. (2024). Autonomous drones and smart containers in medicine: Applying drones to current societal and industrial challenges. Drones 2024, 8, 250.

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE

Open Distribution Networks under Compound Stress: A Multi-Agent Stress-Testing Framework

M.Metin Inan [1], Ashwin S. Pothen[1], Matthieu Lauras[1,2], Benoit Montreuil[1,3]
1. Physical Internet Center, Supply Chain and Logistics Institute H. Milton Stewart School of Industrial & Systems Engineering, Georgia Institute of Technology, Atlanta, United States
2. The Centre of Excellence in Supply Chain and Transportation (CESIT), Kedge Business School, France
3. Coca-Cola Material Handling & Distribution Chair
Corresponding author: minan@gatech.edu

Abstract: *This paper presents a multi-agent stress-testing framework comparing two distribution architectures under compound disruptions: a traditional network with three privately operated distribution centers, and a hyperconnected Physical Internet network accessing thirteen shared Open Certified Storage Centers (OCSCs) through PI lane-tariff pricing and annual capacity reservation. The simulator integrates weekly inventory positioning, production planning, raw-material procurement, and daily customer fulfillment with willingness-to-wait acceptance. The stress-test library contains 341 peak-timing scenario definitions across single-node, concurrent two-node, and regional-zone disruption classes, executed with five random seeds for 1,710 simulation runs. At baseline, the hyperconnected architecture achieves a 95% fill rate and $33.6M profit versus 83% and $27.7M for the traditional architecture: a structural gap driven primarily by the denser OCSC footprint and its effect on customer lead-time acceptance. Under compound disruption, the hyperconnected architecture remains within 1–2% of its own baseline profit across all scenario classes including catastrophic severity, while the traditional architecture loses up to 48% of profit under simultaneous two-node hazards. Disruption-period fill rates exceed 88% for the hyperconnected architecture across every class, compared with 41–65% for the traditional architecture.*

Keywords: *Physical Internet, Hyperconnected Logistics, Supply Chain Readiness, Stress Testing, Open Distribution Networks, Compound Disruptions.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☒ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☒ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☒ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Modes*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

1 Introduction

Supply chains increasingly operate in environments where disruptions propagate through interdependent nodes, lanes, suppliers, inventories, and demand zones rather than affecting isolated entities (Dolgui et al., 2018; Li et al., 2021). At the same time, many distribution networks remain organized around firm-centric, dedicated assets, with limited

ability to access capacity beyond the boundaries of the focal firm. This creates a tension between the system-wide flexibility required for disruption mitigation and the siloed logistics structures that still dominate many distribution operations.

The Physical Internet (PI) offers a long-term vision for resolving this tension (Montreuil, 2011; Montreuil, 2013). PI promotes open, interoperable, and standardized logistics networks in which storage, transportation, and handling resources can be accessed through shared service interfaces rather than exclusive ownership alone (Montreuil et al.,2012). While the conceptual foundations of PI have been extensively developed, the operational consequences of PI-style openness under compound disruptions remain insufficiently quantified. In particular, there is limited evidence on whether hyperconnected networks that access shared distribution capacity can sustain service performance when disruption effects accumulate across node availability, inventory integrity, transportation, and customer service.

This paper addresses that gap through a direct-to-consumer furniture distribution case inspired by the current network structure of an industry partner. The partner operates in the furniture sector and serves geographically dispersed customer demand through a firm-centric distribution system. Building on prior work in proactive supply chain readiness assessment and hyperconnected supply chains (Pothen et al., 2023; Inan et al., 2024), we compare two architectures. The first, referred to as the Base architecture, represents a traditional dedicated network with three privately operated distribution centers (DCs). The second, referred to as the Hyperconnected architecture, replaces the dedicated downstream structure with access to thirteen shared Open Certified Storage Centers (OCSCs) through an Open Distribution Network (ODN). OCSCs are PI-certified nodes for storing, handling, and shipping PI containers. Transportation in this architecture is priced through a density-driven Physical Internet lane tariff.

The paper makes three contributions. First, it operationalizes ODN-enabled hyperconnectivity through explicit capacity reservation and density-driven tariff pricing for shared storage and transportation services. Second, it introduces a stress-test taxonomy that distinguishes independent operational hazards from single-cause regional disasters, allowing each disruption class to be modeled according to its causal structure. Third, it develops a geographically grounded multi-agent simulator that links annual OCSC reservation, weekly inventory positioning, weekly production planning, periodic raw-material procurement, and daily customer fulfillment with willingness-to-wait acceptance.

2 Literature Review and Research Gap

The Physical Internet literature argues for a transition from fragmented, firm-centric logistics systems toward open, interoperable networks with standardized services, interfaces, and routing protocols (Montreuil, 2011; Montreuil, 2013). Within this paradigm, warehousing and distribution assets are no longer treated only as privately owned infrastructure, but as components of a broader logistics web that can be accessed through common service logic. Existing work has quantified the economic and environmental benefits of horizontal collaboration (Ballot and Fontane, 2008), studied demand-supply alignment under hyperconnected production options (Pothen et al., 2023) and explored PI concepts in extreme environments such as humanitarian logistics (Grest et al., 2021). However, much of this work remains either conceptual, architectural, or

focused on nominal efficiency. Fewer studies examine how PI-style openness changes operational performance when the distribution system is exposed to structured stress scenarios (Yang et al., 2017).

A large resilience literature studies how supply chains absorb, adapt to, and recover from disruptions, often emphasizing robustness, recovery time, and post-disruption restoration. More recent work has emphasized preparedness and readiness, shifting attention from recovery after failure to whether a system is already positioned to remain within acceptable performance bounds when disruptions occur (Chang and Shinozuka, 2004; Inan et al., 2024). Stress testing extends this perspective by exposing a system to controlled adverse conditions and evaluating the resulting performance distribution across a designed scenario library (Ivanov and Dolgui, 2022; Dolgui et al., 2025). Our work adopts this stress-testing perspective but applies it to a comparative architectural question: whether an ODN-enabled hyperconnected distribution system degrades less severely than a dedicated distribution system under matched disruption conditions.

Disruption interactions are increasingly recognized as a defining feature of real supply chains. Prior work on ripple effects has shown that local shocks can propagate through capacity imbalances, inventory depletion, transportation delays, and synchronized regional failures (Li et al., 2021). Yet many compound-disruption studies still assume closed logistics networks, fixed recovery options, or simplified network-control mechanisms. As a result, they provide limited insight into how access to shared distribution capacity changes the propagation and mitigation of disruption consequences. This is particularly important for PI research, because the value of openness is not only that more nodes exist, but that capacity, transportation, and inventory positioning can be governed through a more flexible operating structure.

A related stream studies flexible capacity sourcing, reservation contracts, and real options for logistics or storage capacity (Park and Kim, 2014; Yang et al.,2017). These studies show that contractual flexibility can improve performance when dedicated assets are insufficient or geographically mismatched. However, reservation models are often studied separately from daily fulfillment, transportation pricing, customer lead-time acceptance, and disruption propagation. This paper adds an ODN interpretation by embedding capacity reservation inside a multi-agent operating system: annual OCSC reservations represent PI-consistent shared-capacity commitments, while density-driven PI lane tariffs represent a transport mechanism through which external network density affects operational cost.

Taken together, these streams leave a clear gap. PI research motivates open logistics networks; resilience research motivates stress testing; and capacity-flexibility research motivates shared external access. What remains underdeveloped is an integrated operational framework that connects these ideas and evaluates them under controlled compound disruptions. This paper addresses that gap.

3 Problem Setting and Case Study

We study a direct-to-consumer flat-pack furniture manufacturer operating one production plant in the El Paso region. The case setting is inspired by an industry partner in the furniture sector and represents a national U.S. distribution network with approximately 6,000 units of weekly demand. Experiments use three product types representing

common, regular, and unique furniture products, with different prices, storage volumes, demand shares, supplier lead-time exposure, and willingness-to-wait curves. Customer demand is represented at the county level, allowing the simulator to capture how distribution-node geography affects promised delivery lead times and customer acceptance.

The downstream architecture differs across the two cases. The Base architecture represents the current firm-centric structure: three privately operated distribution centers located near Nashville, Salt Lake City, and El Paso. Capacity is committed through a fixed monthly DC cost plus a variable unit-month charge, and plant-to-DC replenishment uses traditional full-truckload (FTL) and less-than-truckload (LTL) pricing. The Hyperconnected architecture replaces the dedicated DC footprint with thirteen shared OCSCs accessed through Open Distribution Network logic. Reserved OCSC capacity is priced per unit-month with no fixed facility commitment, and plant-to-OCSC as well as inter-OCSC flows are priced through a PI lane tariff. Both architectures face identical demand realizations, plant technology, supplier lead times, product economics, and customer-facing parcel pricing.

The geographic consequence of this difference is visible in Figure 1: the Base architecture leaves large portions of the Northeast, Northwest, and Southeast facing 5–7 day promised lead times, while the Hyperconnected architecture achieves 1–3 day coverage across nearly the entire continental United States.

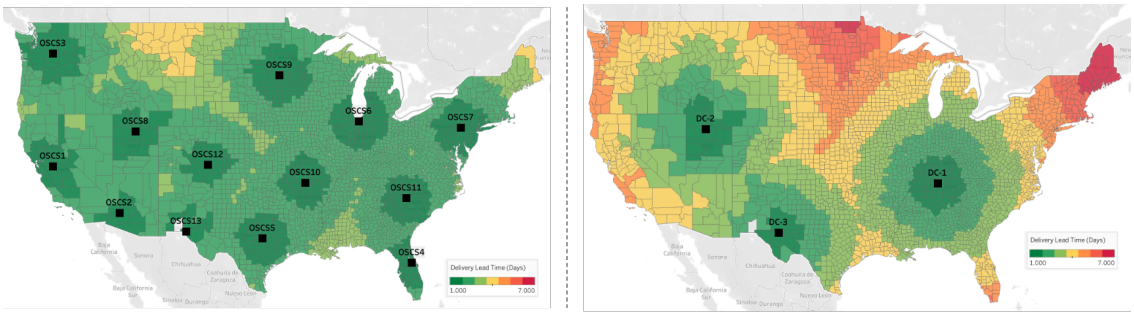


***Figure 1**: Customer-facing delivery lead time by county under the two architectures.*

4 Simulator Architecture and Decision Logic

4.1 Multi-agent architecture

The simulator is organized as a three-layer multi-agent system, illustrated in Figure 2. The Decisional Layer contains five agents. The Distribution agent solves weekly inventory positioning and, in the Hyperconnected architecture, lateral OCSC transfers. The Manufacturing agent plans production and consumes raw materials. The Supply agent places raw-material orders based on the production plan and supplier lead times. The Fulfillment agent routes daily customer orders to the closest stocked node that satisfies willingness-to-wait acceptance. The ODN Orchestrator represents the external ODN operator: it publishes lane tariff prices and per-node reservation bounds and accepts or rejects capacity requests from the Distribution agent. In the reported experiments, annual beginning-of-year reservation is active; during-year extension requests are reserved for future work. These agents do not operate with perfect information or centralized coordination: each agent observes the current system state, accesses shared forecasts and

targets through the Cognitive Layer, and issues decisions within its own scope. The Cognitive Layer provides the shared infrastructure that makes coordinated planning possible: a rolling demand forecast, inventory targets, performance signals, and access to the optimization solvers. The Physical Layer represents the actual supply chain: suppliers, the production plant, distribution and fulfillment nodes, transport assets, and customers.

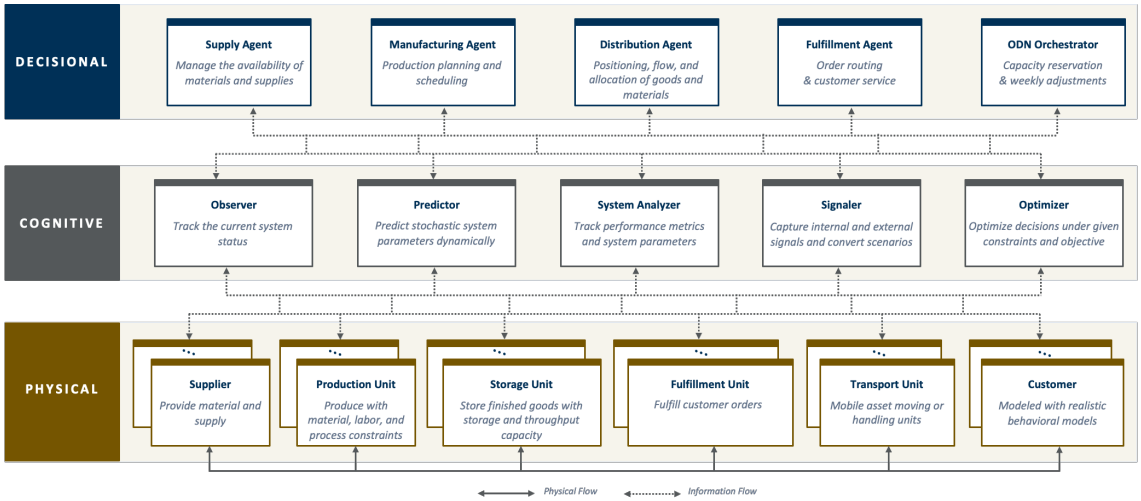


***Figure 2**: Three-layer multi-agent simulator architecture.*

4.2 Temporal cadence and information flow

The simulator runs on a daily clock for 364 days (52 weeks). Decisions operate at four cadences: annual (OCSC reservation), monthly (capacity activation), weekly (positioning, production, procurement), and daily (order generation, fulfillment, shipment processing, KPI recording).

Each week begins by refreshing the rolling demand forecast using only demand observed in prior completed weeks and updating inventory targets accordingly. The positioning model then schedules finished-goods movements over the coming seven days. The production model commits the current week's launch. Procurement orders are placed based on the production plan and confirmed supplier availability.

This ordering means neither model sees future demand directly. Both plan from the realized state, with forecast uncertainty explicitly represented; so, when disruptions hit, the agents react through the same planning loop they use at baseline, with no privileged access to disruption schedules or recovery timelines.

4.3 Decision Logic and PI Operating Mechanisms

Annual OCSC reservation. At the start of the simulation year, the Hyperconnected architecture solves a network-level reservation model to determine how much capacity to commit at each OCSC for each of the twelve planning months. The model minimizes deviations from forecast-driven monthly capacity targets at both the node and network levels, subject to per-node minimum and maximum reservable bounds and binary activation decisions.

Weekly production planning. Each Monday, the Manufacturing agent solves a joint multi-product production model over a 24-week rolling horizon. The model determines how many units of each product to launch into production this week and in upcoming weeks, subject to three classes of constraint: shared plant process capacity (each product requires time on two of five processes), raw-material availability within each supplier's lead-time window (raw materials have 7-, 14-, or 28-day supplier lead times, and the model cannot plan production that depends on raw materials not yet contractually secured), and system-wide storage volume. The objective penalizes unmet demand at the product's lost-sales rate, always more severely than any inventory positioning shortfall; so projected unmet demand is always prioritized over target-buffer shortfalls.

Weekly time-expanded positioning. The Distribution agent solves a seven-day time-expanded network model each Monday, scheduling plant-to-node shipments across each day of the coming week. For the Hyperconnected architecture, the model also schedules OCSC-to-OCSC lateral transfers: moving inventory between shared nodes when local demand or a disruption creates a geographic imbalance. This lateral repositioning capability is the primary operational mechanism through which the Hyperconnected architecture absorbs single-node failures: when an OCSC becomes unavailable, its demand can be redirected to adjacent OCSCs, and inventory can be repositioned toward them in the same weekly planning cycle. The Base architecture has no lateral repositioning mechanism and far fewer alternative service points; when a DC fails, affected demand must be served from more distant DCs, increasing rejection risk.

5 Stress-Test Scenario Taxonomy and Experimental Design

The stress-test library is organized around two causal models for compound disruption. The first treats facility outages as independent operational hazards (fires, labor strikes, IT outages, or equipment failures) affecting one facility or coinciding at two. The second treats disruption as a single regional trigger, such as severe weather or a regional grid failure, whose footprint affects a geographic risk pool. Regional triggers are anchored to the Base DC regions; under the same trigger, the Hyperconnected architecture loses the OCSC pair occupying the corresponding risk pool.

The three active scenario classes are summarized in Table 1. S-N captures single-node independent hazards. C-N2 captures two independent node hazards, either simultaneous or staggered. R-Z captures a single regional trigger affecting a matched risk pool. Each affected node operates in Mode A (logistics block, no inventory damage) or Mode B (logistics block plus finished-goods inventory damage at event start). Mode A is exercised at Mild severity only; Mode B at Mild, Moderate, Severe, and Catastrophic severities. Table 2 summarizes the four severity tiers.

Table 1: *Scenario classes used in the stress-test library*

Class	Interpretation	Architecture exposure
S-N	Single-node independent operational hazard	3 DCs; 13 OCSCs
C-N2	Two independent node hazards, simultaneous or staggered	3 DC pairs; 9 OCSC pairs
R-Z	A regional trigger affecting one geographic risk pool	3 Base regions; 3 matched Hyperconnected regions

Disruption start is fixed at the seasonal demand peak, day 150, the most adverse timing because demand is high, buffers are stressed, and recovery slack is limited. For C-N2, the inter-event gap varies across three stagger tiers: simultaneous, 14 days, and 30 days. R-Z is single-trigger and always simultaneous. Combining these dimensions yields 341 distinct scenario definitions and 1,710 simulation runs including baseline after five-seed replication.

Each architecture-scenario combination is simulated with five random seeds. When both architectures face a comparable scenario, the same seed generates identical stochastic demand and background noise, enabling paired comparison. Primary outputs are annual fill rate, disruption-period fill rate, annual profit, full cost decomposition, stockout units, willingness-to-wait rejection units, and lost-sales value. Loss-vs-baseline metrics use the matched-seed baseline for the same architecture and include maximum daily fill-rate drop, total annual units lost, and time-to-recover defined as days from disruption end until the seven-day rolling fill rate returns to at least 95% of the matched baseline. Catastrophic scenarios may be right-censored when recovery is not observed within the 364-day horizon.

Table 2: *Severity tiers used in the stress-test library*

Severity	Duration	Mode-B inventory loss	Interpretation
Mild	30 days	25%	Recoverable one-month disruption
Moderate	90 days	50%	Multi-month operating disruption
Severe	150 days	70%	Long disruption with major damage
Catastrophic	-	100%	Node unavailable for rest of simulation

6 Results

6.1 Baseline performance and structural service gap

The Base architecture achieves an annual fill rate of 0.826 and a profit of $27.7M, while the Hyperconnected architecture achieves an annual fill rate of 0.951 and a profit of $33.6M. This corresponds to a 12.5 percentage-point service advantage and a $5.9M profit advantage for the Hyperconnected architecture.

The baseline results are important because they show that the service gap is not primarily caused by optimization noise or random operational frictions. The gap is structural: the denser OCSC footprint places inventory closer to demand, shortens promised lead times, and improves customer acceptance, particularly for time-sensitive common products where acceptance drops sharply past two days.

The cost decomposition further clarifies the tradeoff. Total operating cost is similar across architectures: $33.0M for Base and $33.9M for Hyperconnected. However, the composition differs. The Base architecture pays $3.27M in fixed and variable DC capacity costs, while the Hyperconnected architecture pays $0.46M in OCSC reservation cost. This capacity-cost advantage is partially offset by higher internal transport cost in the Hyperconnected architecture, driven by additional plant-to-OCSC and inter-OCSC movements. The decisive difference is revenue capture: the Base architecture incurs $10.0M in lost-sales value, compared with $3.2M for the Hyperconnected architecture.

Thus, the Hyperconnected architecture converts a comparable operating-cost base into substantially higher revenue and profit by reducing lead-time-driven customer rejection.

6.2 Performance under compound disruption

Figure 3 summarizes annual fill rate and profit distributions across the stress-test scenario classes. Two patterns emerge. First, the Hyperconnected architecture remains close to its baseline service level across all disruption classes. Its median annual fill rate remains in the 0.92–0.93 range under single-node hazards, concurrent two-node hazards, and regional-zone disasters.

Second, the Base architecture degrades more substantially and exhibits wider dispersion. Its median annual fill rate drops from 0.826 at baseline to approximately 0.79 under single-node hazards, 0.76 under regional-zone disasters, and 0.76 under concurrent two-node hazards. The lower tail is especially severe under concurrent two-node scenarios, where the worst seed-scenario combinations fall below 0.50 fill rate.

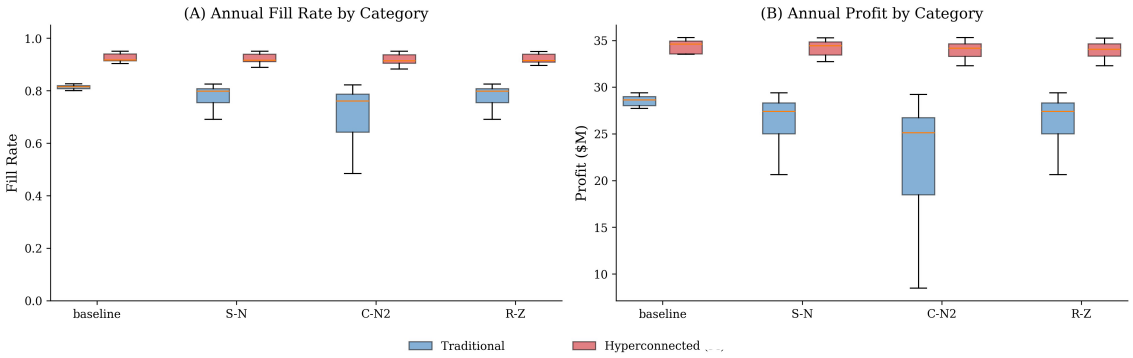


Figure 3: *Annual fill rate and profit distributions by scenario class.*

Profit retention confirms the same pattern. Under catastrophic single-node hazards, the Base architecture retains 78.1% of baseline profit, while the Hyperconnected architecture retains 99.1%. Under catastrophic concurrent two-node hazards, the Base architecture retains only 52.0% of baseline profit, while the Hyperconnected architecture retains 97.9%. Across all reported class–severity combinations, Hyperconnected profit retention remains between 97.9% and 99.8%, whereas Base profit retention ranges from 97.7% down to 52.0%.

6.3 Disruption-period response and loss decomposition

The architectural gap becomes sharper when performance is measured only during active disruption windows. Table 3 reports disruption-period fill rates by scenario class. The Hyperconnected architecture maintains disruption-period fill rates between 0.887 and 0.896 across all classes. The Base architecture falls to 0.645 under single-node and regional-zone scenarios and to 0.495 under concurrent two-node scenarios. In other words, during the actual stress window, the Hyperconnected architecture maintains close to 90% service, while the Base architecture can lose roughly half of demand under the most severe compound conditions.

The loss decomposition provides additional insight. At baseline, most Base lost sales are willingness-to-wait rejections rather than physical stockouts. This means that the Base

network often has product somewhere in the system but cannot offer it quickly enough to retain the customer.

Table 3: *Disruption-period fill rate by scenario class, mean over within-class variants*

Class	Base	Hyperconnected
S-N	0.645	0.896
C-N2	0.495	0.887
R-Z	0.645	0.892

Under concurrent two-node disruption, however, the stockout share rises substantially as the remaining DC network becomes physically depleted. This shows that severe stress changes the nature of the Base failure mode: the architecture is not only too far from demand, but also increasingly unable to maintain physical availability at the remaining service nodes.

The Hyperconnected architecture behaves differently. Its lost sales remain relatively stable in composition across baseline and disruption cases. Because the OCSC footprint provides more service points and lateral repositioning options, individual node outages do not translate as quickly into systemwide stockouts. Lost sales remain dominated by willingness-to-wait effects rather than empty-shelf effects, indicating that physical availability is better preserved under stress.

7 Conclusion and Future Work

This paper develops a consequence-driven stress-testing framework comparing a traditional DC-based distribution architecture against a hyperconnected ODN-enabled alternative under compound disruptions, using a multi-agent simulator with scenario injection.

Three findings stand out. First, the Hyperconnected architecture has a structural service advantage even before injected disruption occurs. Its denser OCSC footprint places inventory closer to demand, increasing customer acceptance and producing a 12.5 percentage-point baseline fill-rate advantage over the Base architecture. This gap reflects network geography rather than operational tuning: only 18% of demand reaches a Base DC within two days, versus 97% under the Hyperconnected architecture. Second, the architectural gap widens under stress. The Hyperconnected architecture retains between 97.9% and 99.8% of its baseline profit across all scenario class–severity combinations. The Base architecture falls as low as 52.0% profit retention under catastrophic concurrent two-node hazards. During active disruption windows, the Hyperconnected architecture maintains fill rates near 89%, while the Base architecture drops to 49.5% under concurrent two-node stress. Third, disruption changes not only how much the Base architecture loses, but how it loses. At baseline, 94% of Base lost sales are willingness-to-wait rejections and only 6% are physical stockouts: the architecture fails primarily because it cannot reach customers fast enough, not because it lacks inventory. Under concurrent two-node disruption, the stockout share rises to 32% as the remaining DC network becomes physically depleted. The Hyperconnected architecture maintains a stable 17–18% stockout share across all conditions, indicating that its redundancy preserves physical availability even under stress.

The present scope focuses on node and regional distribution stress at peak seasonal demand. Future work will extend the framework along five dimensions: inventory autonomy policy, ODN market regime, disruption timing sensitivity, during-year ODN extension dynamics, and shipment and storage cost sensitivity.

References

- Ballot E., F. Fontane (2008): Rendement et efficience du transport: un nouvel indicateur de performance. Revue Française de Gestion Industrielle, v27, no2, 41–55.
- Chang S.E., M. Shinozuka (2004): Measuring improvements in the disaster resilience of communities. Earthquake Spectra, v20, no3, 739–755.
- Dolgui A., D. Ivanov, B. Sokolov (2018): Ripple effect in the supply chain: an analysis and recent literature. International Journal of Production Research, v56, no1–2, 414–430. https://doi.org/10.1080/00207543.2017.1387680
- Dolgui A., D. Ivanov, D. Simchi-Levi (2025): Stress tests for supply chains: towards resilience and viability. International Journal of Production Research, v63, no9, 3254–3258. https://doi.org/10.1080/00207543.2025.2483113
- Grest M., M.M. Inan, Y.M. Cohen, A.V. Barenji, M. Dahan, M. Lauras, B. Montreuil (2021): Design of a simulation-based experiment for assessing the relevance of the Physical Internet concept for humanitarian supply chains. Proceedings of IPIC 2021 — 8th International Physical Internet Conference, 10 p.
- Inan M.M., B. Montreuil, M. Lauras, M. Dahan (2024): Assessing and enhancing readiness in hyperconnected supply chains. Proceedings of IPIC 2024 — 10th International Physical Internet Conference, 12 p. https://doi.org/10.35090/gatech/11078
- Ivanov D., A. Dolgui (2022): Stress testing supply chains and creating viable ecosystems. Operations Management Research, v15, 475–486. https://doi.org/10.1007/s12063-021-00194-z
- Li Y., K. Chen, S. Collignon, D. Ivanov (2021): Ripple effect in the supply chain network: forward and backward disruption propagation, network health and firm vulnerability. European Journal of Operational Research, v291, no3, 1117–1131.
- Montreuil B. (2011): Towards a Physical Internet: meeting the global logistics sustainability grand challenge. Logistics Research, v3, no2–3, 71–87. https://doi.org/10.1007/s12159-011-0045-x
- Montreuil B., R.D. Meller, E. Ballot (2012): *Physical Internet foundations*. IFAC Proceedings Volumes, v45, no6, 26–30.
- Montreuil B. (2013): Physical Internet Manifesto, v1.10. www.physicalinternetinitiative.org
- Park S.I., J.S. Kim (2014): A mathematical model for a capacity reservation contract. Applied Mathematical Modelling, v38, no5–6, 1866–1880.
- Pothen A.S., M.M. Inan, B. Montreuil, M. Lauras, F. Benaben, Y. Xie (2023): Demand-supply alignment in supply chain networks with access to hyperconnected production options. Proceedings of IPIC 2023 — 9th International Physical Internet Conference, 494–504. https://doi.org/10.35090/gatech/8030
- Yang Y., S. Pan, E. Ballot (2017): Mitigating supply chain disruptions through interconnected logistics services in the Physical Internet. International Journal of Production Research, v55, no14, 3970–3983. https://doi.org/10.1080/00207543.2016.1223379

IPIC 2026
12th International
Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE

Interactive Service and Blockchain Traceability in a Cyber-Physical Internet Service Network

Jialing Lin [1], Mengdi Zhang [1], Zhiheng Zhao [1, 2,*], George Q. Huang [1, 2]
1. Department of Industrial and Systems Engineering, The Hong Kong Polytechnic University, Hong Kong, China
2. Research Institute for Advanced Manufacturing, The Hong Kong Polytechnic University, Hong Kong, China
Corresponding author: Zhiheng Zhao (zhiheng.zhao@polyu.edu.hk)

Abstract: *Cyber-Physical Internet (CPI) service networks increasingly rely on digital platforms to coordinate heterogeneous service nodes, user demand, and information flows across organizational boundaries. This paper develops an analytical framework to study pricing and governance in such a network under two settings: a benchmark platform model and a blockchain-enabled interactive-service model. The analysis examines how data value affects demand coordination, how interactive service improves requirement matching, and how blockchain traceability reshapes competition between nodes with different credibility levels. The results show that data value improves network performance in both settings, but digital enhancement reallocates demand more selectively toward the higher-credibility node. Blockchain traceability therefore functions not only as a transparency mechanism, but also as a governance instrument that changes allocation and competitive balance within the network. The study offers managerial implications for platform design, service support, and traceability deployment in CPI-enabled systems.*

Keywords: *Physical Internet; Cyber-Physical Internet; service network; interactive service; blockchain traceability; data value*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan: ☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories), ☐ Transportation Equipment, ☐ PI Networks, ☐ System of Logistics Networks, ☐ Vertical Supply Consolidation, ☐ Horizontal Supply Chain Alignment, ☒ Logistics/Commercial Data Platform, ☐ Access and Adoption, ☒ Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐ Flash Video, ☐ In-Person presentation*

1 Introduction

The Physical Internet and its cyber-physical evolution are reshaping how logistics, service, and information systems are organized across open and interconnected networks. Existing studies emphasize that Physical Internet systems rely on interconnection, interoperability, and shared protocols, while recent CPI research further shows that digitalization is the key enabling force behind the transition from PI to Cyber-Physical Internet service systems (Montreuil, 2011; Pan et al., 2017; Wu et al., 2025). In such environments, digital platforms increasingly coordinate heterogeneous service nodes, users, and information flows through shared interfaces and real-time data exchange, making network governance a central issue rather than a purely transactional one (Gawer, 2022; Wu et al., 2024).

This transformation has strong practical relevance. In a Cyber-Physical Internet service network, participating nodes often differ in credibility, historical records, responsiveness, and service quality. As a result, demand allocation depends not only on price, but also on whether the platform can reduce information asymmetry, improve requirement clarification, and provide credible support for service selection and fulfilment. Two mechanisms are especially important in this context. The first is interactive service at the customer interface, which can improve requirement clarification, solution communication, quotation support, and response speed. The second is blockchain traceability, which can strengthen the visibility and credibility of transaction records and fulfilment information across network participants (Agrawal et al., 2021; Ahmed and MacCarthy, 2023; Cui et al., 2023; Saberi et al., 2019).

Existing studies have examined several parts of this problem, but not their integration. Recommender-systems research shows that richer information and feedback can improve matching quality and reduce search frictions (Adomavicius and Tuzhilin, 2005; Park et al., 2012). Platform and operations research further shows that data, learning, and service-side incentives can reshape pricing and coordination decisions in digital environments (Taylor, 2018; den Boer and Keskin, 2022; Choi et al., 2018). Meanwhile, blockchain research emphasizes that traceability is not only a transparency tool but also a governance mechanism that affects trust, verification, and allocation in distributed operations (Agrawal et al., 2021; Ahmed and MacCarthy, 2023; Cui et al., 2023; Saberi et al., 2019).

Our research aims to develop an analytical framework for a Cyber-Physical Internet service network under two settings: a benchmark platform model and a blockchain-enabled interactive-service model. Specifically, this study addresses three research questions. First, how does data value affect pricing, demand coordination, and participant profitability in a Cyber-Physical Internet service network? Second, how does interactive service reshape matching quality and demand allocation across heterogeneous service nodes? Third, how does blockchain traceability change the competitive position of nodes with different credibility levels and affect the governance logic of the network?

This paper makes three contributions. First, it develops a parsimonious analytical framework for pricing and governance in a Cyber-Physical Internet service network with heterogeneous nodes. Second, it clarifies how interactive service and data value jointly reshape demand allocation and participant profitability in a platform-mediated environment. Third, it shows how blockchain traceability strengthens the position of the more credible node and changes the allocation logic of the network, thereby highlighting the governance role of trusted information infrastructure in CPI-enabled systems.

The remainder of the paper is organized as follows. Section 2 reviews the relevant literature. Section 3 presents the problem setting and analytical models. Section 4 reports the main analytical results. Section 5 discusses managerial insights. Section 6 concludes the paper.

2 Literature Review

Research on the Physical Internet and its cyber-physical extensions emphasizes interconnection, interoperability, and the redesign of logistics and service systems around open networks rather than isolated entities. Existing studies highlight interconnected logistics services, heterogeneous actors, and digital coordination mechanisms as key enablers of integrated and responsive network operations (Montreuil, 2011; Pan et al., 2017). More recent CPI research further argues that digitalization is the main enabling force behind the transition from PI to CPI and formalizes CPI as a layered, open, and scalable coordination architecture, while applied work has shown how CPI can support infrastructure integration through cyber, physical, and internet dimensions (Wu et al., 2025; Wu et al., 2024). Related digital-platform

research also shows that platforms are not merely transaction intermediaries but governance structures through which interfaces, data, and rules shape coordination and value capture (Gawer, 2022).

A second stream of literature focuses on interactive service, data, and matching. Foundational recommender-systems research shows that personalization, feedback, and richer user information can improve user-item matching and reduce search frictions (Adomavicius and Tuzhilin, 2005; Park et al., 2012). Related platform and operations research further shows that data, demand learning, and service-side incentives significantly affect pricing and coordination decisions in digital environments (Taylor, 2018; den Boer and Keskin, 2022). At a broader operations-management level, big data analytics has been identified as a strategic capability for improving operational decisions, information use, and coordination quality (Choi et al., 2018). Taken together, these studies imply that interactive service and platform-held data should be treated not only as operational support tools, but also as strategic levers that influence demand allocation and economic performance.

A third stream studies blockchain traceability as a mechanism for transparency, verification, and governance in distributed operations. Existing work shows that blockchain-based architectures can strengthen visibility, authentication, and information sharing across multi-tier operational systems, while the scope and value of traceability depend on network structure, incentive alignment, and the granularity of verifiable records (Agrawal et al., 2021; Ahmed and MacCarthy, 2023; Cui et al., 2023). Broader blockchain research in supply chain management also shows that blockchain has important implications for transparency, traceability, and coordination, but that its benefits depend on adoption conditions and system design (Saberi et al., 2019). This suggests that blockchain traceability should not be understood only as an information technology, but also as a governance mechanism that can reshape trust, bargaining positions, and allocation across network participants.

Taken together, the literature provides useful insights into PI/CPI coordination, data-enabled matching, and blockchain-based traceability. However, these streams remain weakly integrated. Existing PI- and CPI-oriented studies focus mainly on network architecture, interconnection, and infrastructure integration; service-platform research often abstracts from trusted traceability; and blockchain studies rarely examine how interactive service and data value jointly reshape competition between heterogeneous nodes. This paper addresses that gap in a unified analytical setting.

3 Problem Setting and Models

We consider a Cyber-Physical Internet service network in which a platform coordinates heterogeneous service nodes and downstream users through a shared digital interface. To examine how digital enhancement changes network outcomes, two settings are studied: a benchmark platform model and a blockchain-enabled interactive-service model. The notation used in the analytical models is summarized in Table 1.

Table 1: Notation used in the analytical models

Notation	Definition
U_1, U_2	User utilities from Service Options I and II
V_1, V_2	Zero-utility thresholds for Service Options I and II
θ	Baseline credibility parameter between the two service nodes
C_1, C_2	Service costs of Nodes I and II
D_1, D_2	Demands for Service Options I and II
W_1,W_2	Upstream terms of Nodes I and II

P_1, P_2	Downstream prices of Service Options I and II
π_{N_1}, π_{N_2}	Profits of Service Nodes I and II
π_P	Profit of the platform
π_I	Profit of the interactive service operator
Z	Intensity of interactive service
F	Fixed operating cost of the interactive service operator
f	Total blockchain traceability cost
δ	Traceability effect of blockchain
S	Per-transaction data value

3.1 Benchmark platform model

We first consider a benchmark platform setting in which two heterogeneous service nodes provide competing service options through a platform. Service Node I represents a relatively less-established node with lower credibility, whereas Service Node II represents a more established node with stronger credibility and more reliable historical records. The platform accesses service capacity from the two nodes at upstream terms W_1 and W_2, sets downstream prices P_1 and P_2, and derives a per-transaction data value S from platform-mediated interactions, transaction feedback, and data-enabled coordination. The benchmark structure is illustrated in Figure 1.

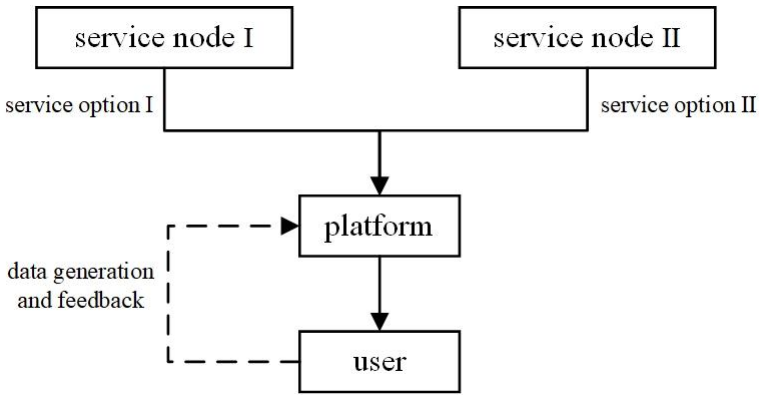


Figure 1: Benchmark Cyber-Physical Internet service network

Let user valuation V be uniformly distributed on $[0,1]$. The baseline credibility difference between the two service nodes is captured by $\theta \in (0,1)$, where a larger θ indicates a smaller credibility gap. The utilities associated with the two service options are defined as $U_1 = \theta V - P_1$ and $U_2 = V - P_2$. Users choose the service option that provides the higher nonnegative utility; otherwise, they do not purchase. Let V_{12} denote the indifferent user between the two service options. Solving $U_1 = U_2$ yields $V_{12} = \frac{P_2 - P_1}{1-\theta}$. The zero-utility thresholds are $V_1 = \frac{P_1}{\theta}$, $V_2 = P_2$. To focus on the competitive two-node regime, we restrict attention to the interior region in which both service options remain active, that is, $0 < V_1 < V_{12} < 1$. Under this condition, users choose Service Option I when $V \in [V_1, V_{12}]$ and Service Option II when $V \in [V_{12}, 1]$. The corresponding demands are therefore $D_1 = V_{12} - V_1$ and $D_2 = 1 - V_{12}$.

The platform's profit consists of the margins earned on both service options and the value created from data utilization:

$$\pi_P = (P_1 - W_1)D_1 + (P_2 - W_2)D_2 + S(D_1 + D_2) \quad (1)$$

The profits of Service Node I and Service Node II are

$$\pi_{N_1} = W_1 D_1 - C_1 D_1 \quad (2)$$

$$\pi_{N_2} = W_2 D_2 - C_2 D_2 \quad (3)$$

The benchmark model is solved as a Stackelberg game with backward induction. The two service nodes first determine their upstream terms, after which the platform sets downstream prices. Demand is then realized according to users' utility comparison. This benchmark model captures a basic Cyber-Physical Internet service network in which pricing, demand allocation, and profitability are jointly shaped by price competition, node credibility differences, and platform data value.

3.2 Blockchain-enabled interactive-service model

We next consider an enhanced setting in which the platform is supported by two additional mechanisms: interactive service and blockchain traceability. Specifically, an interactive service operator assists the platform through requirement clarification, solution communication, quotation support, and real-time response. Let Z denote the intensity of interactive service, and let $\beta \in [0,1]$ capture the relative effectiveness of this support for Service Option I. In addition, a blockchain traceability layer is introduced to improve the visibility and credibility of service records, transaction histories, and fulfillment information. Let $\delta \in (0,1)$ denote the traceability effect. The enhanced structure is illustrated in Figure 2.

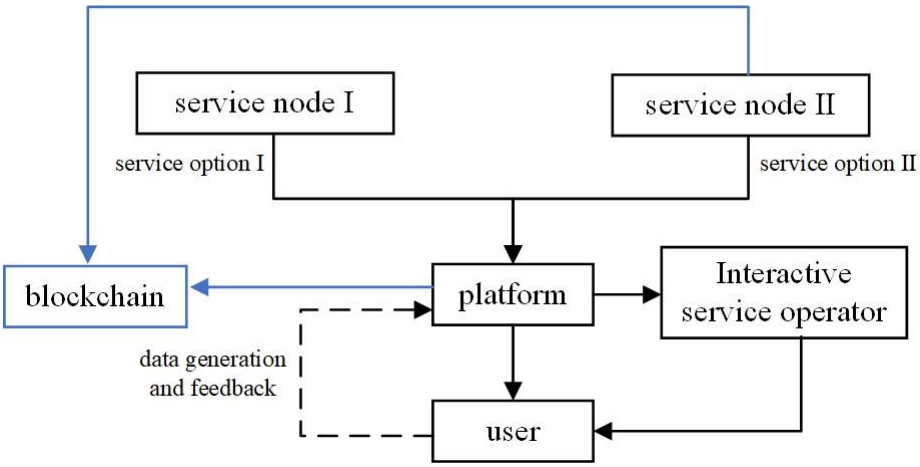


Figure 2: Blockchain-enabled interactive-service Cyber-Physical Internet service network

Under this setting, users' utilities become $U_1 = (\theta - \delta)V - P_1 + \beta Z$ and $U_2 = (1 + \delta)V - P_2 + Z$. Compared with the benchmark model, interactive service directly improves users' perceived value, while blockchain traceability changes the credibility structure of the two service options by strengthening trusted information disclosure. Let V_{12} again denote the indifferent user between the two service options. Solving $U_1 = U_2$ gives $V_{12} = \frac{P_1 - P_2 + (1-\beta)Z}{\theta - 2\delta - 1}$. The zero-utility thresholds are $V_1 = \frac{-P_1 + \beta Z}{\delta - \theta}$ and $V_2 = \frac{P_2 - Z}{\delta + 1}$. As in the benchmark case, we focus on the competitive two-node regime in which both service nodes remain active, that is, $0 < V_1 < V_{12} < 1$. Under this condition, users choose Service Option I when $V \in [V_1, V_{12}]$and Service Option II when $V \in [V_{12}, 1]$. The resulting demands are $D_1 = V_{12} - V_1$ and $D_2 = 1 - V_{12}$.

The platform receives a share λ of the downstream operating return, while the interactive service operator receives the remaining share $1 - \lambda$. Blockchain adoption entails a total traceability cost f, of which the platform bears proportion $\gamma$and Service Node II bears $1 - \gamma$. The platform's profit is

$$\pi_P = [(P_1 - W_1)D_1 + (P_2 - W_2)D_2 + S(D_1 + D_2)]\lambda - \gamma f \quad (4)$$

and the profit of the interactive service operator is

$$\pi_I = [(P_1 - W_1)D_1 + (P_2 - W_2)D_2 + S(D_1 + D_2)](1 - \lambda) - kZ^2 - F \quad (5)$$

where $kZ^2$is the convex cost of interactive service and $F$is a fixed operating cost.

The profits of the two service nodes are

$$\pi_{N_1} = W_1 D_1 - C_1 D_1 \quad (6)$$

$$\pi_{N_2} = W_2 D_2 - C_2 D_2 - (1 - \gamma) f \quad (7)$$

The enhanced model is also formulated as a Stackelberg game and solved by backward induction. Service nodes first determine their upstream terms, after which the platform and the interactive service operator determine downstream decisions, including prices and service intensity. This setting captures a digitally enhanced Cyber-Physical Internet service network in which pricing, demand allocation, and profitability are jointly shaped by data value, interactive service, and blockchain traceability.

4 Analytical Results

4.1 Results in the benchmark platform model

The benchmark model yields an interior equilibrium in the competitive two-node regime. The key comparative-statics results are summarized as follows.

Proposition 1. In the benchmark platform model, a higher data value S increases demand, downstream prices, upstream terms, and profits for all participating actors: $\frac{\partial D_1}{\partial S} > 0$, $\frac{\partial D_2}{\partial S} > 0$, $\frac{\partial P_1}{\partial S} > 0$, $\frac{\partial P_2}{\partial S} > 0$, $\frac{\partial W_1}{\partial S} > 0$, $\frac{\partial W_2}{\partial S} > 0$, $\frac{\partial \pi_P}{\partial S} > 0$, $\frac{\partial \pi_{N_1}}{\partial S} > 0$, $\frac{\partial \pi_{N_2}}{\partial S} > 0$.

This result shows that data value acts as a coordination amplifier. Greater value from user interaction and transaction feedback strengthens both service options, supports stronger prices, and improves the positions of the platform and both nodes.

Proposition 2. The relative positions of the two nodes depend on the credibility parameter θ:

(i) If $\theta < \frac{2S}{1+S}$, $D_1 < D_2$, if $\theta > \frac{2S}{1+S}$, $D_1 > D_2$;

(ii) If $\theta < \frac{6-S}{2}$, $P_1 < P_2$, if $\theta > \frac{6-S}{2}$, $P_1 > P_2$;

(iii) If $\theta < 2 - S$, $W_1 < W_2$, if $\theta > 2 - S$, $W_1 > W_2$;

(iv) If $\theta < S^2$, $\pi_{N_1} < \pi_{N_2}$, if $\theta > S^2$, $\pi_{N_1} > \pi_{N_2}$.

Thus, data value improves overall performance, but the distribution of gains still depends on node credibility. When the credibility gap is large, Service Node II enjoys stronger demand and profitability; when it narrows, Service Node I becomes more competitive.

4.2 Results in the blockchain-enabled interactive-service model

The enhanced model also yields an interior equilibrium in the competitive two-node regime. Because the closed-form expressions are lengthy, the main analytical properties are summarized through comparative statics.

Proposition 3. In the blockchain-enabled interactive-service model, a higher data value Sreallocates demand toward the more credible node and improves upstream terms: $\frac{\partial D_1}{\partial S} < 0$, $\frac{\partial D_2}{\partial S} > 0$, $\frac{\partial W_1}{\partial S} > 0$, $\frac{\partial W_2}{\partial S} > 0$.

Moreover, both service nodes benefit in profit terms from higher data value, while the stronger node retains the more favorable position: , $\frac{\partial \pi_{N_1}}{\partial S} > 0$, $\frac{\partial \pi_{N_2}}{\partial S} > 0$, $\pi_{N_1} < \pi_{N_2}$.

Downstream pricing also becomes more selective than in the benchmark case. As data value increases, the stronger node's downstream price tends to decline, while the weaker node remains in a lower-price position and adjusts more modestly. Under digital enhancement, the platform uses data not simply to support higher prices, but to sharpen requirement matching and strengthen the attractiveness of trusted nodes.

4.3 Comparative insights across the two settings

A comparison of the two models yields three insights. First, data value plays different roles across the two settings. In the benchmark model, more valuable data expands demand for both options and improves the performance of all participants. In the enhanced model, more valuable data has a more selective effect: it strengthens demand reallocation toward the higher-credibility node while still improving upstream terms and profitability.

Second, interactive service and blockchain traceability alter the logic of platform competition. In the benchmark setting, competition is governed mainly by prices, credibility differences, and platform-held data value. In the enhanced setting, competition is shaped by a broader digital architecture in which requirement clarification, trusted records, and information visibility jointly influence user choice.

Third, blockchain traceability has an asymmetric governance effect. It does not improve the positions of all participants uniformly. Instead, it disproportionately strengthens the node with higher credibility and better verifiable records. From a CPI governance perspective, traceability should therefore be viewed as a strategic mechanism that reshapes demand allocation, bargaining positions, and platform incentives.

5 Managerial Insights

The analytical results yield three managerial insights. First, platform operators should treat data value as a strategic coordination resource. In the benchmark setting, higher data value improves demand, pricing, upstream terms, and profitability for all participants. Even without advanced digital enhancement, a Cyber-Physical Internet platform can therefore create substantial value through better data capture, feedback integration, and data-driven coordination.

Second, interactive service should be designed as a matching mechanism rather than merely as a communication function. In the enhanced setting, customer-interface support improves requirement clarification, reduces user uncertainty, and helps the platform allocate demand more effectively across heterogeneous nodes. Interactive service should therefore be evaluated by its contribution to matching accuracy, conversion efficiency, and coordination quality, not only by its operating cost.

Third, blockchain traceability should be deployed as a governance instrument with asymmetric effects. Stronger traceability disproportionately benefits the more credible node by making trusted information more visible and verifiable. Blockchain adoption does not simply improve transparency in a neutral way; it also reshapes competition and allocation within the network. Managers should therefore calibrate traceability investment to their governance objectives and network structure.

6 Conclusion

This paper develops an analytical framework to study pricing and governance in a Cyber-Physical Internet service network with heterogeneous service nodes. Two settings are examined: a benchmark platform model and a blockchain-enabled interactive-service model. The analysis focuses on how data value, interactive service, and blockchain traceability jointly influence demand allocation, pricing decisions, and participant profitability.

Three conclusions emerge. First, in the benchmark setting, data value functions as a coordination amplifier: it increases demand, supports stronger pricing and upstream terms, and improves profitability for all participants. Second, once interactive service and blockchain traceability are introduced jointly, competition becomes more selective. Digital enhancement no longer expands demand symmetrically; instead, it reallocates demand toward the more credible node while improving coordination efficiency. Third, blockchain traceability has an asymmetric governance effect. By making trusted information more visible and verifiable, it strengthens the position of the higher-credibility node and changes the allocation logic of the service network.

These findings contribute to the understanding of Cyber-Physical Internet platform systems by showing that pricing, interactive support, data utilization, and traceability design should be considered jointly rather than separately. Future research may consider multi-node competition, dynamic pricing, repeated interaction, and richer governance mechanisms such as node admission rules or data-sharing arrangements.

Acknowledgments

We thank the editors and the anonymous reviewers for their fruitful comments and suggestions in improving the quality of this paper. This work was supported by the National Natural Science Foundation of China (No. 52305557), Hong Kong Research Grants Council (No. 15203025, T32-707/22-N, C7076-22GF, R7036-22), Guangdong Basic and Applied Basic Research Foundation (No. 2024A1515011930), Research Institute for Advanced Manufacturing (RIAM) of The Hong Kong Polytechnic University (No. CDLU, CDLM, CDJX).

References

- Adomavicius G., A. Tuzhilin (2005): Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, v17, no6, 734-749, https://doi.org/10.1109/TKDE.2005.99.
- Agrawal T.K., V. Kumar, R. Pal, L. Wang, Y. Chen (2021): Blockchain-based framework for supply chain traceability: A case example of textile and clothing industry. *Computers & Industrial Engineering*, v154, 107130, https://doi.org/10.1016/j.cie.2021.107130.
- Ahmed W.A.H., B.L. MacCarthy (2023): Blockchain-enabled supply chain traceability – how wide? how deep? *International Journal of Production Economics*, v263, 108963, https://doi.org/10.1016/j.ijpe.2023.108963.
- Choi T.-M., S.W. Wallace, Y. Wang (2018): Big data analytics in operations management. *Production and Operations Management*, v27, no10, 1868-1883, https://doi.org/10.1111/poms.12838.
- Cui Y., M. Hu, J. Liu (2023): Value and design of traceability-driven blockchains. *Manufacturing & Service Operations Management*, v25, no3, 1099-1116, https://doi.org/10.1287/msom.2022.1161.

- den Boer A.V., N.B. Keskin (2022): Dynamic pricing with demand learning and reference effects. *Management Science*, v68, no10, 7112-7130, https://doi.org/10.1287/mnsc.2021.4234.
- Gawer A. (2022): Digital platforms and ecosystems: remarks on the dominant organizational forms of the digital age. *Innovation: Organization & Management*, v24, no1, 110-124, https://doi.org/10.1080/14479338.2021.1965888.
- Montreuil B. (2011): Toward a Physical Internet: Meeting the global logistics sustainability grand challenge. *Logistics Research*, v3, no2-3, 71-87, https://doi.org/10.1007/s12159-011-0045-x.
- Pan S., E. Ballot, G.Q. Huang, B. Montreuil (2017): Physical Internet and interconnected logistics services: Research and applications. *International Journal of Production Research*, v55, no9, 2603-2609, https://doi.org/10.1080/00207543.2017.1302620.
- Park D.H., H.K. Kim, I.Y. Choi, J.K. Kim (2012): A literature review and classification of recommender systems research. *Expert Systems with Applications*, v39, no11, 10059-10072, https://doi.org/10.1016/j.eswa.2012.02.038.
- Saberi S., M. Kouhizadeh, J. Sarkis, L. Shen (2019): Blockchain technology and its relationships to sustainable supply chain management. *International Journal of Production Research*, v57, no7, 2117-2135, https://doi.org/10.1080/00207543.2018.1533261.
- Taylor T.A. (2018): On-demand service platforms. *Manufacturing & Service Operations Management*, v20, no4, 704-720, https://doi.org/10.1287/msom.2017.0678.
- Wu H., L. Huang, M. Li, G.Q. Huang (2024): Cyber-Physical Internet (CPI)-enabled logistics infrastructure integration framework in the Greater Bay Area. *Advanced Engineering Informatics*, v60, 102551, https://doi.org/10.1016/j.aei.2024.102551.
- Wu H., M. Li, C. Yu, Z. Ouyang, K.H. Lai, Z. Zhao, S. Pan, S. Wang, R.Y. Zhong, Y.H. Kuo, F. Zhang, W. Huang, Z.J.M. Shen, E. Ballot, G.Q. Huang (2025): Towards Cyber-Physical Internet: A systematic review, fundamental model and future perspectives. *Transportation Research Part E: Logistics and Transportation Review*, v197, 104051, https://doi.org/10.1016/j.tre.2025.104051.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Enabling durable horizontal multilateral collaboration by integrating collaborative business modelling with multi-party optimization

Ruben Fransen[1], Tamara Oukes[1], Jaco van Meijeren[1], Marijn van Adrichem[1]
TNO, The Hague, The Netherlands
Corresponding author: ruben.fransen@tno.nl

Abstract: *This paper examines how durable horizontal collaboration among competing freight transport companies can be enabled by integrating collaborative business modelling (CBM) with multi-party optimization and contribute to the development of the Physical Internet (PI). Focusing on Dutch container road transport, the study addresses persistent barriers to collaboration, including lack of trust, governance challenges, unequal value distribution, and data-sharing concerns. A five-step CBM framework was developed and applied to a real-world use case involving the matching of import and export container flows to reduce empty kilometers. Impact assessment on real world data shows that collaboration increased the share of matched export containers from about 27% to over 52%, reducing trucking mileage by approximately 17,300 kilometers per month, even under time-window constraints. The findings demonstrate that combining organizational design with optimization methods can improve operational efficiency while creating the governance and trust conditions necessary for scalable, long-term horizontal collaboration.*

Keywords: *Horizontal Collaboration, Physical Internet, Collaborative Business Models, Decentral scheduling, Data Space, Secure Multi-Party Computation, Road transport, Container shipping*

Physical Internet (PI) Roadmap Fitness*: ☒ System of Logistics Networks, ☒ Horizontal Supply Chain Alignment, ☒ Logistics/Commercial Data Platform, ☒ Access and Adoption, ☒ Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☒ In-Person presentation*

1 Introduction

Freight transport companies are facing multiple, interconnected challenges. These include ambitious sustainability targets, increasing pressure on security of supply, decreasing availability of physical infrastructure and staff shortages. As a result, transport companies are increasingly looking beyond the boundaries of their own operations to find efficient transport alternatives and improve overall logistics system performance. This need is further intensified by the expected growth in freight volumes and upcoming legislation and policy measures such as the introduction of kilometer-based road pricing in the Netherlands mid-2026.

We see these developments as an opportunity to make a step in the development towards the Physical Internet (PI) and realize durable horizontal, multilateral collaboration in the transport sector as the urgency amongst transport companies to find and realize effective and feasible solutions is very high. Although horizontal collaboration in transport operations has been widely studied and tested in practice, successful implementations remain limited. This is largely due to the complexity of collaboration between competitors, where issues such as lack of trust and unequal distribution of benefits play a major role (Ahmadi et al., 2024; Durán-Micco et al., 2025; Karam, Hussein, et al., 2021). In addition to these organizational challenges, there is also

a significant technological challenge. Effective horizontal collaboration requires technological innovations that support automated, secure and fair decision-making between independent parties.

The overall objective is to enable asset sharing between competing companies by developing and combining organizational and technological innovations in a way that meets user requirements, builds the required trust, preserves companies autonomy and data sovereignty, improves logistics system performance and sustains the required long term collaboration. We see this as a step towards enabling the Physical Internet in which vertical and horizontal collaboration in logistics and transport operations are required.

1.1 Research program Future-Proof Smart Logistics

To deal with the challenge described in the previous section we make use of the high urgency experienced by transport companies and we develop and validate innovations that solve the well-known barriers for long lasting multilateral horizontal collaboration within the TNO program Future-Proof Smart Logistics. Until now transport companies most of the time had only limited interest in solutions for horizontal collaborations where at the moment transport companies indicate they need solutions for horizontal collaboration to stay in business given pressing targets and limited availability of means and people. In line with Physical Internet concept logistics companies are realizing more and more they have to do 'more with less'.

At the same time it is clear there are a number of well-known barriers for multilateral horizontal collaboration that have to be overcome to make sure these kind of collaborations will sustain over time. Therefore, the state-of-the-art established on a number of relevant topics will be advanced in the Future-Proof Smart Logistics program in the following way:

- Collaborative business modelling (CBM) for multilateral and scalable horizontal collaboration where literature focusses on vertical, or bilateral horizontal collaborations (Karam et al., 2021).
- Automatically configured data spaces for trusted, flexible, fast and easily accessible data sharing solutions where currently the configuration of data spaces takes lots of time, efforts and difficult discussions and leads to solutions in silos (EC, 2023).
- Data driven decentral algorithms for logistics planning of shared assets between multiple companies, where decentral algorithms are currently not used for optimization purposes or only for optimization within a single company. In addition, PETs (privacy enhancing technologies) for trusted yield distribution and proof of computation are needed as these have not been developed for these applications (TNO, 2022; Kaanich et al., 2020).
- System-of-systems engineering approach for system integration and alignment combining organizational and technological elements with many (partly) competing stakeholders, where current systems engineering focusses mainly on product development and value chains with a couple of largely dependent stakeholders (Muller, 2011).

1.1.1 Research contribution

In this paper we describe the first step in our research program, the development of a collaborative business model framework applied on a logistics use case to match import and

export flows in road freight transport to reduce the number of empty kilometers for a number of collaborating container transport companies. As part of this framework a quick scan analysis has been applied as an ex-ante impact assessment to determine the potential impact of this kind of horizontal collaboration for this logistics use case. We assumed the identification of a set of feasible, accepted and trusted agreements formalized in a mutual agreement supported by a quick scan analysis showing the potential impact would be crucial for the commitment of the involved transport companies to take next steps towards the implementation of a prototype solution in real-world logistics operation. Although this step is often skipped or only partly executed, our research program confirmed how important it is to take this step in full extent.

2 Research background

2.1 Collaborative business modelling

While individual transport companies have optimized their planning processes extensively, the potential for further efficiency gains to reduce costs and lowering CO_2 emissions within a single company is limited. To achieve additional improvements, the sector is exploring Horizontal Transport Collaboration (HTC) – "bundling transport of competitors, i.e., companies operating on the same level of the supply chain, who offer similar or complementary transportation services" (Vanovermeire et al., 2014). Research has shown that such collaboration can, in theory, improve economic performance, sustainability, operational efficiency, and strategic resilience (e.g., Durán-Micco et al., 2025; Grote et al., 2023; Zenezini et al., 2019). In practice, however, successful examples remain scarce, largely because collaboration among competitors introduces challenges such as trust deficits, hidden coordination costs and unequal benefit distribution (Ahmadi et al., 2024; Durán-Micco et al., 2025; Karam et al., 2021).

Collaborative business models are suggested to be an effective approach to overcome such trust issues, governance barriers and technology and information sharing barriers in HTCs (Vargas et al., 2020; Vargas et al., 2018). Collaborative business models are a structured framework that defines how collaborating transport companies co-create and share value, allocate benefits, costs and responsibilities, and agree on governance arrangements (Rohrbeck et al., 2013). As the literature on collaborative business models focuses on joint value creation and capture while integrating economic and sustainability objectives (Pedersen et al., 2021), it provides an appropriate lens to study how HTCs can overcome their challenges and realize their value creation potential. Specifically, collaborative business models can increase trust between actors by embedding relational governance structures, e.g. delimit information sharing, define decision rights and create joint monitoring systems, that reduce opportunism, knowledge leakage and strategic vulnerability (Audy et al., 2007; Karam, Hussein, et al., 2021; Pérez-Cervera et al., 2021). They can also reduce coordination costs by reframing the joint transport operations as a shared activity systems with standardized process, a joint planning platform and shared performance metrics (Karam, Reinau, et al., 2021; Vargas et al., 2020; Zenezini et al., 2019). Moreover, they can also avoid conflicts about unequal benefit distributions by designing transparent and fair value-sharing mechanisms to clarify who creates what value and who captures which share (Ahmadi et al., 2024; Aydemir et al., 2023; Basso et al., 2025; Durán-Micco et al., 2025).

Although collaborative business models have been proposed as a promising approach to strengthen trust and fairness in HTCs, their practical application remains limited as the focus of research has been on operational and technical aspects of HTCs (Vargas et al., 2020; Vargas

et al., 2018). As relational aspects are crucial for the success of technical solutions for HTCs, here we further expand our knowledge on collaborative business models for HTCs.

2.2 Logistics use case - horizontal collaboration to reduce inefficiencies

In this program a logistics use case is developed together with container transport companies to set up a horizontal collaboration to reduce the number of empty kilometers driven. The specific decision problem for which transport companies in this study aim to collaborate is the matching of import and export trips of container transports. This problem can be formulated as an assignment problem. In a broader range of transport planning problems it is a special case of a pick-up and delivery problem (PDP) where we (currently) focus on the matching of the deliveries and the pick-ups and do not yet take into account the scheduling or routing of vehicles which is part of PDPs. The concept of the container matching problem is explained in the figure below.

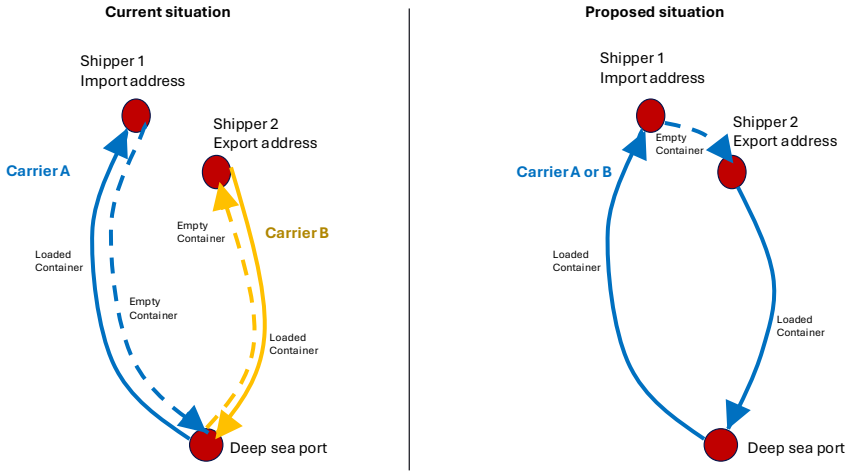


Figure 1 Example of how matching import and export container trips can reduce transport of empty containers and increase transport efficiency.

The matching problem can be quite efficiently solved by The Hungarian method (Kuhn, 1955). Leder et al. 2025 show that the method can even perform in the secure domain solving the problem whilst securely handling sensitive information (Leder, Fransen, Sijpesteijn, Veugen, & Phillipson, 2025).

3 Methodology and results

3.1 Collaborative business model framework

Based on the lessons learned from the literature and the illustrative examples as well as prior experience with collaborative business modelling in other projects, we developed an initial process for collaborative business modelling (CBM). This process was designed for both the horizontal transport collaboration (HTC) between transport companies and the broader value network surrounding them. Its purpose is to support the development and commercialization of the decentralized planning algorithm system.

The collaborative business modelling process evolves from committing the right partners to implementing the model in five steps, see Figure 2 below. In step 1, the right partners are engaged in the business modelling process, including setting a joint ambition and sketching the context. This is followed by step 2 where an ex-ante assessment is done of the impact for the

partners which is concluded with a intention to collaborate. Then in step 3 the actual design of the collaborative business model starts by defining and validating its content, structure and governance. Before this can be implemented often barriers have to be overcome and drivers leveraged, this is done in step 4. The final step is implementing the business model from pilot to demo to actual exploitation. Below each step is explained in detail. In this description the steps appear linear, but in practice this process will never be as sequential as described. Therefore, we also describe the back loops from later steps to the particular steps, describing the situation when it is necessary to reiterate to an earlier step.



Figure 2 Conceptual steps of CBM framework.

3.1.1 Application to the logistics use case

Scoping the collaboration

BoxReload, a previous collaboration initiative among Dutch container road transport companies, facilitated the exchange and matching of transport orders to combine two single-leg trips into one efficient backhaul in deep sea port hinterland. Although the concept showed promise, the initiative failed due to non-technical barriers: insufficient governance arrangements, limited willingness of planners to accept matches, and an unsustainable business case for the software provider.

Six former partners of this previous initiative, together with TNO, still see the value of collaborative planning; or even stronger, it is highly needed. They have therefore initiated the development of the Collaborative Fleet Management System (CFMS), focusing both on a solid technical foundation and on a viable collaborative business model addressing the shortcomings that led to Boxreload's collapse. In 2025, we began this process; below we outline the steps taken to date.

Step 1: Setting the collaborative foundation

TNO took the initiative to identify potential end-users for the CFMS and, through TLN (Transport en Logistiek Nederland), reconnected with former Boxreload partners who were willing to collaborate again on both the technical system and its associated collaborative business model. TLN and this initial partner subsequently involved five additional transport companies, forming a group of six companies committed to exploring renewed horizontal collaboration.

In June 2025, TNO organized a workshop with these transport companies to validate earlier survey results on their current planning challenges. The transport companies confirmed that sharing operational data offers clear gains, such as more efficient use of transport capacity, improved service reliability, higher profitability, and the ability to benefit from each other's networks and expertise. They also highlighted significant pains, including issues of trust, planner autonomy, commercial sensitivities, fair trip allocation, required system integrations, and the risk of negative customer perceptions.

Based on this shared understanding, the transport companies formulated a joint ambition: "use transport capacity more efficiently so that more benefits can be made on company profitability and sustainability". The workshop made clear that the group wished to further pursue and frame this collaborative opportunity.

Step 2: Framing the collaborative opportunity - ex-ante impact assessment on real-world data
To quantify the collaboration potential for these transport companies an ex-ante impact assessment has been executed that shows that the collaboration can increase the share of matched export containers from 27% to 52% which would lead to a reduction of 17,300 km of trucking mileage on a monthly basis. Detailed results are presented in paragraph 3.2.

To formalize their commitment and mark the transition to next steps, the partners prepared statement of intent. This document builds on insights from steps 1 and 2, including informal discussions, joint workshops, and additional individual interviews with five of the six transport companies. It outlines the joint ambition, the scope of collaboration, participation conditions, governance arrangements, and agreements on data, privacy, and KPI-based reporting.

Step 3: Designing the collaborative business model
The collaborative business model for the CFMS is structured around three dimensions: content, structure, and governance.
Content: The value proposition focuses on improving profitability and sustainability by reducing empty mileage through sharing and matching transport orders within the CFMS. Early impact assessments show that this value can already be realized with the current group of transport companies, and the potential will grow as more partners join. Today, value is shared through bilateral agreements, but the partners aim to introduce a system-wide yield distribution mechanism to allocate benefits fairly in the future.
Structure: Currently, the six transport companies jointly handle decision-making, supported by TNO in both technical development and collaboration design. As the collaboration grows, this structure will become unsuitable. Therefore, the partners and TNO are exploring scalable governance arrangements based on clear principles: contributing all shareable trips, providing essential planning data in a standardized format, participating in necessary meetings, maintaining valid licensing, and upholding individual quality and safety obligations.
Governance: The statement of intent defines the initial agreements that guide the collaboration, including ambition, scope, participation conditions, data and privacy arrangements, and KPI-based reporting. These agreements will be regularly evaluated and refined as the CFMS and the collaboration mature, ultimately forming a stable governance framework that new members can adopt upon joining.

Step 4: Building the collaborative roadmap
Transport companies are strongly motivated to improve efficiency due to shrinking margins, and their prior experience with Boxreload provides a basis for collaboration. At the same time, bottlenecks are expected, such as planners rejecting matched trips, concerns about exchanged service quality, and technical challenges in data sharing. These issues will be monitored, evaluated and solved throughout the pilot.

Step 5: Implementing the collaborative business model
The initial collaborative business model will be tested in a pilot using the first CFMS prototype in the remainder of 2026. This pilot will enable partners and TNO to refine both the collaborative business model and the CFMS through iterative, practice-based learning.

3.2 Ex-ante impact assessment potential impact

3.2.1 Approach

To quantify the reduction potential of horizontal collaboration by the partners an ex-ante impact assessment has been executed. Real-world data on container shipments has been collected and prepared for analysis for a full month in the recent past. The potential reduction of collaboration has been identified by comparing the difference between a scenario of current individual operations and a scenario of collaboration between the transport companies. In the individual scenario only import and export containers will be matched on individual company datasets, whilst in the collaborative scenario matches of imports and exports have been optimized on all import and export shipments in the collaboration. To find the best set of matches the Hungarian method has been applied.

The matching of import and export containers is constrained by a number of elements. Firstly the container type should fit for matching, only general purpose containers can be matched. Special containers such as reefers or tank containers cannot be directly transferred from import shippers to export shippers as cleaning or quality inspections are required. Secondly, all the container properties of the import and export container should match container size and container owner (deep sea carrier). Concerning container size there is an exemption rule that a high cube import container can be matched to a standard export container request, but not vice versa. Thirdly, geographically the import location and export should be relatively close to each other for a match to be efficient. Finally delivery (import) and pick-up (export) time windows should match with the driving time between the locations. The time windows are operationally critical because import and export locations have crew scheduled to (un)load the containers. In other words, in case of a container match the delivery of the container to the export client should not be too late (after the time window) and also not be too early that the truck driver has to spent too much time waiting.

These decision rules have been implemented in an analysis framework, such that the best possible matches in the set of real world data can be determined by the Hungarian method. The input for the Hungarian method is a *n x m* matrix with the $cost_{ij}$ of matching import container *i* to export container *j.* In our case base costs are the driving distance between client *i* and *j*, such that the Hungarian method finds a minimal distance matching set from the matrix. To include all matching constraints, we have created a separate matching matrix for each container size and container owner, the hard constraints. For the soft constraints, the time window constraints, penalty costs are introduced into the matching matrix to ensure no matches will be made that cannot make the time window or would incur too much waiting time.

3.2.2 Results

We show results of applying the Hungarian match-making algorithm for 4 different scenarios: firstly comparing individual matching with collaborative matching without time windows and secondly comparing the impact of time window constraints on the number of total matches and the reduction of number of kilometers driven. Main results are described below.

- The analysis was done on one month of container data of the 6 companies containing 4379 import containers and 1220 export containers.
- Collaboration leads to an additional 360 matches in the scenario without time window constraints and 354 matches in the scenario with time window constraints. Showing that 50%-59% of export containers find a suitable match in collaboration compared to only 21%-30% in case of individual operations.
- Saving up to 20,200 kilometers driven without time constraints and respectively 16,700 kilometers per month with time constraints.

- Adding time window constraints reduced the number of total matches from 724 to 605 (-16 %) and reduced the savings in kilometers driven by 17%. Which can be explained by less match-options due to time constraints leading to a lower quality overall matching.

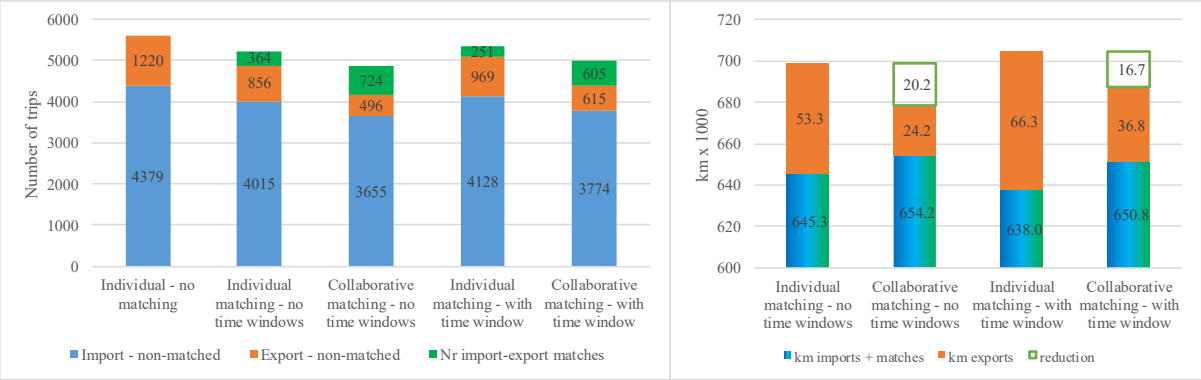


Figure 3 Results of ex-ante impact assessment comparing individual matching with collaborative matching on number matches (left) and savings in trucking mileage.(right)

4 Summary, conclusions and discussion

This paper addresses the persistent challenge of achieving durable horizontal collaboration among competing freight transport companies, despite its well-documented potential to improve efficiency and sustainability. Building on the Physical Internet vision, the study focuses on integrating collaborative business modelling (CBM) with multi-party optimization to enable scalable, trusted asset sharing while preserving company autonomy and data sovereignty. The paper presents a collaborative business model framework specifically designed for multilateral horizontal transport collaboration. This framework structures collaboration development into five iterative steps, ranging from partner commitment and ex-ante impact assessment to implementation and exploitation. The framework is applied to a logistics use case involving Dutch container road transport companies aiming to reduce empty kilometers by matching import and export container flows.

An ex-ante impact assessment was conducted using real operational data from five transport companies. By applying the Hungarian assignment algorithm across company boundaries, the collaborative scenario increased the share of matched export containers from approximately 25% to over 50%. This resulted in a reduction of around 17,000 trucking kilometers per month, even after incorporating strict time window and operational constraints. Beyond quantitative efficiency gains, the study demonstrates how CBM helps address non-technical barriers such as trust, governance, planner autonomy, and fair value distribution. The results indicate that combining business model design with decentralized optimization is essential for making horizontal collaboration both feasible and sustainable in practice. The paper concludes that CBM is a critical enabler for transitioning from ad-hoc collaboration pilots focusing on specific parts of a collaboration with limited attention for agreements, organization and trust to accepted and trusted and therefore durable and scalable horizontal collaborations.

The involved transport companies in the logistics use case have signed a letter of intent to accept and formalize the agreements on the horizontal collaboration developed as part of the CBM process and are fully committed to collaborate on next steps towards a prototype, also partly based on insights of the potential impact from the ex-ante impact assessment showing substantial benefits for the transport companies.

4.1 Discussion

The findings of this study confirm that trust remains the central limiting factor in scaling horizontal collaboration beyond small, well-known groups of transport companies. Importantly, trust cannot be treated as a soft, relational precondition alone, but must be operationalized and embedded explicitly in the collaborative business model. CBM provides a suitable mechanism to do so, as it allows trust to be translated into concrete agreements on data usage, decision rights, value sharing, and governance.

As horizontal collaborations grow, CBM must shift from assumed trust to institutionalized trust mechanisms. This means establishing clear, enforceable rules for data sharing, specifying purpose, conditions, and guarantees for confidentiality. Agreements on how to collaborate are equally essential, ensuring both maximum benefit and sustaining long lasting collaborations. In addition, trust depends strongly on perceived fairness. As collaboration scales, bilateral arrangements and informal compensation mechanisms become insufficient. CBM must therefore evolve towards transparent, auditable and accepted value- and yield-distribution mechanisms that make benefits traceable and defensible for all participants. This is especially important in competitive environments where even small perceived imbalances can quickly erode commitment. CBM thus acts as the bridge between algorithmic optimization outcomes and their acceptance by human decision-makers.

The results from the initial impact assessment show that transport companies need clear evidence of expected benefits before they are willing to invest their time, money, and resources in further development. This analysis not only highlights potential advantages based on agreed business rules but also suggests ways to boost those benefits. In the logistics case studied, it became apparent that including more participants—especially those with a larger number of export trips—would lead to greater overall gains and help balance import and export flows. Another important area for future investigation is how the timing requirements set by shippers, whether very flexible or very strict, can affect the potential benefits. Discussing these requirements with shippers may open up new possibilities for improvement. These topics should be explored further as the next steps toward applying these ideas in real-world logistics operations.

References

- Ahmadi, T., Veen, J. A. A. van der, Venugopal, V., & Kamran, M. A. (2024). Conditions for viable horizontal collaborative transport: Insights from a stylised model. *International Journal of Industrial and Systems Engineering*, *47*(5), 1–35. https://doi.org/10.1504/IJISE.2024.139945
- Audy, J.-F., D'Amours, S., & Rönnqvist, M. (2007). Business Models for Collaborative Planning in Transportation: An Application to Wood Products. Establishing the Foundation of Collaborative Networks, 667–676. https://doi.org/10.1007/978-0-387-73798-0_72
- Aydemir, I., McLeod, F., Grote, M., & Cherrett, T. (2023). Evaluating the Feasibility of a Shared-Fleet Operation in Healthcare Logistics between Public Organisations. Sustainability, 15(21), Article 21. https://doi.org/10.3390/su152115361
- Basso, F., Monardes-Concha, C. A., Lorca, F., Pezoa, R., & Varas, M. (2025). A horizontal collaboration approach for grape transportation in a pisco cooperative. Computers & Industrial Engineering, 200, 110797. https://doi.org/10.1016/j.cie.2024.110797
- EC (2023). European data spaces and the role of data.europe.eu, European Commission.

- Durán-Micco, J., Alaei ,Shafagh, & and Macharis, C. (2025). Evaluating the impact of horizontal collaboration in logistics systems: A simulation-based study. *Transportmetrica B: Transport Dynamics*, *13*(1), 2475204. https://doi.org/10.1080/21680566.2025.2475204
- Grote, M., Cherrett, T., Oakey, A., Martinez-Sykora, A., & Aydemir, I. (2023). Benefits of Shared-Fleet Horizontal Logistics Collaborations: A Case Study of Patient Service Vehicles Collecting Pathology Samples in a Public Sector Healthcare Setting. Future Transportation, 3(1), Article 1. https://doi.org/10.3390/futuretransp3010011
- Kaanich, N., Laurent, M., Belguith, S. (2020). Privacy enhancing technologies for solving the privacy-personalization paradox: Taxonomy and survey, Journal of Network and Computer Applications
- Karam, A., Reinau, K. H., & Østergaard, C. R. (2021). Horizontal collaboration in the freight transport sector: Barrier and decision-making frameworks. European Transport Research Review, 13(1), 53. https://doi.org/10.1186/s12544-021-00512-3
- Karam, A., Hussein, M., & Reinau, K. H. (2021). Analysis of the barriers to implementing horizontal collaborative transport using a hybrid fuzzy Delphi-AHP approach. *Journal of Cleaner Production*, *321*, 128943. https://doi.org/10.1016/j.jclepro.2021.128943
- Kuhn, H. (1955). The hungarian method for the assignment problem. Naval Research Logistics Quarterly.
- Leder, S., Fransen, R., Sijpesteijn, T., Veugen, T., & Phillipson, F. (2025). Secure Collaborative Pickup and Delivery Problem using MPC. International Conference on Computational Logistics. Delft: Springer.Muller, G.J. (2011). Architectural reasoning explained, Embedded Systems Institute.
- Pedersen, E. R. G., Lüdeke-Freund, F., Henriques, I., & Seitanidi, M. M. (2021). Toward Collaborative Cross-Sector Business Models for Sustainability. Business and Society, 60(5), 1039–1058. Scopus. https://doi.org/10.1177/0007650320959027
- Pérez-Cervera, C., Lu, M., Sánchez-Pérez, A., Sáez-Carramolino, L., & Furió-Pruñonosa, S. (2021). Collaborative Service-Dominant Business Model Design for a Just-In-Time rail shuttle service at the port of Valencia. International Journal of Transport Development and Integration, 5(1), 28–40. https://doi.org/10.2495/TDI-V5-N1-28-40
- Rohrbeck, R., Konnertz, L., & Knab, S. (2013). Collaborative business modelling for systemic and sustainability innovations. International Journal of Technology Management, 63(1–2), 4–23. https://doi.org/10.1504/IJTM.2013.055577
- TNO (2022). Talking Trucks: Decentralized Collaborative Multi-Agent Order Scheduling for Self-Organizing Logistics.
- Vanovermeire, C., Sörensen, K., Van Breedam, A., Vannieuwenhuyse, B., & Verstrepen, S. (2014). Horizontal logistics collaboration: Decreasing costs through flexibility and an adequate cost allocation strategy. International Journal of Logistics Research and Applications, 17(4), 339–355. https://doi.org/10.1080/13675567.2013.865719
- Vargas, A., Fuster, C., & Corne, D. (2020). Towards Sustainable Collaborative Logistics Using Specialist Planning Algorithms and a Gain-Sharing Business Model: A UK Case Study. Sustainability, 12(16), Article 16. https://doi.org/10.3390/su12166627
- Zenezini, G., Gonzalez-Feliu, J., Mangano, G., & Palacios-Arguello, L. (2019). A Business Model Assessment and Evaluation Framework for City Logistics Collaborative Strategic Decision Support. Collaborative Networks and Digital Transformation, 552–561. https://doi.org/10.1007/978-3-030-28464-0_48

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE
IMT Mines Albi École Mines-Télécom

Physics of Decision: A Data-Driven Discovery of Physical Internet Dynamics

Benoit Morvan[1] and Nafe Moradkhani[2], Guillaume Martin[1], Victor Romero[1], Benoit Montreuil[3], Frederick Benaben[1,3]
1. CGI, IMT, Albi, France
2. CESIT, Kedge Business School, Talence, France
3. ISyE, Georgia Institute of Technology, Atlanta, USA
Corresponding author: benoit.morvan@mines-albi.fr

Abstract: *This study falls within the Physics of Decision (POD) framework to enable, through a data-driven approach, the modeling of risk and opportunity scenarios. This study aims to recover the trajectory equation of a discrete simulation engine to model the represented system using a continuous dynamic equation. We opt for a continuous dynamic equation to reduce the computation time of a simulator. To uncover the dynamic law governing our system, we will use a data-driven algorithm, which, from a library of nonlinear terms, recovers a governing differential equation for the system. To enable the learning of an entire network, we will use our node-based data-driven engine rather than applying it to the entire network as a whole. This atomization approach isolates global laws applicable to each node, trains on a network, and remains generalizable to new configurations.*

Keywords: *Physics of Decision, SINDy, Surrogate Model, Resilience, System Dynamics*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:☐ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),☐ Transportation Equipment, ☐ PI Networks, ☐ System of Logistics Networks, ☐Vertical Supply Consolidation, ☐Horizontal Supply Chain Alignment, ☐Logistics/Commercial Data Platform, ☐Access and Adoption, ☐Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐Flash Video, ☐In-Person presentation*

1 Introduction

1.1 Motivation

Climate events, geopolitical crises, and pandemics have all highlighted the critical need for decision-makers to shift toward highly resilient supply chains. This need is particularly relevant in the context of the Physical Internet (PI) paradigm. While the hyper-connected, modular architecture of the PI drives unprecedented efficiency, it also creates complex network topologies, where localized instabilities can rapidly propagate system-wide. Despite this urgency, industry report indicates that 92% of companies will still lack full supply chain resilience by 2026 (Gartner, 2026), risking massive operational and financial losses.
To address this structural vulnerability, Simchi-Levi (2022) emphasized that *'We Need A Stress Test for Critical Supply Chains"*. Building on this, Dolgui et al. (2025) defines a stress test as *'a resilience analysis through a computer or social simulation designed to determine the supply chain ability to withstand disruption and recover in different crisis scenarios'*.
However, executing network-wide stress remains a major computational challenge (Ben Rabia et al., 2023). Classical simulation models, such as Discrete Event Simulation (DES) or Agent-Based Modeling (ABM), are often data-heavy, hard to maintain and computationally expensive

for real-time monitoring. They require significant expertise to operate, resulting in limited flexibility for agile decision-making during crises.
To overcome, the limitation of classical discrete modeling, data-driven system identification offers a promising alternative (Duan et al., 2019). This paper proposes a methodology to discover the underlying dynamic laws governing PI networks directly from operational data. By capturing the system's governing mechanisms through continuous differential equations, we aim to provide a fast, scalable surrogate model capable of simulating new disruption scenarios and testing proactive resilience strategies.

1.2 Physics of Decision as a decision-centered problem statement

The Physics of Decision (POD) provides a conceptual framework to define the evolution of a system and quantify systemic instability. Under the POD paradigm, a managed system is viewed as a physical object navigating within a multidimensional space. The system's evolution is depicted as a continuous trajectory, where risks and opportunities act as physical forces altering is path (Benaben et al., 2022).
To formalize this, POD distinguishes between two environments: the *Description space* and the *Performance space* (Figure 1). The *Description space* encapsulates the defining variables that characterize the system and its operational parameter, including deliberate decision levers (e.g., nominal capacities, process efficiency) and exogenous conditions (e.g., uncertain demand, hazard intensities). A specific point in this space represents the system's context. Concurrently, the *Performance space* maps the system's operational outcomes and Key Performance Indicators (KPIs) at any given time.

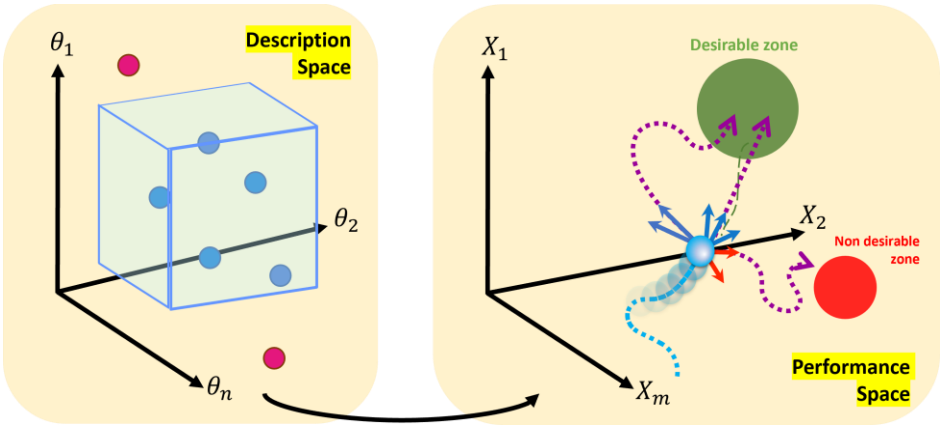


Figure 1: System in Description space and Performance space

This paradigm assumes the managed system can be represented by a continuous-time space formulation:

$$\frac{d\boldsymbol{x}}{dt} = F\big(\boldsymbol{x}(t), \boldsymbol{\theta}(t), \boldsymbol{u}(t)\big), \quad \boldsymbol{y}(t) = C\big(\boldsymbol{x}(t)\big)$$

where $\boldsymbol{x}(t)$ is the internal state vector (e.g., inventory levels), $\boldsymbol{\theta}(t)$ represents the system parameters within the *Description space* and $\boldsymbol{u}(t)$ defines the exogenous inputs within the *Description space*. *F* is unknown nonlinear evolution dictating the system's dynamics, and *C* maps the internal spaces to the performance metrics $\boldsymbol{y}(t)$. In this formulation, a *force* is interpreted as any shift in the system parameters $\boldsymbol{\theta}(t)$ or input $\boldsymbol{u}(t)$ that alters the future trajectory of the performance $\boldsymbol{y}(t)$.
The central problem is that classical simulation models to not provide a fast, mathematically explicit, and interpretable link between decision levers (the context $\boldsymbol{\theta}$) and performance trajectories. As a result, the strategic design phase often remains opaque and non-interactive for

decision-makers. The remainder of this paper addresses these issues by constructing an atomized, continuous model of a logistic network, serving as a rapid surrogate engine for resilience stress-testing.

1.3 Contribution

To bridge the gap between static resilience analysis and dynamic, real-time decision support in Physical Internet networks, this paper introduces a data-driven system identification methodology grounded in the Physics of Decision (POD) framework. Our primary objective is to discover the underlying dynamic laws, acting as force fields, that govern the logistical network's trajectory. More specifically, our contribution can be summarized in three points:

1. **Atomized Continuous Formulation for Topological Invariance:** We shift from macroscopic network-level modeling to an atomized approach, where each PI-nodes is modeled as an autonomous agent governed by local, continuous ordinary differential equations (ODEs). This decentralized formulation provides a strict topological invariance. It allows the dynamic laws to be learned on small-scale training networks and directly generalized to large-scale, complex network topologies without requiring retraining.
2. **Physics-Informed SINDy discovery:** Standard system identification relies on generic polynomial expansions that fail to capture the hard constraints of logistics networks. We propose a Parametric Physics-Informed adaption of the SINDy algorithm. By embedding continuous logistical operators, such as smooth minimum functions, into the candidate library, the algorithm successfully discovers the true mathematical link between the operational parameters ($\boldsymbol{\theta}$) and the system state ($\boldsymbol{x}$).
3. **Resilience Stress-Testing Engine:** By successfully abstracting microscopic discrete events into continuous governing equations, we provide a fast-to-compute surrogate model. This model acts as a decision-support engine, empowering managers to stimulate novel perturbation scenarios (e.g., sudden capacity loss or demand spikes) and proactively evaluate the network's resilience within the POD performance space.

2 State of the Art

2.1 Data-Driven Modeling and Surrogate Models in Logistics

Modeling complex, nonlinear, and highly interconnected networks present a significant computational challenge. Simulators accurately capture local behaviors but become computationally intractable for real-time decision-making due to the curse of dimensionality. To bypass this limitation, surrogate models act as lightweight mathematical approximations. By abstracting the discrete loading and routing events into aggregated continuous flows, the network's behavior can be modeled using a set of ordinary differential equations (ODEs), enabling rapid bifurcation analysis and scenarios exploration at a fraction of the computational cost.

To learn these continuous dynamics from data, several methodological paradigms have recently emerged, each offering different trade-offs between expressivity, interpretability, and computational efficiency.

Deep learning approaches, such as Neural ODEs and Graph Neural Networks, offer immense expressive power for handling high-dimensional, nonlinear, and spatio-temporal logistical data. Deep Learning has emerged as a powerful alternative particularly for handling high-dimensional, nonlinear, and data-rich environment (Richey Jr. et al., 2023). While they can fit

complex, nonstationary dynamics where the underlying physics are unknown, these neural surrogates typically operate as black box. To mitigate this opacity, hybrid architecture like Interpretable Neural System Dynamics (INSD) attempt to combine concept learning with causal graphs (D'Elia et al., 2025).

Sparse Identification of Nonlinear Dynamics (SINDy) has emerged as a powerful method for discovering governing equations directly from observational datasets (Brunton et al., 2016). By assuming that the true dynamics are driven by only a few fundamental terms, SINDy selects a sparse set of active functions from a predefined candidate library using regularized regression. This approach has been applied to socio-technical systems, such as intertwined supply chains (El Harraki et al., 2024) or traffic network (Wei et al., 2025). While SINDy provides an interpretable surrogate model, its integration with Model Predictive Control (MPC) further unlocks prescriptive capabilities, allowing dynamic reallocation of network capacity in response to disruption (Lu et al., 2025). However, standard SINDy faces significant limitations regarding scalability and library design. Without prior domain knowledge to prune the candidate library, the algorithm is prone to a combinatorial explosion and a higher risk of false discoveries.

This specific limitation motivates methodological contribution of our paper, by integrating a Physics-Informed structure into the SINDy library, we embed explicit logistical constraints (e.g., capacity saturation) to guide the regression.

2.2 System Learning in logistic network for resilience purpose

Recent literature emphasizes the critical role of stress tests in evaluating supply chain resilience (Ivanov et al., 2025). To operationalize these stress tests, various artificial intelligence techniques have been increasingly deployed to identify hidden supplier dependencies and predict network perturbations using historical datasets. For example, Kosasih et al. (2022) utilize Graph Neural Networks (GNN) to predict hidden links within supply chain topologies. Similarly, to address the challenge of evaluating the propagation of catastrophic disruptions when historical data is scarce, Liu et al. (2021) employ dynamic Bayesian networks combined with probability intervals and simulated annealing.

Parallel to probabilistic approaches, another stream of research leverages continuous dynamical systems to model network interactions. Cuong et al. (2022) applied deterministic Lotka-Volterra systems to describe cooperative interactions between ports, where the learned system acts as a controller through estimated parameters. The same systems as been adapted to data-driven supply chain contexts using SINDy to manage sourcing decisions under ripple effects (El Harraki et al., 2024).

While these approaches successfully identify structural vulnerabilities or model macro-level stability, they often struggle to provide a fast, high-fidelity, node-level simulation required for dynamic, real-time stressing.

2.3 Physics of Decision: A Force-Driven Framework for Dynamic Resilience

Traditional supply chain risk management often relies on static resilience metrics, such as Time-To-Survive (TTS) and Time-to-Recover (Ivanov et al., 2025), or evaluates disruptions through discrete probabilistic scenarios. To transition toward dynamic risk management, the Physics of Decision (POD) framework offers a fundamental paradigm shift. POD proposes a conceptual architecture where managing a collaborative system consists of controlling its continuous evolution, represented as a trajectory, withing a multidimensional *Performance space* (Moradkhani et al., 2022; Moradkhani et al., 2024).

Within this framework, uncertainties are conceptualized as physical forces acting upon the system. Detrimental forces (risks, such as localized capacity drop or sudden demand spike) push the system's trajectory away from its target zone. Conversely, beneficial forces (opportunities, such as the deployment of mitigation policies or capacity sharing) act as corrective forces to steer the trajectory back into an acceptable state (Morvan et al., 2025). Consequently, resilience is redefined dynamically: it is not longer just a post-event recovery metric, but the system's active capacity to navigate its *Performance space* and maintain its trajectory despite exogenous perturbing forces.

While POD provides a robust conceptual architecture, establishing a quantitative, data-driven link between the conceptual decision variables and the resulting trajectory has remained a major computational challenge. This is precisely where our application of the SINDy algorithm bridges the gap.

3 Problem formulation

3.1 Atomized Supply Chain Formulation

Conventional network-level modeling frameworks frequently encounter the curse of dimensionality, particularly in highly heterogenous, dynamic and decentralized systems such as the Physical Internet (PI). To ensure scalability and flexibility, we introduce a supply chain atomization paradigm. Rather than conceptualizing the logistics network as a centralized, monolithic entity, our approach decomposes the system into autonomous, self-governing entities, wherein each node *i* operates as an independent agent constrained solely by its localized physical governing law, denoted as f_i (Figure 2).

This formulation endows the logistic network with topological invariance. By isolating the intrinsic laws of each node, the mathematical model can be trained on a small network and subsequently generalized to a large-scale network without retraining. We assume that the local dynamics of any node *i* can be represented as a continuous-time ordinary differential equation (ODE):

$$\frac{d\boldsymbol{x}_i}{dt} = f_i(\boldsymbol{x}_i, \boldsymbol{\theta}_i, \boldsymbol{u}_i)$$

where $\boldsymbol{x}_i$ is the internal state of the node *i* (e.g., inventory level), and $\boldsymbol{\theta}_i$ represents the operational parameters (e.g., capacity, efficiency). The exogenous input vector $\boldsymbol{u}_i$ encompasses the incoming flows and outcoming flows.

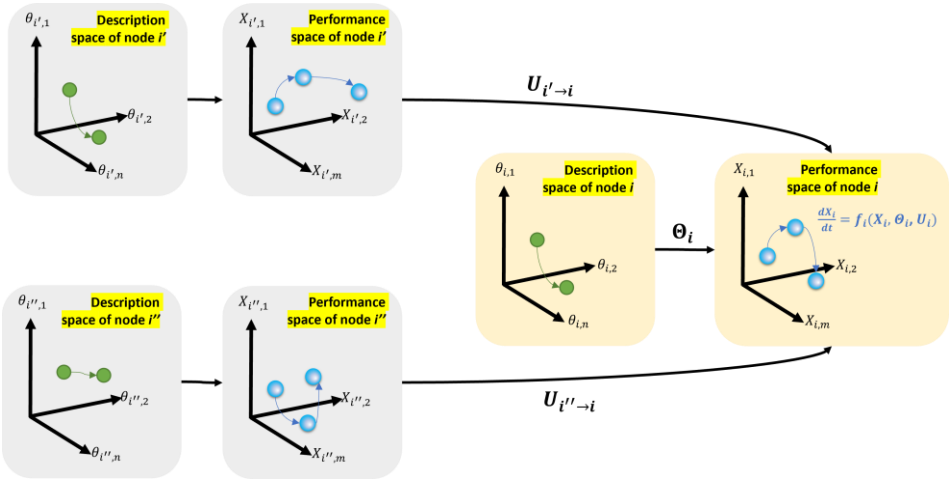


Figure 2: Atomization of the logistic network

3.2 Sparse Identification of Nonlinear Dynamics (SINDy)

Given a dataset of state of derivatives $\frac{dx}{dt}$ and corresponding states, SINDy (Brunton et al., 2016) construct a comprehensive library matrix $\Phi(X)$ containing candidate nonlinear functions (e.g., polynomials, trigonometric terms). The system evolution is then approximated as:

$$\frac{dX}{dt} \approx \Phi(\mathrm{X})\ \Xi$$

where Ξ is a matrix coefficients. To ensure the model remains interpretable and physically plausible. SINDy utilizes sparse regression algorithms (such as Sequentially Thresholded Least Squares (STLS) or Lasso) to force the coefficients of inactive terms to zero. The remaining non-zero elements of Ξ reveal the explicit mathematical equations governing the system.

3.3 Physics-Informed Library Construction

Standard SINDy algorithms rely on polynomial or trigonometric libraries to approximate $f(X)$. However, logistic dynamics are fundamentally governed by limits, such as bounded capacities, finite demand, and discrete routing policies. Approximating these sharp logistical bounds via high-degree polynomials often leads to overfitting, numerical instability, and a complete loss of physical interpretability.

To overcome this limitation and accurately project the system's causal dynamics into the Performance space, we introduce a Physics-Informed SINDy approach. The core principle is to guide the learning of the system's trajectory by embedding logistical domain knowledge directly into the candidate library. For a given node *i*, let $\boldsymbol{x}_i$ be its continuous state, $\boldsymbol{\theta}_i$ its operational parameters (e.g., nominal capacity, process efficiency), and $\boldsymbol{u}_i$ the exogenous inputs (e.g., upstream flows, downstream demand).

Rather than relying on naïve mathematical expansions, we generate physically plausible hypotheses by defining continuous operators that reflects known logistical policies, denoted as γ. These policies translate discrete managerial rules into continuous, differentiable functions $F_{\mathrm{i},\gamma}(\boldsymbol{x}_i, \boldsymbol{\theta}_i, \boldsymbol{u}_i)$. For instance, a transition function can be formulated to represent the bounding between an internal production potential and an external market demand. To ensure the mathematical problem remains tractable for the optimization solver, these sharp discrete boundaries are implemented using continuous relaxation (e.g., smooth minimum or maximum operators).

The complete library

$$\Phi_i = \left[1, \boldsymbol{u}_i, \boldsymbol{x}_i, \dots, F_{\mathrm{i},\gamma}(\boldsymbol{x}_i, \boldsymbol{\theta}_i, \boldsymbol{u}_i), \dots\right]$$

By applying sparse regression, the algorithm effectively evaluates competing physical hypotheses. The model naturally isolates the true governing dynamic law driving the network's performance trajectory, while forcing the coefficients of spurious or non-physical combinations to zero. This ensures that the discovered equations remain interpretable and explicitly linked to the causal mechanics of the Physical Internet.

4 Experimental results and validation

4.1 Use Case: PI-Network

To validate our data-driven engine, we must evaluate its capability to discover continuous governing laws from realistic, discrete operations. To achieve this, we developed a discrete-time simulator to gather operational data and assess our data-driven engine's ability to recover underlying continuous laws. In the Physical Internet (PI), the flow of goods in fundamentally

driven by the discrete routing of PI-containers between autonomous PI-hubs. Therefore, the Agent-Based Model (ABM) simulates these PI-containers as discrete entities.
The simulated logistics network is conceptualized as a direct graph composed of three distinct types of nodes:

- **Sources Nodes:** These act as the entry points of the network. They generate an exogenous inflow of PI-containers based on a time-varying production rate (P_i).
- **Transit PI-hubs:** These represent the core infrastructure of the Physical Internet. They act as intermediate process, and forward containers to downstream neighbors.
- **Clients Nodes (Sinks):** These represent the final destinations or market sinks. They pull inventory from the network based on a specific, exogenous downstream demand (D_i).

The internal transitions within a PI-hubs are constrained by localized operational rules. The logic governing the movement of stock is based on a strict capacity-constrained flow mechanism. At each time step, a PI-hub attempts to clear its current inventory (x_i) by pushing it to downstream nodes. However, this theoretical outflow potential is strictly bounded by node's operational parameters, specifically its nominal processing capacity (C_i) and its current operational efficiency (η_i).
Consequently, the routing logic operates at the intersection of a push and pull regime. A hub pushed inventory up to its maximum throughput capability, but the actual volume successfully transferred is ultimately limited by the available capacity or the explicit demand of the receiving downstream nodes. If the available stock exceeds the processing capacity, or if downstream nodes are saturated/ disrupted, the PI-hubs accumulate locally, so with an accumulation of the state variable x_i.

4.2 Experimental Protocol and Model Instantiation

Data generation. To train and validate the surrogate model, we generated trajectories using the discrete ABM simulator with a fixed step $\Delta T = 1$ day over a horizon of $T = 80$ days. To ensure a comprehensive exploration of the Description and Performance spaces, we generated N_{net} = 100 distinct random network scenarios of $n = 10$ nodes each. For each node i, the operational parameters $\theta_i = [\eta_i, C_i]$ (efficiency and nominal capacity) were sampled as constant piecewise-constant step-functions, while exogenous inputs $u_i = [P_i, D_i]$ (production and downstream demand) vary dynamically. To simulate realistic measurement uncertainty inherent to discrete counting processes, a Poisson noise term with intensity of 10% is applied to the observed stock trajectories. Initially inventory levels were set to zero ($x_{i,0} = 0$) during training. The parameter sampling ranges are summarized in Table 1.

Table 1: Parameter sampling ranges (per day)

Symbol	Description	Range
η_i	Efficiency of node *i*	[0.0 ; 1.2]
C_i	Capacity of node *i*	[2.0 ; 150.0]
D_i	Demands of node *i*	[0.0 ; 40.0]
P_i	Production of node *i*	[0.0 ; 500.0]

Physics-Informed Library Instantiation. Following the methodology in Section 3.3, we instantiate three candidate libraries of increasing physical specificity to empirically validate our library design choices.

The baseline library L_0 relies solely on polynomial expansions of the state variable and system context: $\phi_0 = [\, P_i, Inflow_i, x_i, x_i^2, x_i^3\ x_i C_i]$. This formulation makes no assumption about the governing physics.

The partial physics library L_1 introduces a first logistical operator encoding the capacity-constrained push regime. PI-hubs attempt to clear their inventory up to their maximum throughput. To ensure solver differentiability, we replace the discrete min operator with a smooth approximation $SM(a,b) = \frac{1}{2}\left(a + b - \sqrt{(a-b)^2 + \varepsilon}\right)$ where $\varepsilon \ll 1$ acts as a smoothing parameter. These yields: $\phi_1 = [1, P_i, Inflow_i, x_i, SM(x_i, \eta_i * C_i)]$. However, L_1 ignores the pull regime, the actual outflow is not only bounded by internal capacity but also by the downstream demand D_i.

The Physics-Informed library L_2 addresses this limitation by embedding the complete capacity-demand interaction as a nested smooth-minimum operator: $\phi_2 = [1, P_i, Inflow_i, x_i, x_i^2, x_i C_i, SM(SM(x_i, \eta_i * C_i), D_i)]$, encoding the full push and pull decision policy in a single differentiable function. The additional polynomial terms x_i^2 and $x_i C_i$ are included as distractors candidates to verify that the algorithm correctly rejects non-physical combinations in favor of the governing operator.

4.3 Trajectory Reconstruction

To evaluate the fidelity of the discovered governing equations, we conduct two complementary experiments. First, we assess the ability of three candidate libraries: L_0 (polynomial baseline), L_1 (partial physics), and L_2 (Physics-Informed), to reconstruct inventory trajectories across all node types. Second, we evaluate the surrogate's robustness under an unseen disruption scenario. Figure 3 illustrates the inventory trajectories of 10 nodes across four layers of a 60-nodes PI-network (10 sources, 20 hubs of layer 1, 15 hubs of layer 2, 15 clients), trained on small-scale random networks of 10 nodes.

Figure 3 demonstrates the accuracy of our data-driven engine. The graph illustrates the inventory trajectories of 10 distinct nodes, encompassing pure push transit hubs, sources nodes, and pull-driven client nodes, within a 60 nodes network topology. The Physics-Informed library L_2 achieves a sMAPE of 12% against the discrete ABM reference, compared to 41% for L_1 and 173% for L_0. The explicit encoding of the push/pull transition operator $SM(SM(x_i, \eta_i * C_i), D_i)$ is the physical ingredient. Crucially, this generalization is obtained without retraining, the governing law learned on 10-node random topologies transfers directly to the 60-node hierarchical PI network, validating the topological invariance of our atomized formulation.

To evaluate, the surrogate's capability as a resilience decision-support engine, we inject a sudden capacity loss of 90% on the node 10 between $t = 20$ and $t = 40$ days. As shown in Figure 3, the surrogate with the library L_2 accurately captures the shock waves propagating through the network.

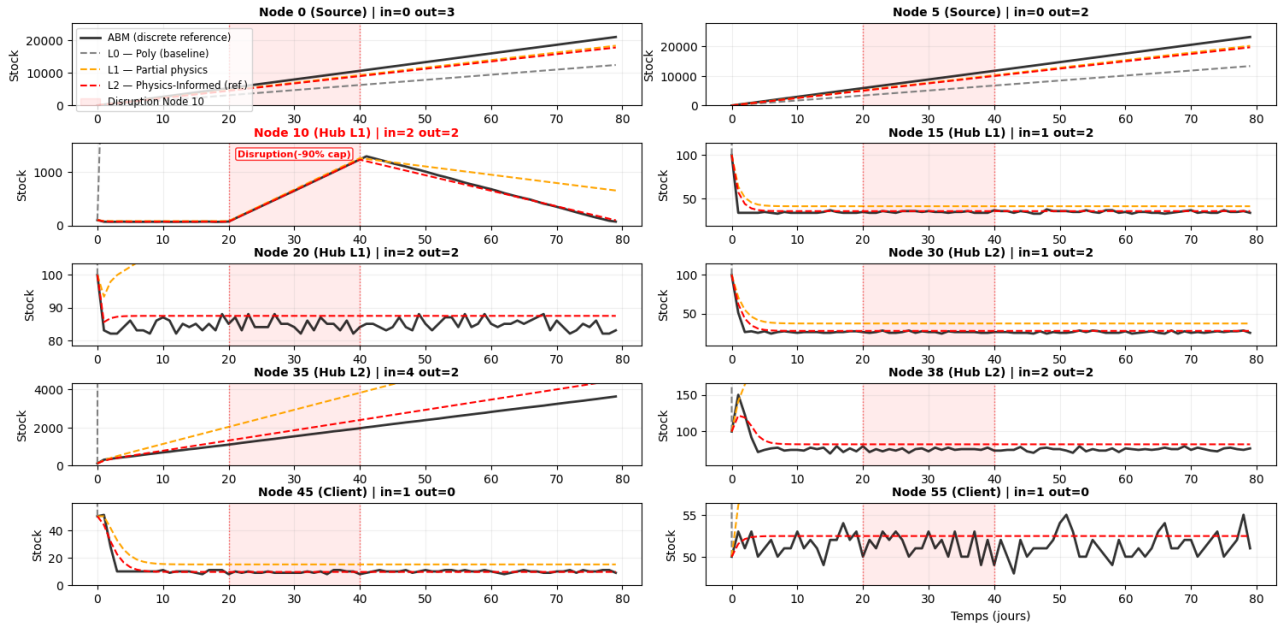


Figure 3: Bridging Micro-Stochasticity and Macro Dynamics

5 Conclusion and discussion

Physical Internet (PI) networks, driven by their hyper-connected and modular architecture, are inherently susceptible to the rapid propagation and amplification of localized disruptions. Consequently, decision-makers require robust tools to evaluate the network-wide scale and impact of such perturbations. While traditional simulation engines accurately replicate system mechanics, they are notoriously slow to implement, lack of flexibility, and are computationally prohibitive when evaluating thousands of risk scenarios.

To bridge this computational gap, we introduced a data-driven engine, inspired by SINDy algorithm. Grounded in the Physics of Decision (POD) framework, our approach conceptualizes systemic performance as a continuous trajectory within a multidimensional space, perturbed by exogeneous risks and opportunities. From this premise, we aimed to recover the governing differential equations of the system directly from operational data. By atomizing the logistic network into node-level ODEs, we subsequently developed a Physics-Informed SINDy to learn these localized continuous laws from a discrete Agent-Based Model (ABM). Integrating custom functions into the candidate library enabled us to explicitly enforce strict logistical bounds, such as nodal capacity constraints. The final section clearly demonstrates the effectiveness of SINDy as a tool for learning differential equations associated with a generic network. We have been able to demonstrate its usefulness in generating perturbation scenarios. This surrogate model presents inherent mathematical limitations. One limitation of the algorithm is its sensitivity to noise resulting from the integration of the differential equation. Several studies have attempted to adapt the algorithm to reduce this sensitivity. Then, the approach is highly dependent on the manual formulation of the candidate library. Sparse regression is extremely sensitive to the terms included, collinearity among candidates' functions can lead to the selection of spurious, non-physical terms, ultimately causing trajectory divergence in unexplored scenarios. To overcome this library dependency, future iterations could leverage advanced neural architectures, such as Kolmogorov-Arnold Network (KAN), Symbolic Regression or even Nested-SINDy, which dynamically construct function rather relying on a predefined set of static terms.

Looking forward, the extraction of explicit ODEs via SINDy unlocks the potential to apply optimal control methods, shifting the paradigm from predictive to prescriptive resilience. The integration of Model Predictive Control (MPC) with SINDy has already proven in dynamic environments. Applying this control logic to our POD trajectory would allow autonomous hubs to dynamically deploy corrective forces (e.g., dynamic capacity allocation) to actively steer the network back to decision-maker's target stability zone.

References

- Brunton S. L., J.L. Proctor, J.N. Kutz (2016): *Discovering governing equations from data by sparse identification of nonlinear dynamic system.* PNAS, 113(15):3932–3937, 2016.
- Duan Y., J.S. Edwards, Dwivedi Y.K. (2019): *Artificial intelligence for decision making in the era of big data – evolution, challenges and research agenda.* International Journal of Information Management.
- Simchi-Levi D., E. Simchi-Levi (2020). *We Need A Stress Test For Critical Supply Chains*. Harvard Business Review.
- Liu M., Z. Liu, F. Chu, F. Zheng, C. Chu (2021). *A new robust dynamic bayesian network approach for disruption approach for disruption risk assessment under the supply chain ripple effect.* International Journal of Production Research.
- Benaben F., L. Faugere, B. Montreuil, M. Lauras, N. Moradkhani, T. Cerabona, J. Gou and W. Mu (2022). *Instability is the norm! A physics-based theory to navigate among risks and opportunities.* Enterprise Information Systems.
- Kosasih E.E., A. Brintrup (2022). *A machine learning approach for predicting hidden link in supply chain with graph neural networks.* International Journal of Production Research.
- Cuong T.N., H-S. Kim, S-S. You, D.A. Nguyen (2022). *Seaport throughput forecast and post-covid-19 recovery policy using effective decision-making strategy: A case study of vietnam ports.* Computers Industrial Engineering.
- Moradkhani N., F. Benaben (2022). *Multi-criteria performance analysis based on physics of decision*. IEEE Transactions on Services Computing.
- Ben Rabia M.A., A. Bellabdaoui (2023). *Collavorative intuitionistic fuzzy-AHP to evaluate simulation-based analytics for freight transport.* Expert Systems with Applications.
- Richey Jr. R.G., S. Chowdhury, B. Davis-Sramek, M. Giannakis, Y.K. Dwivedi (2023). *Artificial intelligence in logistics and supply chain management. A primer and roadmap for research.* Journal of Business Logistics.
- El Harraki I., M.Z. Abedin, A. Belhadi, S. Kamble, K. Zkik, M. Oudani (2024). *Data-driven control and prey-predator model for sourcing decisions in the low-carbon intertwined supply chain.* Business Stategy and the Environment.
- Moradkhani N., F. Benaben (2024). *A force-inspired paradigm for performance-based decision support.* Journal of Industrial Information Integration.
- Ivanov D. (2025). *Literature review on the stress tests developed and applied in non-food supply chains.* Publications Office of the European Union.
- D'Elia R., A. Termine, F. Flammini (2025). *Towards explainable decision support using hybrid neural models for logistics terminal automation.* ArXiv.
- Wei Z., Y. Zhou, L. Du, Y. Zhang (2025). *Discover physically analyzable governing nonlinear ordinary differential equations of traffic network flow dynamics.* Expert Systems with Applications.
- Lu Q-L., S. Raphael, M. Sadrani, C. Antoniou (2025). *Dynamic network capacity allocation using model predictive control with sparse identification of nonlienar dynamics.* IEEE Transactions on Intelligent Transportation Systems.
- D. Ivanov (2025). *When is the supply chain resilient? Customer and operational perspectives.* International Journal of Production Research.

- Morvan B., N. Moradkhani, G. Martin, V. Romero, B. Montreuil, F. Benaben (2025). *Predictive vision of physics of decision: Modelisation of virus propagation with force fields paradigm.* In 2025 Winter Simulation Conference (WSC).

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

KEDGE BUSINESS SCHOOL

Optimizing container capacity for light cargo parcel delivery

Luis Marques[1], Yann Bouchery[1], Olivier Labarthe[1], and Nafé Moradkhani[1]

1. KEDGE Business School - CESIT, Talence, France

Corresponding author: luis.marques@kedgebs.com

Abstract

Light cargo vehicles (LCVs) such as cargo bikes offer a sustainable alternative to vans for urban last mile delivery, but their limited capacity requires efficient parcel encapsulation to facilitate the modal shift from traditional transport. In line with the Physical Internet vision of standardized, modular π-containers, logistics service providers increasingly rely on parcel containers to organize deliveries and modal shifts efficiently. This paper introduces the Parcel Container Capacity Sizing (PCCS) problem, a strategic optimization problem that jointly determines the optimal capacity of parcel containers and a territory design minimizing long term costs. These costs encompass district design costs, trip costs, and distance costs, the latter two being estimated through continuous approximation formulas extending Daganzo's classical models to account for parcel container capacity constraints. We formulate the PCCS problem as a non-linear integer program and propose dual and primal heuristics to respectively compute lower and upper bounds on the optimal cost. Evaluated on 42 real-world instances from La Poste (French postal service) covering the Toulouse area, the heuristics achieve near-optimal quality on short planning horizons (average gap below 6% up to two weeks) but produce large average gaps (up to 58.66%) as the planning horizon increases. An advantage of these heuristics is their tractability in only a few seconds, whereas MIP solvers exhaust available memory before the branch-and-bound begins. A sensitivity analysis on 52-week instances across 25 cost combinations shows that the parcel container capacity increases with distance costs and decreases with trip costs on the tested instances, while the number of districts increases with both distance and trip costs, providing La Poste with actionable guidance tailored to their specific cost structure.

Keywords: *Light cargo logistics, Last mile logistics, Containerization, Encapsulation, Capacity Sizing, Physical Internet*

1. Introduction

E-commerce has driven a rapid surge in urban parcel volumes, heightening pressure on cities that already house over 55% of the world population, a number projected to reach 68% by 2050 (United Nations, 2019). Meeting this demand sustainably is one of the defining logistical challenges of the coming decades.

Parcel delivery in urban areas, part of what's known as *last mile logistics*, presents considerable challenges, as cities are densely populated and highly susceptible to air, noise, and light pollutions. Cities are responsible for 70% of global pollutant emissions World Economic Forum (2020), and urban freight accounts for 6 to 18% of traffic and 21% of CO_2 in those. With recent studies projecting the 100 largest cities to see a 36% increase in delivery vehicles and a 32% rise in pollutant emissions by 2030, reducing the environmental footprint of the last mile delivery has therefore become a pressing societal concern.

Light cargo vehicles, e.g., cargo bikes and cargo trikes, and the Physical Internet framework together offer a promising path toward more sustainable urban logistics. We focus on the intersection of these two directions.

Light cargo logistics. *Light cargo vehicles* (LCVs) are vehicles whose capacity is lower than a van's and are typically human-powered and electrically assisted. Their limited onboard capacity makes them best suited for short urban routes and requires multiple trips per workday rather than a single long route. The term *light cargo logistics* is used to denote logistics using LCVs.

On the one hand, due to LCVs being small and light, their pollutant emissions are much lower when compared to vans, and their noise pollution is close to zero. Replacing traditional vans with LCVs yields

substantial environmental gains: a London case study found a 62% reduction in CO_2 emissions per parcel and a 54% reduction in kilometers traveled (Leonardi et al., 2011). Their small footprint also eases congestion and parking pressure in dense urban areas.

On the other hand, their lower range and energy capacity make LCVs unsuitable for long distances: simulations in Porto suggest a practical limit of around 2 km (Melo and Baptista, 2017). For this reason and the lack of specialized infrastructure, the routes traveled by LCVs are significantly altered compared to those traveled by vans: whereas vans allow for a large single route to fill the workday, LCVs require using several smaller routes.

Because of these constraints, parcels are typically prepositioned at urban *satellites* by conventional vehicles, then transferred to LCVs for last mile delivery. This modal shift, illustrated in Figure 1 for La Poste, requires significant labor, time, and real estate, making its efficiency a key lever in light cargo logistics.



Figure 1: Modal shift by La Poste.

Physical Internet. The Physical Internet (PI) (Montreuil et al., 2012) envisions a global logistics network modeled on the digital Internet, where goods travel in standardized, modular π-*containers* through an open, interconnected network. The PI paradigm directly addresses three important inefficiencies: heterogeneous parcel shapes that complicate automation, manual sorting at modal transfer points, and the lack of interoperability across operators. π-containers, which play the role of physical data packets, encapsulate parcels in uniform units to enable automated handling and seamless mode transitions. This paper explores parcel containers as a practical instantiation of π-containers for LCV-based last mile delivery.

Although many light cargo operators already deploy parcel containers as illustrated by La Poste in Figures 2a and 2b, the question of the optimal parcel container capacity remains open. Containerization is well-studied in maritime logistics (Notteboom and Rodrigue, 2008; Slijper, 1975), but its strategic design for last mile light cargo operations has received little attention. This paper addresses this gap by introducing the Parcel Container Capacity Sizing (PCCS) problem and developing solution methods validated on real-world La Poste instances.

2. Related literature

The PCCS problem lies at the intersection of two literature streams: containerization in last mile logistics and territory design. We briefly review each and identify the gap this work addresses.

Containerization in last mile logistics. Containerization is well established in middle mile logistics, where it reduces parcel heterogeneity, enables automated handling, and streamlines loading at consolidation points. Last mile logistics benefit from a broad range of vehicles, e.g., vans, trucks, LCVs, but this heterogeneity creates interoperability challenges: without standardized containers, parcels must be individually handled at each modal shift point.



(a) Parcel container.



(b) Parcel containers in a cargo bike.

Figure 2: Parcel containers at La Poste*, the French postal service.*

On the one hand, van and truck based last mile logistics typically foregoes detachable containers. Parcels are loaded individually at distribution centers, with the vehicle itself acting as the effective container unit.

On the other hand, LCVs often carry a single large container named cargo container in which parcels are contained. Recent work addresses these light cargo systems from routing, standardization, and cargo container swapping perspectives (Mohri et al., 2024; Conway et al., 2012; Cairns and Sloman, 2019).

To the best of our knowledge, far less attention is given to the use of parcel containers in light cargo logistics, which can be seen as subdivisions of the cargo container into multiple smaller units. A work of particular interest, by Moradkhani et al. (2025), maps the principle of encapsulation in a Physical Internet context with last mile encapsulation through a case study using parcel containers (named *bags* in their paper) with cargo bikes.

Containerized parcel delivery also appears in systems based on public transport, where parcels are transported via tram or bus before LCV handoff; see Amazon in Frankfurt (About Amazon, 2024), GLS in Madrid (Jiménez, Marta, 2024), and La Poste in Strasbourg and Bordeaux (La Poste, 2024; Dupouy et al., 2025). In these systems, container capacity is not a central design decision, since containers are large and primarily used for pre-LCV consolidation. Moreover, territory design plays no role, as routes follow fixed transit lines.

Territory design approaches. Territory design problems partition a service area into *districts*, each assigned to a single LCV in our context. They take as input a fine grain partition of the service area into small geographical units named *basic units* (BUs) and aggregate them to form balanced, contiguous, and compact districts. For more detailed surveys on territory design problems, we refer the reader to Kalcsics et al. (2005) and Laporte et al. (2019, Chapter 25). In this paper, we consider only the contiguity measure for the territory design.

These approaches facilitate sorting parcels at the city distribution center, as not all parcels may be known when sorting begins. Knowing a parcel's destination suffices to determine the vehicle that will deliver it, since a vehicle delivers only the parcels of a single district. These approaches also foster familiarity between end customers and couriers, which is correlated with higher customer satisfaction (Kovacs et al., 2014).

Introducing parcel containers into a district-based framework adds a critical capacity decision: each parcel container should serve a spatially coherent cluster of end customers. Too small a capacity splits nearby customers' parcels across containers, degrading routing efficiency; too large a capacity leaves significant unused space. Optimal sizing therefore depends jointly on district structure and demand density.

The literature reveals a gap at the intersection of these research streams: while the PI provides the conceptual framework for standardized containers, and territory design approaches address urban delivery complexity, no existing work combines these elements with explicit consideration of parcel container capacity sizing for light cargo logistics. This paper aims to bridge this gap with the following contributions:

- Formalization of the PCCS problem, combining parcel container capacity sizing and territory design.
- Continuous approximation formulas to estimate tour lengths and costs with parcel containers.
- A non-linear integer programming formulation with dual and primal heuristics to compute lower and upper bounds.

The remainder of this paper is organized as follows. Sections 3 and 4 define the PCCS problem and formulate it as a non-linear integer program. Section 5 presents the dual and primal heuristics. Section 6 reports computational experiments on real-world La Poste instances. Section 7 concludes.

3. Problem statement

The PCCS problem jointly determines a parcel container capacity c and a territory design minimizing long term costs. These costs comprise:

- *District costs*: $u \in \mathbb{R}^+$ per district, covering the amortized LCV purchase cost.
- *Trip costs*: $s_k \in \mathbb{R}^+$ for the k-th LCV trip, covering driver salary and penalties in case of too many trips that need to be covered by overtime or additional means of transport.
- *Distance costs*: $m \in \mathbb{R}^+$ per unit of distance, covering maintenance and energy.

District costs are strategic costs incurred once every few years; trip and distance costs are operational costs and averaged over the planning horizon.

Hypothesis and simplifying assumptions. In order to model the PCCS problem, we use several hypotheses and simplifying assumptions. First, we express capacities in number of parcels, exploiting the near-uniform shoe-box size of most parcels (Merchán et al., 2024) that allows us to directly connect the capacity of the parcel container in number of parcels to its volume. This simplifies the packing of the parcels by considering only their number instead of their three-dimensional volumes.

Second, each district is served by the closest satellite and satellites are decoupled. This reduces the problem to a single-satellite setting, and we focus exclusively on the LCV portion of the delivery.

Lastly, parcel containers for each BU are filled sequentially, i.e., parcels are sorted into the current open container until it is full, then a new one is opened.

Note that, in practice, end customers' demand is uncertain. In operational research, a common way to handle such uncertainty is scenario-based optimization, where a finite set of possible demand realizations, the so-called scenarios, is generated and used within the optimization model to represent variability. In this paper, we consider instances with historical demand for a set of days, treat each time step as a scenario and optimize over their average.

Problem definition. An instance of the PCCS problem is characterized by a tuple $(\mathcal{T}, V, \mathcal{C}, \mathcal{B}, D)$, where $\mathcal{T}$ is the set of time steps, $V \in \mathbb{N}_{>0}$ is the homogeneous capacity of the LCVs, $\mathcal{C}$ is the set of capacities the parcel containers can take, $\mathcal{B}$ is the set of BUs, and D is the maximum number of districts that can be designed. For each time step $t \in \mathcal{T}$ and BU $b \in \mathcal{B}$, we denote N_b^t the demand of BU b at time step t. For the sake of clarity, we denote N_c the number of parcel containers with capacity c that the LCV can carry, i.e., $N_c = \lfloor V/c \rfloor$.

A solution to the PCCS problem is characterized by a tuple (c, ϕ), where $c \in \{1, \dots, V\}$ is the capacity of the parcel containers and $\phi : \mathcal{B} \to \mathcal{D}$ is a function assigning a BU to a district. The PCCS problem involves determining a solution (c, ϕ) such that the long term costs are minimized.

For each capacity $c \in \mathcal{C}$, let x_c be a binary decision variable equal to 1 if capacity c is chosen for the parcel containers, and 0 otherwise. For each BU $b \in \mathcal{B}$ and district $d \in \mathcal{D}$, let $y_{b,d}$ be a binary decision variable equal to 1 if BU b is assigned to district d, and 0 otherwise. Let also $\mathcal{R}$ be the set of BU to district assignments that form contiguous districts. An integer program (IP) to solve the PCCS problem is of the form

$$\begin{aligned}
\text{minimize} \quad & \alpha && && \text{(1a)}\\
\text{subject to} \quad & \sum_{c \in \mathcal{C}} x_c = 1, && && \text{(1b)}\\
& \sum_{d \in \mathcal{D}} y_{b,d} = 1, && \forall b \in \mathcal{B}, && \text{(1c)}\\
& \boldsymbol{y} \in \mathcal{R}, && && \text{(1d)}\\
& \boldsymbol{x} \in \{0,1\}^{|\mathcal{C}|}, && && \text{(1e)}\\
& \boldsymbol{y} \in \{0,1\}^{|\mathcal{B}|\times|\mathcal{D}|}, && && \text{(1f)}
\end{aligned}$$

where $\alpha \in \mathbb{R}^+$ denotes the value of the long term costs associated with decisions $\boldsymbol{x}$ and $\boldsymbol{y}$.

The objective function in (1a), detailed in Section 4, minimizes the total long term costs associated with the chosen capacity and territory design. Constraint (1b) ensure that only one capacity is chosen for the parcel containers. Constraints (1c) ensure that each BU is assigned to a district, and constraints (1d) ensure that each district is contiguous. We model these constraints using the flow-based formulation of Shirabe (2009). Finally, constraints (1e)-(1f) state that all variables are binary.

The core constraints of (1) are linear and involve few variables. The difficulty lies in the objective α that may require additional decision variables and constraints to formulate. Routing costs depend on LCV routes, which we estimate by extending the continuous approximation models of Daganzo (1984) and Robust et al. (1990) to account for parcel container capacity and the sequential delivery constraint stating that the parcels of the current open container must be delivered before opening another parcel container.

4. Objective function of the PCCS problem

In this section, we detail each cost component in turn, the continuous approximation used and its translation into the IP.

District costs. Districts costs are modeled exactly by introducing binary variables $\boldsymbol{z} \in \{0,1\}^{|\mathcal{D}|}$. For each district $d \in \mathcal{D}$, z_d is equal to 1 if district d has been designed, i.e., if it has at least one assigned BU, and 0 otherwise. In other terms,

$$y_{b,d} \leq z_d, \quad \forall b \in \mathcal{B}, d \in \mathcal{D}, \tag{2}$$

giving

$$\alpha_1 = u \cdot \sum_{d \in \mathcal{D}} z_d. \tag{3}$$

Trip costs. Trip costs are approximated by estimating the number of trips necessary for LCVs to serve each district. Let

$$R^t_{d,c}(\boldsymbol{y}) = \frac{\sum_{b \in \mathcal{B}} \left\lceil \frac{N^t_b}{c} \right\rceil y_{b,d}}{N_c} \tag{4}$$

be an estimation of the number of trips necessary to serve district d with parcel containers of capacity c under the territory design given by variables $\boldsymbol{y}$ and let K^t be an upper bound on $R^t_{d,c}(\boldsymbol{y})$ for all $d \in \mathcal{D}$, $c \in \mathcal{C}$, and all territory designs $\boldsymbol{y}$. Binary variables $g^t_{d,k}$ indicate whether district d requires at least k trips at time step t. Constraint (5) forces $g^t_{d,k} = 1$ when needed by inspecting whether $R^t_{d,c}(\boldsymbol{y})$ is larger than k. Note that setting variables $\boldsymbol{g}$ to 0 is unnecessary since their objective coefficient is positive.

$$g^t_{d,k} \geq \sum_{c \in \mathcal{C}} \frac{R^t_{d,c}(\boldsymbol{y}) - k}{K^t} \cdot x_c - 1. \tag{5}$$

Trips costs are then formulated as

$$\alpha_2 = \frac{1}{|\mathcal{T}|} \cdot \sum_{t \in \mathcal{T}} \sum_{d \in \mathcal{D}} \sum_{k=1}^{K} s_k g^t_{d,k}. \tag{6}$$

Distance costs. We use the following methodology to approximate the distance traveled by an LCV: (i) we continuously approximate the delivery of the parcels of each parcel container in its corresponding BU, (ii) we create a capacitated vehicle routing problem (CVRP) using the parcel containers as end customer and estimate the optimal value of this problem with another continuous approximation. We name the first costs *intra-BU costs*, and the second costs *inter-BU costs*.

The routing associated with the parcels of a parcel container is similar to that of a traveling salesman problem (TSP), since the LCV has to visit each end customer without capacity constraints. Therefore, we build on the continuous approximations of Daganzo (1984) for the TSP, which are

$$k_1 \cdot \sqrt{A} \cdot \sqrt{n}, \tag{7}$$

where A is the area of the service zone, n is the number of end customers, and k_1 is a coefficient related to the shape of the service zone.

Applied to the PCCS problem and due to the sequential filling of parcel containers, BU b at time step t yields $\lfloor N^t_b / c \rfloor$ full containers and possibly one partial container with $N^t_b \bmod c$ parcels. Summing the TSP approximation over all parcel containers gives the intra-BU costs α_3 in (8).

$$\alpha_3 = \frac{m}{|\mathcal{T}|} \cdot \sum_{t \in \mathcal{T}} \sum_{c \in \mathcal{C}} \sum_{b \in \mathcal{B}} k_1 \cdot \sqrt{A(b)} \cdot \left(\left\lfloor \frac{N^t_b}{c} \right\rfloor \cdot \sqrt{c} + \sqrt{N^t_b \bmod c} \right) \cdot x_c. \tag{8}$$

Inter-BU routing is more similar to that of a CVRP, since the LCV has a limit on how many parcel containers it can carry. To approximate this component, we build on the continuous approximations of Robust et al. (1990) for the CVRP, which are

$$\left(k_1 + \frac{k_2 \cdot n}{\tilde{c}^2} \right) \cdot \sqrt{A} \cdot \sqrt{n} \tag{9}$$

where A is the area of the service zone, n is the number of end customers, $\tilde{c}$ is the vehicle capacity, and k_1 and k_2 are coefficients related to the shape of the service zone. The additional $k_2 n / c^2$ term captures the detour length from capacity constraints. Summing the CVRP approximation over all districts gives the inter-BU costs α_4 in (10).

$$\alpha_4 = \frac{m}{|\mathcal{T}|} \cdot \sum_{t \in \mathcal{T}} \sum_{d \in \mathcal{D}} \left(k_1 + k_2 \cdot \sum_{c \in \mathcal{C}} \frac{\sum_{b \in \mathcal{B}} \left\lceil \frac{N^t_b}{c} \right\rceil \cdot y_{b,d}}{N^2_c} \cdot x_c \right) \cdot \sqrt{\sum_{b \in \mathcal{B}} A(b) \cdot y_{b,d}} \cdot \sqrt{\sum_{c \in \mathcal{C}} \sum_{b \in \mathcal{B}} \left\lceil \frac{N^t_b}{c} \right\rceil \cdot y_{b,d} \cdot x_c}. \tag{10}$$

Long term costs. Finally, we define the long term costs as

$$\alpha = \alpha_1 + \alpha_2 + \alpha_3 + \alpha_4. \tag{11}$$

The resulting IP has the advantages of being decomposable by capacity $c \in \mathcal{C}$, as there are no coupling constraints for different capacities and there are not many possible capacities in $\mathcal{C}$. In fact, decomposing the IP by capacity results in having a set of linear constraints. However, the objective function is non-linear and non-convex/non-concave even when the problem is decomposed by capacity, precluding off-the-shelf MILP solvers to solve the problem.

5. Solution approach

A key structural property of the IP (1) is its decomposability by capacity: fixing $c \in \mathcal{C}$ reduces the problem to a territory design subproblem over $\boldsymbol{y}$ alone. This property is exploited by both heuristics we describe in this section, which iterate independently over all capacities and combine the results.

Dual heuristic. To bound the optimal objective from below without solving the integer program, we compute a lower bound by iterating over all capacities $c \in \mathcal{C}$ and, for each, summing three independent components. The overall bound is the minimum across capacities. We explain in the two following paragraphs the computation of these components for fixed capacity $c \in \mathcal{C}$.

District and trip costs. For each candidate number of districts $p \in \{1, \dots, D\}$, the district cost is $p \cdot u$. To obtain a dual bound on trip costs, we distribute for each time step $t \in \mathcal{T}$ $\lceil \sum_{b \in \mathcal{B}} \lceil N_b^t / c \rceil / N_c \rceil$ trips as evenly as possible across p districts to minimize the sum of non-decreasing trip costs s_k; the minimum over $\{1, \dots, D\}$ is retained.

Distance costs. For intra-BU distance costs, note that α_3 depends only on demands and on the containers capacity. Therefore, we compute these costs exactly. For inter-BU distance costs, we compute a territory design independent dual bound from two inequalities holding for any territory design. Let $A(d)$ (resp. N_d^t) denote the area (resp. the number of parcel containers at time step t) of any district d. First, the Cauchy-Schwarz inequality gives, for each time step t, $\sum_{d \in \mathcal{D}} \sqrt{A(d) N_d^t} \geq \sum_{b \in \mathcal{B}} \sqrt{A(b) \lceil N_b^t / c \rceil}$. Second, the superadditivity of $z \mapsto z \cdot \sqrt{z}$ combined with the monotonicity of $z \mapsto \sqrt{z}$ gives $\sum_{d \in \mathcal{D}} \sqrt{A(d)} \cdot N_d^t \cdot \sqrt{N_d^t} \geq \sum_{b \in \mathcal{B}} \sqrt{A(b)} \cdot \lceil N_b^t / c \rceil \cdot \sqrt{\lceil N_b^t / c \rceil}$. Both yield per-BU bounds summed over b and t, requiring no knowledge of the territory design.

Primal heuristic. To obtain a feasible solution and an upper bound without solving the integer program, we propose a BFS-based greedy construction heuristic. For each candidate capacity, the algorithm selects a BU as seed, assigns it to a new district, and initializes a BFS queue with its unassigned neighbors. For each BU dequeued, the algorithm identifies districts adjacent to it, i.e., districts that already contain a neighbor of that BU, and compare the incremental cost of joining the cheapest adjacent district against the cost of opening a new district, and take the cheaper option subject to the district budget. By construction, each district consists of mutually adjacent BUs, enforcing contiguity. The incremental cost accounts for the district fixed cost if the district was previously empty, the additional trip costs across all time steps, and the change in inter-BU routing cost via (10). The intra-BU cost, which is independent of the territory design, is accumulated once at the beginning. The procedure runs independently for each candidate capacity, and the best feasible solution across all capacities is returned as the primal bound.

The solution returned by the primal heuristic, its objective value, and the dual bound computed by the dual heuristic can then serve as warm start parameters for general purpose MIP solvers. However, preliminary results show that the improvements obtained by the additional step are limited, as even when providing the heuristic results as warm start, the solver we tested (SCIP Version 10.0.2) results in memory usages larger than 64 Go before the branch-and-bound begins. Therefore, we focus on numerical experiments using the heuristics only in the next section.

6. Numerical experiments

In this section, we report preliminary experiments using dual and primal heuristics on real-world La Poste instances, and provide a sensitivity analysis of the impact the cost parameters have on the computed solutions.

Experimental setup. We evaluate the heuristics on real-world delivery data from La Poste, covering a territory in the Toulouse area. Instances are grouped into seven temporal benchmarks, *1-day*, *1-week*, *2-weeks*, *4-weeks*, *12-weeks*, *24-weeks*, and *52-weeks*, reflecting single and multi-period planning horizons,

and two demand groups, *small* and *large*, based on total daily parcel volume. This yields 42 instances: 20 for 1-day instances, 4 instances each for the 1-week to 24-weeks benchmarks, and 2 for 52-week instances. The LCV capacity is set to 100 parcels and the maximum number of districts to 10. Large-demand instances contain on average 1103 to 14701 parcels per day; small-demand instances contain on average 437 to 1598 parcels per day. The heuristics were run on MacOS with an Apple M4 CPU and 16 Go of RAM. For the sensitivity analysis, we vary the distance cost m across five levels (0.38, 1.0, 1.5, 2.0, and 2.5 EUR/km) and a trip cost scale s across five levels ($0.25\times$, $0.5\times$, $1\times$, $2\times$, and $4\times$ the baseline trip costs $[80, 80, 160, 500]$ EUR), yielding 25 cost combinations run on the 52-week instances. These values were chosen in accordance with typical French costs and salaries.

Results: heuristic quality. We first detail the results obtained by the heuristic algorithms. Table 1 summarizes these outcomes. The reported gap is $(\text{primal} - \text{dual})/\text{primal}$, i.e., the relative distance between the primal bound and the dual bound; the reported time is the combined wall-clock time of the dual and primal heuristics. On single-day instances, performance is near-optimal for both demand groups (avg. gap 3.08% for large and 1.46% for small). Small-demand runs complete in under 0.1 seconds; large-demand instances complete in under 4 seconds.

Table 1: Heuristic results on La Poste instances.

Instance group	Demand	Instances	Avg. gap (%)	Avg. time (s)
1-day	large	10	3.08	3.3
1-day	small	10	1.46	0.0
1-week	large	2	4.17	3.7
1-week	small	2	28.64	0.1
2-weeks	large	2	5.45	4.0
2-weeks	small	2	46.89	0.2
4-weeks	large	2	24.99	2.4
4-weeks	small	2	46.74	0.1
12-weeks	large	2	54.38	0.4
12-weeks	small	2	58.66	0.4
24-weeks	large	2	30.00	5.5
24-weeks	small	2	58.44	0.6
52-weeks	large	1	37.33	8.2
52-weeks	small	1	42.84	7.7

For multi-period instances, gaps tend to grow with the planning horizon. For large-demand instances, gaps remain low up to 2 weeks (4.17% and 5.45%) before rising sharply at longer horizons, with a non-monotonic pattern: the average gap reaches 54.38% at 12 weeks before falling to 30.00% at 24 weeks, then rising to 37.33% at 52 weeks. For small-demand instances, gaps grow rapidly from 1.46% at 1 day to 28.64% at 1 week and reach 43–59% for horizons of 4 weeks and beyond. We attribute this loosening to the dual bound structure: the inter-BU lower bound is derived per time step independently, whereas the primal must commit to a single territory design simultaneously accommodating all demand patterns across the planning horizon.

Sensitivity analysis: capacity as a function of costs. A key advantage of our approach is that it can be applied to any LSP's cost structure and deliver a tailored capacity and territory design. To illustrate this, Table 2 reports the mean parcel container capacity $\bar{c}$ observed for the two 52-week instances across the 25 combinations of costs.

Two main trends emerge from Table 2. First, $\bar{c}$ increases with the distance cost m: higher per-kilometer costs amplify the routing efficiency gain of larger containers, and the trend is clearest at low trip costs (the $0.25\times$ row increases monotonically from 2.8 to 4.7). Second, $\bar{c}$ decreases as the trip cost scale s rises, most visibly in the $m = 2.0$ column ($4.3 \to 2.2$ as s ranges from $0.25\times$ to $4\times$). This direction is explained by the sparse demand in the La Poste territory (approximately 2 parcels per BU per day): increasing c reduces containers per BU but also reduces $N_c = \lfloor V/c \rfloor$, so the LCV carries fewer containers per trip

Table 2: Mean parcel container capacity $\bar{c}$ selected by the primal heuristic as a function of distance cost m and trip cost scale s.

$s \backslash m$	0.38	1.0	1.5	2.0	2.5
$0.25\times$	2.8	3.4	3.7	4.3	4.7
$0.5\times$	1.9	2.9	2.8	3.6	4.0
$1\times$	1.9	2.8	2.7	3.0	3.8
$2\times$	1.7	2.6	2.3	2.3	2.8
$4\times$	1.7	1.9	1.9	2.2	2.1

and must make more trips overall. When trips are costly, small containers are therefore preferred. At $s = 4\times$ this effect dominates, yielding $\bar{c} \leq 2.2$ regardless of m.

Across the 25 cost combinations, the mean number of districts ranges from 2.0 at low m and s to 8.4 at high m and s, with an overall average of 4.1 out of a maximum of 10. The district count increases with both m and s: at $s = 4\times$ every column yields an average above 6 (approaching the maximum), while at $s = 0.25\times$ the average stays below 3.0 regardless of m. These results illustrate that the PCCS framework provides actionable guidance: given any LSP's cost structure, it rapidly computes a tailored capacity and territory design.

7. Conclusion

This paper addressed the problem of capacity sizing parcel containers to enhance the efficiency of light cargo and last mile logistics. We formulated the Parcel Container Capacity Sizing (PCCS) problem, a strategic optimization problem at the intersection of light cargo logistics and territory design, with close links to the Physical Internet paradigm.

The contributions of this paper are threefold. First, we formalized the PCCS problem, which jointly determines parcel container capacity and territory design to minimize long term costs encompassing distance, trip, and district costs. Second, we developed continuous approximation formulas extending Daganzo's classical models to account for parcel container capacity constraints and the sequential delivery requirement. Third, we proposed dual and primal heuristics that achieve near-optimal quality on short planning horizons (average gap below 6% up to two weeks) but produce large average gaps (up to 58.66%) as the planning horizon increases. Further, these heuristics remain tractable in seconds for the largest instances, whereas MIP solvers exhaust available memory before branch-and-bound begins.

From a managerial perspective, the PCCS framework provides LSPs with actionable guidance calibrated to their own cost structure. Our sensitivity analysis shows that the parcel container capacity increases with distance cost (reaching $\bar{c} \approx 3$–5 at 2.0–2.5 EUR/km) and decreases with trip cost, while the number of districts increases with both cost parameters. Rather than relying on intuition or ad-hoc practices, managers can use this approach to evaluate the long term implications of different container capacities and territory designs.

Several directions for future research merit investigation. First, the dual bound loosens significantly at longer planning horizons, and tightening it, for instance through scenario-aggregation techniques, would extend the range of horizons for which the heuristics provide reliable quality guarantees. Second, empirical validation of the continuous approximation models against actual routing data would strengthen practical relevance. Third, extending the model to heterogeneous container sizes or LCV capacities could provide additional flexibility for territories with uneven demand density. Finally, sample average approximation methods could determine how many time steps are required to obtain a reliable estimation of the long term costs.

References

About Amazon (2024). 'LastMileTram' in Germany to deliver packages with zero-tailpipe emissions in last mile. https://www.aboutamazon.eu/news/sustainability/lastmiletram-in-germany-to-deliver-packages-with-zer

o-tailpipe-emissions-in-last-mile.

Cairns, S. and Sloman, L. (2019). Potential for e-cargo bikes to reduce congestion and pollution from vans in cities. *Transport for Quality of Life Ltd*.

Conway, A., Fatisson, P.-E., Eickemeyer, P., Cheng, J., and Peters, D. (2012). Urban micro-consolidation and last mile goods delivery by freight-tricycle in manhattan: Opportunities and challenges. In *Conference proceedings, Transportation Research Board 91st Annual Meeting*.

Daganzo, C. F. (1984). The Distance Traveled to Visit N Points with a Maximum of C Stops per Vehicle: An Analytic Model and an Application. *Transportation Science*, 18(4):331–350.

Dupouy, C., Clautiaux, F., Klibi, W., and Labarthe, O. (2025). Multi-stakeholder Insights on Integrating Public Transport for Urban Freight Deliveries. In *IPIC 2025 - 11th International Physical Internet Conference*, Hong Kong, China. The Hong Kong Polytechnic University, Georgia Institute of Technology.

Jiménez, Marta (2024). Así funciona el transporte de paquetería en Metro de Madrid: primeras pruebas piloto con GLS. https://logistica.cdecomunicacion.es/ultima-milla/145095/video-asi-funciona-el-transporte-de-paqueteria-en-metro-de-madrid.

Kalcsics, J., Nickel, S., and Schröder, M. (2005). Towards a unified territorial design approach—applications, algorithms and gis integration. *Top*, 13(1):1–56.

Kovacs, A. A., Golden, B. L., Hartl, R. F., and Parragh, S. N. (2014). Vehicle routing problems in which consistency considerations are important: A survey. *Networks*, 64(3):192–213.

La Poste (2024). La livraison de colis par le tramway de Strasbourg, une expérimentation pour un centre-ville plus fluide. https://www.lapostegroupe.com/fr/actualite/la-livraison-de-colis-par-le-tramway-de-strasbourg-une-experimentation-pour-un-centre-ville-plus-fluide.

Laporte, G., Nickel, S., and Saldanha Da Gama, F., editors (2019). *Location Science*. Springer International Publishing, Cham.

Leonardi, J., Browne, M., and Allen, J. (2011). Véhicules électriques et réseau de plates-formes de fret à londres. , *Seminaire SELEST*.

Melo, S. and Baptista, P. (2017). Evaluating the impacts of using cargo cycles on urban logistics: Integrating traffic, environmental and operational boundaries. *European Transport Research Review*, 9(2):30.

Merchán, D., Arora, J., Pachon, J., Konduri, K., Winkenbach, M., Parks, S., and Noszek, J. (2024). 2021 Amazon Last Mile Routing Research Challenge: Data Set. *Transportation Science*, 58(1):8–11.

Mohri, S. S., Mohammadi, M., and Van Woensel, T. (2024). Designing zero-emissions containerized last-mile delivery systems: A case study for melbourne. *Transportation Research Part C: Emerging Technologies*, 159:104492.

Montreuil, B., Meller, R. D., and Ballot, E. (2012). Physical Internet Foundations. *IFAC Proceedings Volumes*, 45(6):26–30.

Moradkhani, N., Bouchery, Y., Lebeau, F., Labarthe, O., and Klibi, W. (2025). Toward modular equivalence: Mapping last-mile encapsulation to the physical internet container ecosystem.

Notteboom, T. and Rodrigue, J.-P. (2008). Containerisation, box logistics and global supply chains: The integration of ports and liner shipping networks. *Maritime economics & logistics*, 10(1):152–174.

Robust, F., Daganzo, C. F., and Souleyrette II, R. R. (1990). Implementing vehicle routing models. *Transportation Research Part B: Methodological*, 24(4):263–286.

Shirabe, T. (2009). Districting modeling with exact contiguity constraints. *Environment and Planning B: Planning and Design*, 36(6):1053–1066.

Slijper, M. T. (1975). Containerisation: An appraisal. *International Journal of Physical Distribution*, 5(3):165–175.

United Nations (2019). *World Urbanization Prospects: The 2018 Revision*. United Nations.

World Economic Forum (2020). *The Future of the Last-Mile Ecosystem*. World Economic Forum.

A. Container utilization loss

The baseline PCCS formulation accounts for the effect of parcel-container capacity through the number of containers required by each basic unit, namely $\lceil N_b^t/c \rceil$, and through the intra-BU distance approximation, where the last partially filled container is implicitly captured by $N_b^t \bmod c$. This captures the effect of capacity on routing and trip requirements. However, it does not explicitly measure the degree to which the selected capacity leads to well-utilized parcel containers.

This issue is relevant because parcel containers are not only routing units, but also encapsulation units. A capacity may reduce distance or trip costs while still producing many partially filled containers. In such cases, the selected capacity may be efficient from a routing perspective but less satisfactory from an operational encapsulation perspective. To capture this effect without changing the structure of the model, we could introduce a simple container-utilization loss term.

For each time step $t \in \mathcal{T}$, basic unit $b \in \mathcal{B}$, and candidate parcel-container capacity $c \in \mathcal{C}$, define the average fill rate as

$$\rho_{b,c}^t = \frac{N_b^t}{c \lceil N_b^t/c \rceil}. \tag{12}$$

The value $\rho_{b,c}^t = 1$ corresponds to perfectly filled containers, whereas lower values indicate residual unused capacity caused by partially filled containers. The associated utilization loss can be written as

$$1 - \rho_{b,c}^t = 1 - \frac{N_b^t}{c \lceil N_b^t/c \rceil}. \tag{13}$$

An additional cost component can then be defined as

$$\alpha_5 = \frac{\gamma}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \sum_{b \in \mathcal{B}} \sum_{c \in \mathcal{C}} \left(1 - \frac{N_b^t}{c \lceil N_b^t/c \rceil} \right) x_c, \tag{14}$$

where $\gamma \in \mathbb{R}_+$ is a penalty coefficient associated with residual unused container capacity.

If this component is considered, the long-term cost becomes

$$\alpha = \alpha_1 + \alpha_2 + \alpha_3 + \alpha_4 + \alpha_5. \tag{15}$$

This extension preserves the binary structure of the original formulation. No additional assignment variable is required: once a capacity c is selected, the utilization loss is directly determined from historical demand. Moreover, the extension preserves the decomposability by capacity, since α_5 depends only on the selected capacity and the observed demand values.

The term adds a complementary operational interpretation to the PCCS problem. Very small containers may increase the number of containers to be handled, while very large containers may leave more residual unused capacity in the last container of each basic unit. Penalizing unused capacity therefore encourages capacities that are not only efficient in terms of trips and distance, but also better aligned with the objective of effective parcel encapsulation.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE
alice Alliance for Logistics Innovation through Collaboration in Europe

From Fragmented Digitalization to Physical Internet Coordination: Achieving Concrete Data Interoperability

Liz Araceli Cristaldo[1] and Marion Cottet [1]

1. Alliance for Logistics Innovation through Collaboration in Europe (ALICE), Bruxelles, Belgium

Corresponding author: marion.cottet@etp-alice.eu

Abstract: *Data interoperability is essential for coordination in urban mobility and logistics systems in Europe, yet in practice, data remains difficult to share and reuse across actors, systems and platforms. Despite extensive policy efforts and numerous initiatives, data interoperability remains fragmented, limiting effective coordination between transport systems. This paper explores this challenge through the lens of the Physical Internet, which emphasizes standardized and interconnected systems to enable coordinated logistics networks. Based on a combined review of European policies, standards, and data space initiatives, complemented by an exploratory survey of practitioners, the paper identifies persistent gaps across technical, semantic, governance, and operational dimensions. The results show that interoperability challenges are not primarily due to a lack of solutions, but to misalignment across these dimensions, reflected in inconsistent data definitions, unclear governance rules, and limited reuse of data models and APIs. Based on these findings, the paper highlights priority areas for action to better align data interoperability efforts with Physical Internet principles and support more coordinated and scalable urban mobility and logistics systems.*

Keywords: *Data interoperability, European Data Spaces, Urban, Logistics, Mobility, Standardization, Physical Internet.*

1 Introduction

Mobility and logistics systems involve numerous organizations and stakeholders whose operations increasingly rely on digital information exchange and data-driven processes (Pan et al., 2021). In this context, the ability to generate, exchange, interpret, and reuse data across organizational and territorial boundaries, referred to here as *data interoperability*, becomes essential for coordinating urban mobility and logistics systems.

At the same time, European urban mobility and logistics systems face growing pressure to become more sustainable, efficient and resilient. Policy frameworks emphasize the role of data in supporting decision-making, improving performance monitoring, and addressing the diverse and often conflicting needs of stakeholders, while simultaneously supporting the achievement of multiple local policy goals. However, despite these ambitions, effective data sharing between public and private actors remains difficult to implement, limiting the use and reuse of available data (Expert Group for Urban Mobility, 2024). Existing data ecosystems remain fragmented, with heterogeneous data models, and developed through systems that do not communicate effectively with each other (Delaplace, 2024).

In parallel, the Physical Internet (PI) has emerged as a paradigm for coordinating mobility and logistics through standardized units, shared interfaces, and open protocols, enabling seamless interaction between actors and systems (Montreuil, 2011; Wu et al., 2025). More recent research highlights that PI evolves as a cyber-physical paradigm, requiring close integration between physical flows and digital information layers (Crainic and Montreuil, 2016; Wu et al., 2025).

From this perspective, without the ability to consistently share and reuse data across actors and systems, coordination in European urban mobility and logistics remains limited and cannot scale.

Against this background, this paper addresses the following research question: *how are data interoperability frameworks experienced in practice by European mobility and logistics stakeholders, and how do these challenges affect the realization of Physical Internet?*

To answer this question, an analysis incorporating practitioner perspectives is conducted to identify which interoperability barriers are perceived in practice as most critical for enabling coordination across actors, transport modes, and territories. The results provide a basis for prioritizing actions toward more operational and scalable data interoperability frameworks in the European context.

The remainder of this paper is structured as follows. Section 2 introduces the PI perspective on data and standardization. Section 3 reviews European policy, standardization, and data space landscape. Section 4 presents practitioner perspectives to assess and prioritize the main data interoperability gaps. Finally, Section 5 discusses these findings' implications, concludes this paper, and outlines directions for future research.

2 The Physical Internet Perspective on Data Interoperability

The PI provides a conceptual framework for coordinating logistics and mobility systems by drawing inspiration from the architecture of the Digital Internet (Montreuil, 2011; Wu et al., 2025). It proposes a shift from fragmented and vertically integrated transport systems toward an open, distributed, and interconnected network, where flows can move seamlessly across infrastructures and service providers. The objective is to improve efficiency, sustainability, and resilience by enabling greater consolidation, sharing of resources, and coordination across actors (Montreuil, 2011).

A central principle of PI is modularity. Goods are encapsulated into standardized, modular units that can be handled and routed across different networks and operators (Crainic and Montreuil, 2016; Montreuil, 2011). This modular structure allows physical flows to become independent from specific carriers or infrastructures, enabling flexible and dynamic routing across a network of interconnected services. However, achieving such flexibility requires more than physical standardization.

PI also relies on standardized interfaces and protocols, like those of the Digital Internet, enabling independent actors to interact within a common system without requiring centralized control. In this context, coordination emerges from the ability of these actors to exchange information and align their operations based on common rules and communication mechanisms. As a result, PI is increasingly understood as a cyber-physical paradigm, where digital information flows synchronize physical operations across actors, locations, and time scales (Wu et al., 2025).

Within this paradigm, digital systems support routing decisions, coordinate operations, and enable collaboration across actors. These functions depend on the ability of information systems to exchange, interpret, and reuse data across organizational and technical boundaries (Crainic and Montreuil, 2016; Wu et al., 2025). As a result, data interoperability becomes a core coordination layer of PI-enabled systems, particularly in urban mobility and logistics systems, where multiple transport modes, operators, and services coexist and interact within the same physical space.

In such environments, effective coordination relies on the ability of independent systems to exchange data through standardized formats, interfaces, and communication protocols, while remaining operationally autonomous (Pan et al., 2021). However, current data environments remain fragmented, with differences in data models, standards, and governance frameworks limiting data sharing and reuse, constraining coordination, and reducing system efficiency (Bakri et al., 2023; Laborda et al., 2023).

From a PI perspective, these limitations highlight that data interoperability is not only a technical issue, but a structural condition for urban mobility and logistics coordination. Without interoperable data infrastructures, PI remains confined to isolated project pilots or local applications and cannot scale across networks and territories (Delaplace, 2024; Wu et al., 2025).

3 European Context: Policies, Standards, and Data Spaces

Over the last decade, Europe has developed a wide range of policies, regulatory instruments, and technical initiatives aimed at improving data availability, effective data sharing, and interoperability. This section reviews the main components of this landscape, including emerging data space initiatives, standardization efforts, and policy frameworks.

3.1 Data Spaces initiatives

At the technical level, initiatives such as GAIA-X and the International Data Spaces Association (IDSA) provide key building blocks for federated data sharing, including identity management, access control, usage control, and trust mechanisms. IDSA proposes a reference architecture and interoperability protocols, while GAIA-X focuses on a federated trust framework and governance rules. Together, these initiatives support interoperability across technical and organizational layers.

However, important limitations remain. While IDSA explicitly addresses semantic interoperability through shared vocabularies and models, and GAIA-X enables domain-specific extensions, neither provides fully specified data models tailored to specific sectors. As a result, domain-level semantics must often be developed within individual data spaces, which limits consistent data interpretation across systems (Bakri et al., 2023; Delaplace, 2024).

In addition, challenges persist regarding liability and accountability. While these initiatives define frameworks for access control, identity, and usage policies, they do not fully resolve how responsibility should be allocated when shared data is reused, combined, or used for decision-making (Picht, 2020; Vellinga and Hailevich, 2024).

3.2 Standardization frameworks

Standardization efforts at European and international levels, led by organizations such as European Committee for Standardization (CEN) and International Organization for Standardization (ISO), provide essential building blocks for interoperability by defining data models, exchange specifications, and, in some cases, process frameworks across transport and logistics domains (Bakri et al., 2023). These standards contribute to a shared technical foundation and can facilitate data exchange within specific domains, particularly when consistently implemented. For instance, standards such as DATEX II for traffic and road exchange or NeTEx for public transport provide structured data models that support interoperability among stakeholders.

However, in practice, their effectiveness is often limited by heterogeneous implementation across actors (Gharehgozli et al., 2019). Standards may leave room for interpretation, meaning the same concept may be understood differently depending on context. For example, what constitutes a “delivery event” (e.g., vehicle arrival at a distribution center or proof of delivery). Sometimes, cities and operators include fields to adapt the standard to their own operational needs, however these custom additions are not always documented or shared, which limits reuse by other stakeholders.

These differences mean that even when common standards are adopted, variations in interpretation, context, and usage may lead to inconsistencies in how data is produced and consumed.

As a result, while standardization is a necessary condition for interoperability, it is not sufficient on its own to ensure seamless data exchange across the full urban mobility and logistics ecosystem. Achieving effective interoperability requires additional alignment at the semantic, organizational, and operational levels.

3.3 European policies, initiatives and funded projects

At a broader level, the European Interoperability Framework (EIF) provides a common reference for interoperability across public administrations in Europe. It defines interoperability as the ability of organizations to exchange and use information to achieve shared objectives, and highlights that it requires alignment across multiple layers, including legal, organizational, semantic, and technical dimensions (European Commission. Directorate General for Informatics., 2017). While the EIF is not specific to urban mobility and logistics, it establishes a general structure that supports many European digital initiatives and helps frame interoperability as a systemic, multi-layer challenge rather than a purely technical issue.

One key regulatory instrument in the transport domain is the electronic Freight Transport Information (eFTI[1]) Regulation, which focuses on the digitalization of freight transport information exchanges between economic operators and public authorities. Its primary objective is to replace paper-based documents with electronic data, harmonize reporting obligations across Member states, and reduce administrative burdens while ensuring regulatory compliance (Delaplace, 2024). However, its scope remains limited to regulatory reporting (business-to-administration) and does not yet address operational coordination between logistics actors.

In parallel, the European Commission has launched the European Mobility Data Space (EMDS[2]) as part of the broader European Data Strategy. EMDS aims to provide a federated and trusted environment for sharing mobility and logistics data across public and private stakeholders. Its objectives include improving data accessibility, supporting innovation, and enabling new services across passenger and freight transport (Bakri et al., 2023; Laborda et al., 2023). While EMDS explicitly addresses challenges related to interoperability, data sovereignty, and governance, its operational implementation and integration with existing initiatives are still ongoing. The project is expected to run until October 2026.

Numerous European R&I projects (e.g. DELPHI[3], DISCO[4], DECARBOMILE[5], GREEN-LOG[6], UNCHAIN[7], URBANE[8]) are exploring data sharing and interoperability in urban logistics through data spaces, either by building their own infrastructure or aligning with data space principles and architectures. Although developed in specific local contexts, some also explore opportunities for scalability and integration. For instance, URBANE solutions have been integrated into DISCO's Urban Freight Data Space, pointing to a possible pathway towards cross-project interoperability. Yet such approaches remain an exception rather than the norm. Furthermore, the long-term sustainability of these solutions beyond project funding remains a significant challenge.

3.4 Identified Challenges and Gaps

The literature and policy landscape highlight several persistent gaps that continue to limit data interoperability in urban mobility and logistics, despite significant efforts through data space frameworks, standardization initiatives, and European projects. This subsection synthesizes these barriers and analyzes how they translate into operational challenges.

First, technical interoperability remains constrained by semantic heterogeneity. Data models, vocabulary, and metadata structures vary across platforms and initiatives, even when addressing

[1] https://transport.ec.europa.eu/transport-themes/logistics-and-multimodal-transport/efti-regulation_en
[2] https://transport.ec.europa.eu/transport-themes/smart-mobility/creating-common-european-mobility-data-space_en
[3] https://cordis.europa.eu/project/id/101104263
[4] https://cordis.europa.eu/project/id/101103954
[5] https://cordis.europa.eu/project/id/101069806
[6] https://cordis.europa.eu/project/id/101069892
[7] https://cordis.europa.eu/project/id/101103812
[8] https://cordis.europa.eu/project/id/101069782

similar use cases such as shipment tracking, transport planning, or multimodal coordination. This limits the reuse of technical components such as APIs and data models beyond their original context. While such approaches often accelerate experimentation within individual projects, they also result in solutions that are tailored to specific pilots, or local contexts, thereby limiting data interoperability and reuse across initiatives.

Second, significant challenges relate to governance and trust. Data access conditions, usage rights, and reuse mechanisms remain fragmented and inconsistently defined, creating uncertainty for stakeholders and discouraging data sharing. In practice, it is often unclear under what conditions data can be reused, by whom, and for which purposes, particularly in multi-actor and cross-border environments, where it is often unclear which actor is accountable in case of errors or misuse. In addition, the lack of clearly defined and widely adopted trusted intermediaries or third-party governance mechanisms further complicates data sharing, as no neutral entity consistently ensures compliance, trust, and coordination across stakeholders.

Third, policy and standardization efforts, while supporting digitalization and data sharing, do not fully ensure data interoperability in practice. Instruments such as eFTI focus on specific regulatory use cases, particularly business-to-administration data exchange. At the same time, broader European frameworks and standardization initiatives provide general principles and guidelines without defining how data interoperability should be put into practice across the different systems and actors.

Finally, implementation and scalability remain major barriers. Solutions developed in pilot projects are often tailored to specific local contexts and face difficulties in scaling across cities, regions, or countries due to differences in governance structures, digital maturity, procurement practices, and legacy systems. As a result, data-sharing initiatives remain largely project-specific.

Despite the growing body of work on data spaces and interoperability frameworks, limited attention has been paid to how these challenges are experienced and addressed in practice. In particular, empirical evidence on the operational barriers, trade-offs, and strategies adopted by urban mobility and logistics operators, when engaging in data sharing initiatives, remains limited.

To address this gap, the next section presents practitioner perspectives based on an exploratory survey conducted among stakeholders involved in data-sharing initiatives. This empirical approach provides insights into how interoperability challenges are perceived and managed in operational contexts.

4 Practitioner Perspectives: Survey Insights on Data Interoperability

4.1 Methodology

To complement the conceptual and policy analysis, an exploratory survey was conducted among practitioners involved in the design, deployment, or operation of data-sharing solutions in European urban mobility and logistics. The survey was open for 6 weeks (from the beginning of February to mid-March 2026) and gathered 55 responses from a diverse set of stakeholders, including public authorities, R&I projects, transport and logistics operators, and digital solution providers.

Respondents cover a balanced representation of operational, technical, and governance roles in data-sharing ecosystems, including data users (29.6%), policy or coordination roles (25.9%), platform or infrastructure providers (18.5%), and data providers (9.3%). An additional 16.7% selected ‘Other’, reflecting roles such as data space operators, research and innovation managers, or data processing and coordination functions.

The survey combined multiple-choice questions, Likert-scale assessments, and open-ended questions, structured around key dimensions of data interoperability, including semantic alignment, technical reuse, governance and trust, and policy and standardization.

The objective of the survey was to validate the gaps identified in the previous section, understand how they are experienced in practice, identify which limitations are perceived as most critical for enabling data interoperability, and highlight points of convergence and divergence between policy ambitions, standardization efforts, and operational realities.

4.2 Main Results

Survey results are structured around key dimensions of data interoperability, including technical interoperability, governance and trust, policy and standardization, implementation and scalability. They are as follows:

Technical Interoperability

Technical interoperability emerges as a major bottleneck in practice. 64.8% of respondents agree or strongly agree that differences in data meaning and terminology hinder interoperability. Notably, 25.9% of respondents report insufficient experience to assess this issue.

Qualitative responses further explain how these challenges manifest in practice. Respondents highlight that semantic differences lead to additional data transformation, manual alignment, and interpretation efforts, increasing costs and slowing down integration.

Beyond purely technical aspects, respondents also point to broader organizational and contextual barriers, including differences in business practices, legacy systems, and varying levels of data maturity. Some responses highlight the need for common vocabularies and standardized data models, while others stress the importance of coordination mechanisms and clear use case definition.

A minority of respondents (9.3%) disagree or strongly disagree that semantic differences significantly hinder interoperability, suggesting that in some contexts other factors, such as data availability, communication issues, or the lack of standardized protocols and governance frameworks, may be more critical.

Additional challenges encountered relate to data quality and system constraints, including missing, incomplete, or inconsistent data across sources, as well as differences in spatial and temporal resolution. In addition, legacy systems, versioning constraints, and cybersecurity considerations further increase the complexity of reusing data across systems. Limited training and support for end users, as well as highly customized and insufficiently maintained solutions, also slow adoption and reinforce fragmentation.

Semantic heterogeneity limits integration across systems, particularly in multi-actor environments. These challenges arise from differences in terminology, data structures, and definitions across stakeholders, requiring significant preprocessing. As a result, the impact of semantic interoperability varies depending on system maturity, data availability, and existing levels of standardization.

Governance and Trust

Fragmentation is also evident in how data-sharing mechanisms are implemented. When asked whether mechanisms to control data access, reuse, and governance (e.g., consent, usage rules, sovereignty) are implemented in fragmented or project-specific ways, 94.4% of respondents agree or strongly agree, while only 5.6% report insufficient experience to assess the issue.

Beyond technical issues, governance and trust-related challenges are identified as major barriers to data sharing. When asked whether a lack of clarity on data governance rules and usage rights limits trust in data sharing, 88.9% of respondents agree or strongly agree, while 11.1% remain neutral.

When analyzing operational practices, only 11.1% of respondents consider the conditions for data access, usage, and reuse to be clear and well defined. A majority (59.3%) report that they are only partially clear, while 14.8% consider them unclear or poorly defined, and another 14.8% indicate that the question is not applicable to their context.

When asked which aspects are mainly unclear, respondents identify usage rights and restrictions (30%) as the primary source of uncertainty, followed by reuse and redistribution conditions (22.5%), data ownership (20%), and liability and responsibilities (17.5%). Issues related to duration of access (5%) and other context-specific concerns are mentioned less frequently.

Responsibility allocation further illustrates this lack of structure. When asked who should bear operational responsibility for decisions made based on shared data no clear consensus emerges. Responses are distributed across multiple actors: 12% assign responsibility to data providers, 16% to data users, 24% support shared responsibility between these two, 12% to public authorities, and 6% to platforms or intermediaries. In addition, 16% of respondents indicate that responsibility should be defined case by case through contracts, while 14% consider it depends on the specific use case.

Some respondents further underline that barriers often perceived as technical barriers are closely linked to governance and organizational factors, such as unclear data access conditions, or reluctance to share commercially sensitive data.

Governance and trust remain major barriers to data interoperability, as data-sharing mechanisms are fragmented, responsibilities are unclear, and conditions for data access and reuse are insufficiently defined. This lack of clarity and alignment undermines trust and limits stakeholders' willingness to share data, highlighting that data interoperability challenges emerge not only from technical issues, but from the interaction between governance, organizational, and economic factors.

Policy and Standardization

When asked whether existing European policies effectively support data interoperability, qualitative responses (approximately 50% of the sample) highlight several recurring gaps.

Respondents first point to the lack of clear and operational implementation guidelines. While policies such as the Data Governance Act, Data Act, or EMDS define high-level principles, they often do not provide concrete guidance, tools or support for implementation in real operational contexts.

Fragmentation across governance levels also emerges as a major issue. Respondents report inconsistencies between European, national, and local regulations, as well as differences in interpretation (e.g., GDPR), leading to heterogeneous implementations across cities and countries.

In addition, uncertainty related to GDPR and data sensitivity is another recurrent concern, particularly in distinguishing between personal and operational data. While GDPR is recognized as essential, it is often perceived as restrictive or difficult to interpret in logistics contexts.

Finally, respondents highlight a lack of alignment between policies, standards, and operational practices. Existing frameworks are not sufficiently connected to technical standards, business needs, or real-world use cases.

While policy and standardization frameworks provide strategic direction, they remain insufficiently operational, fragmented across governance levels, and poorly aligned with implementation needs, which limits their effective adoption and constrains data interoperability in practice.

Implementation and Scalability

Qualitative responses highlight the importance of business incentives and value creation in making data interoperability a reality. Technical solutions alone are insufficient if data providers and users

do not perceive clear operational or economic benefits from data exchange. Without appropriate business models or incentives, stakeholders have limited interest in making their data available. Some responses further suggest that challenges often attributed to technical issues are, in fact, rooted in misaligned business interests. When incentives are aligned and value is clearly demonstrated, stakeholders are more willing to address semantic differences and related issues.

Data reuse rates further highlight these limitations. When asked how often data models or APIs are reused beyond their original purpose or context, most respondents report limited reuse: 66.7% indicate that reuse occurs only occasionally, rarely, or never, while only 3.7% report systematic reuse. In addition, 29.6% of respondents indicate that this question is not relevant in their context.

This limited reuse is reflected in the weak connection between pilot projects and standardization processes. When asked whether results from pilots or projects contributed to formal standards or specifications, only 24.1% of respondents reported such contributions, while 50% indicated that the question was not applicable to their context.

Qualitative responses provide further insight into how such contributions occur in practice. Reported contributions are predominantly indirect, taking place through participation in European research and innovation projects, pilot implementations, or ecosystem initiatives such as iSHARE, IDSA, or GAIA-X-related developments. Only a limited number of responses refer to direct contributions to formal standards (e.g., ISO or DATEX II).

While projects play an important role in developing interoperability solutions, their solutions are rarely consolidated, formalized, or translated into widely adopted standards. This reflects both a weak connection to standardization pathways and a lack of viable business models and incentives to support sustained participation.

5 Results Discussion, Conclusions and Research Avenues

This paper examines data interoperability as a key condition for coordination in European urban mobility and logistics systems. In this context, it addresses the following question: *how are data interoperability frameworks experienced in practice by European mobility and logistics stakeholders, and how do these challenges affect the realization of Physical Internet?*

To answer this question, the paper conceptualizes data interoperability as the alignment of five interdependent layers: (1) policy frameworks defining the conditions for data sharing, (2) governance mechanisms defining access, responsibilities, and incentive structures, (3) semantic standards ensuring consistent data interpretation, (4) technical interfaces enabling data exchange, and (5) operational practices enabling data reuse across services and territories, including stakeholder capabilities and incentives.

By combining European policies analysis, standardization and data spaces initiatives review, along with practitioner perspectives, it shows that data interoperability remains limited not due to a lack of solutions, but due to misalignment across these dimensions. While many European initiatives actively address these layers individually, they remain insufficiently connected to enable scalability across systems.

Survey results confirm that these challenges are experienced in practice. Semantic heterogeneity continues to limit data integration across systems, while governance and trust issues significantly affect stakeholders' willingness to share data. In parallel, policy and standardization frameworks remain insufficiently operational and not enough aligned with implementation needs, and project-based approaches struggle to scale due to weak links with standardization processes and a lack of sustainable business models.

These findings indicate that data interoperability cannot be treated as a purely technical or policy objective, but as a systemic coordination and integration matter involving technical, governance, and economic dimensions. In this sense, data interoperability is not simply an enabler, but a prerequisite for scalable and efficient decentralized coordination.

From a PI perspective, this misalignment is particularly significant, as coordination across actors, systems, and territories relies on the ability to exchange, interpret and reuse data consistently. When this is not ensured, physical operations remain fragmented and cannot scale, limiting the implementation of PI in practice.

Addressing this gap therefore requires moving from fragmented data interoperability efforts to a more systemic, PI-aligned, coordination of data practices, governance structures, and operational processes (Figure 1). In practice, achieving concrete data interoperability means enabling (1) shared semantic models that support consistent interpretation of transport and logistics data, (2) clear governance frameworks defining data access, usage rights, and responsibilities, (3) interoperable interfaces that allow data exchange and reuse beyond individual use cases, and (4) viable business models that create value and incentivize stakeholders to share and reuse data.

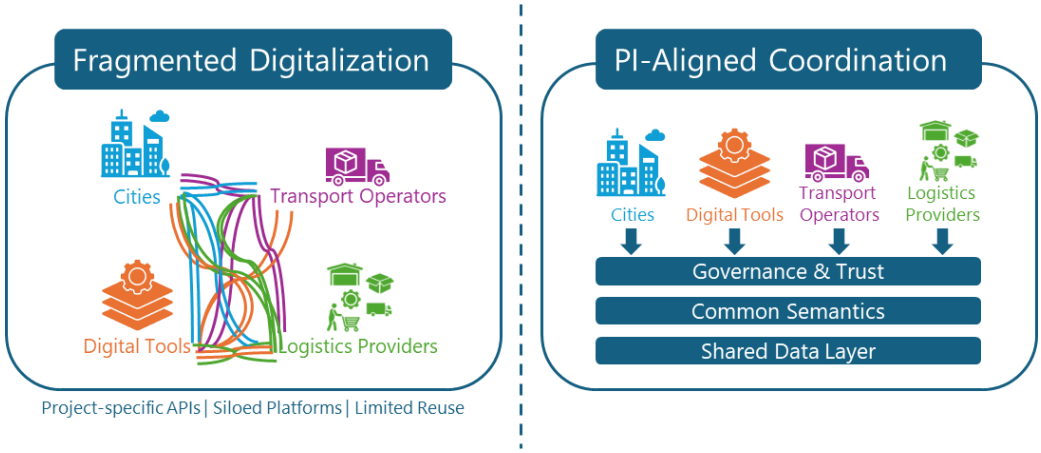


Figure 1. From fragmented digitalization to a PI-aligned coordination

As an exploratory study, this analysis presents some limitations. The survey sample size (55 respondents) remains relatively limited, although it was specifically targeted at stakeholders actively involved in urban mobility and logistics R&I or data-sharing initiatives. In addition, respondents exhibit varying levels of experience and familiarity with data interoperability, as reflected in the share of answers indicating insufficient experience to assess certain issues. This heterogeneity highlights that data interoperability challenges also relate to differences in knowledge and understanding across stakeholders.

This suggests that the current ecosystem remains in a transitional phase, characterized by uneven levels of maturity across stakeholders and limited integration between projects, standards, and policy frameworks.

Building on these findings, several priority directions emerge to achieve. First, greater efforts are needed to support semantic alignment across stakeholders, not only through standardization but also through shared interpretation frameworks and practical guidance. Second, governance mechanisms must be clarified and strengthened, including clearer definitions of roles, responsibilities, and data usage conditions, as well as the development of trusted coordination structures. In this regard, initiatives focused on stakeholder alignment and capacity building, such as the Digital Urban Freight (DUF) Community[9], can contribute to aligning practices and improving the collective understanding

[9] https://www.etp-logistics.eu/alice-urban-freight-community/

of data interoperability in operational contexts, by reducing knowledge gaps, fostering collaborative environments, and supporting the convergence of practices. Third, stronger links should be established between pilot projects, standardization bodies, and policy frameworks to ensure that operational insights are effectively consolidated and translated into widely adopted standards. Finally, the development of viable business models and incentive structures is essential to support long-term engagement in data sharing and interoperability initiatives.

Overall, this paper contributes to bridging the gap between existing interoperability approaches and the realities experienced by stakeholders. Future research could further explore mechanisms to support cross-domain data interoperability and assess how emerging governance and coordination models can be scaled across different contexts.

References

Bakri, T., Bastiaansen, H., Benmayor, A., Blaya, M., Böge, G., Cases, J.-F., Eisenrauch, A., Engh, E., Epardeau, D., Federl, S., Fernandes, E., van der Hoeven, G., King, L., Kirstein, L., Kraft, V., Lenz, G., Lilja, J., Morgan, J., Nagel, I., Öörni, R., Pretzsch, S., Scholliers, J., Schulz, H., Spijkers, N., 2023. Towards a common European mobility data space: Perspectives, recommendations and building blocks (Deliverable). PrepDSpace4Mobility.

Crainic, T.G., Montreuil, B., 2016. Physical Internet Enabled Hyperconnected City Logistics. Transp. Res. Procedia 12, 383–398.

Delaplace, J., 2024. Digitalization of the logistics - Cartography of the data sharing and exchange initiatives in Europe and recommendations for the introduction of the "Shared Data Language" developed by Geodis into this landscape. Paris, France.

European Commission. Directorate General for Informatics., 2017. New European interoperability framework: promoting seamless services and data flows for European public administrations. Publications Office, LU.

European University Institute, 2023. Creating a common European mobility data space. Publications Office, LU.

Expert Group for Urban Mobility, 2024. Data Sharing for Zero Emission Urban Logistics.

Gharehgozli, A., De Vries, H., Decrauw, S., 2019. The role of standardisation in European intermodal transportation. Marit. Bus. Rev. 4, 151–168.

Laborda, J., Dani, E., Bates, L., Ahtes, J., Touma, R., 2023. Unlocking the Future of Mobility with European Data Spaces. EIT KIC Urban Mobility S.L.

Montreuil, B., 2011. Toward a Physical Internet: meeting the global logistics sustainability grand challenge. Logist. Res. 3, 71–87.

Pan, S., Trentesaux, D., McFarlane, D., Montreuil, B., Ballot, E., Huang, G.Q., 2021. Digital interoperability in logistics and supply chain management: state-of-the-art and research avenues towards Physical Internet. Comput. Ind. 128, 103435.

Picht, P.G., 2020. Towards an Access Regime for Mobility Data. IIC - Int. Rev. Intellect. Prop. Compet. Law 51, 940–976.

Vellinga, N., Hailevich, E., 2024. The Legal Framework for Sharing Mobility Data: On the Road to an EU Mobility Data Space.

Wu, H., Li, M., Yu, C., Ouyang, Z., Lai, K., Zhao, Z., Pan, S., Wang, S., Zhong, R.Y., Kuo, Y.-H., Zhang, F., Huang, W., Shen, Z.-J.M., Ballot, E., Huang, G.Q., 2025. Towards cyber-physical internet: A systematic review, fundamental model and future perspectives. Transp. Res. Part E Logist. Transp. Rev. 197, 104051.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Place your logo here

Multimodal Itinerary Optimization for PI-Container Networks: Leveraging Inland Waterways for Sustainable Freight Consolidation

Gabriel Zambrano-Rey[1], Monica-Juliana Perez[2], Tarik Chargui[2], Yves Sallez[2] and Damien Trentesaux[2]

1. LISPEN, Arts et Métiers Institute of Technology, F-59046, Lille, France
2. LAMIH, UMR CNRS 8201, University Polytechnic Hauts-de-France, 59313, Mont Houy, France

Corresponding author: gabriel.zambrano_rey@ensam.eu

Abstract: *The steady rise in transportation demand intensifies environmental degradation, noise, congestion, and land-use conflicts. Inland waterway transport (IWT) offers a compelling response through its capacity for freight consolidation and shared asset utilization, core Physical Internet (PI or π) principles. Consequently, IWT emerges as a strategic mode for PI implementation by enabling massification of flows through interconnected hub networks. However, research gaps remain in operationalizing IWT-based PI-hub networks through optimized routing strategies. This paper proposes a multimodal container itinerary optimization approach for the Hauts-de-France IWT network, integrating the future Seine–Nord Europe Canal with road transport as a complementary mode. The optimization minimizes costs including operational expenses and environmental/social externalities, demonstrating sustainable freight consolidation in fluvial-centered PI-hub networks.*

Keywords: *Physical Internet, bimodal transportation, itinerary planning, genetic algorithm.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☒ *PI Networks,* ☐ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☐ *In-Person presentation*

1 Introduction

Transport is a cornerstone of European industry and global supply chains. In 2023, the European freight-forwarding market generated USD 58.9 billion and is projected to reach USD 86.6 billion by 2030. Road freight is expected to remain the dominant mode in the EU-27, reaching an estimated 2,027 billion tons-km by 2030, 24.5% more than in 2015, and accounting for 67.6% of all freight activity. In contrast, Inland Waterway Transportation (IWT) has maintained a modest and stable share of 1.6–1.8% (116 billion tons-km) (Eurostat, 2025). Despite this limited share, IWT contributes directly to the transport-related Sustainable Development Goals of the 2030 Agenda, and sustainability assessments must consider environmental, social, and economic dimensions (Barros et al., 2022). The EU's Sustainable and Smart Mobility Strategy therefore aims to increase inland waterway and short-sea shipping volumes by 25% by 2030 and 50% by 2050 (European Environment Agency, 2024).

Although IWT offers superior energy efficiency and lower CO_2 emissions compared with road and rail (Barros et al., 2022), its performance depends on several interrelated factors. Ports have evolved into multimodal logistics hubs with value-added activities, but their planning, yard

configuration, equipment, and barge-handling processes impose capacity limits and transshipment delays that affect reliability (Alias et al., 2022; Siswanto et al., 2023). The IWT fleet also requires careful tactical planning (SteadieSeifi et al., 2014), fleet composition, service lines, and routing must be optimized to accommodate regular and spot demand while maximizing profit (Bilegan et al., 2022), in addition to empty container repositioning (Alfandari et al., 2019). Finally, the waterway infrastructure is shaped by regulatory and operational constraints, including lock opening hours, capacity, availability (i.e., the water level), and investment needs for improved connectivity and interoperability (Binsfeld et al., 2025)

Multimodal transport integrates road, rail, water, and air in coordinated chains, and containerized multimodal logistics has become the dominant approach for improving overall efficiency. Its central challenge is to minimize logistics costs while selecting routes and modes that satisfy customer requirements (Hao & Yue, 2016). Within such systems, IWT behaves as an open subsystem, continuously interacting with its environment and with other transport modes (Calderón-Rivera et al., 2024).

While multimodal transport traditionally seeks to minimize costs through coordinated chains, the Physical Internet (PI) (Montreuil et al., 2010) provides the operational framework to transform IWT from an isolated mode into a backbone of hyperconnected, interoperable, and collaborative networks. Its modular PI-containers enable shipment consolidation and dynamic routing (Pan et al., 2021), while its systemic approach can reduce the number of circulating trucks, improve capacity usage, and lower emissions and congestion (Gallo et al., 2025). Since the PI envisions synchronized corridors and hubs, container flows must move seamlessly across intermodal services and operators. Yet, despite the alignment between PI principles and the sustainability potential of inland waterways, their integration remains insufficiently explored. Understanding how IWT can be embedded into PI-enabled multimodal networks is therefore a key research need.

To address these challenges, this paper proposes a structured roadmap designed to incrementally overcome the main operational, environmental, and organizational barriers that hinder the modal shift of container flows toward inland waterways. The roadmap consists of three successive phases: (1) an optimization phase, (2) a simulation-based optimization phase, and (3) a reactive, multi-dimensional, and decentralized phase. This paper reports on the first phase and presents preliminary results. In this phase, a Genetic Algorithm (GA) is developed to optimize π-container itineraries by coordinating road-fleet departure times with barge service schedules. The GA is designed to generate a feasible itinerary for each client demand, from origin to destination, within a shared and multi-operator network. The method minimizes a single weighted objective combining internal economic costs, environmental externalities, and deviations from requested delivery times. A real-valued encoding scheme is introduced to allocate client demand across available barge services, enabling the algorithm to manage barge capacities, promote fair access to inland waterway transport, and encourage consolidation. Routing feasibility is ensured through Physical Internet-compatible constraints, including a maximum number of allowed transshipments per container. The approach is validated on the inland waterway network of the Hauts-de-France Region, demonstrating its capacity to generate efficient, consolidation-oriented multimodal routes under synthetic demand conditions.

2 Related literature

The foundational framework for multimodal logistics planning is rooted by the Freight Routing Problem (FRP), which encompasses critical tactical decisions regarding service selection, mode integration, and commodity routing. As noted by (SteadieSeifi et al., 2014), routing problems are traditionally bifurcated into static and dynamic models; while the former assumes constant parameters over a time horizon, the latter acknowledges temporal variations in demand or

capacity. Historically, the literature has addressed these challenges through Deterministic and Exact Optimization Methods, primarily utilizing Mixed Integer Linear Programming (MILP) to minimize costs or emissions. For instance, (Chen et al., 2019) employed non-linear integer programming and the normalized normal constraint method to resolve container shortages in rail-road-waterway networks. Similarly, (Oguntola et al., 2024) demonstrated the modern integration of Machine Learning predictions as deterministic inputs for MILP models, solved via Gurobi to optimize sustainable multi-echelon supply chains. Also MILP models have been used to optimize container ship routes, as proposed by (Li & Yang, 2023), including the repositioning of empty containers, and considering navigational restrictions (i.e, water depth and bridge height) and aiming to maximize profits. The authors proposed two MILP models under splittable and unsplittable demand, applied on the Yangtze River for validation. (Shobayo et al., 2026) develop an integrated optimization model focused on cargo consolidation as a solution to the problem of suboptimal container filling rates. The results showed that consolidation can reduce overall transport costs by up to 4 % and attract as much as 42 % more container volumes to IWT. Similarly, (Ypsilantis & Zuidwijk, 2019) proposed a compact MILP formulation that aims to support the decisions on the fleet size and mix selection, the outing of the fleet over the network and the assignment of container flows to these IWT services.

Despite the precision of exact methods, the inherent volatility of global networks has necessitated advanced Methods for Handling Uncertainty.

The literature highlights three dominant paradigms: Stochastic Programming (SP), Fuzzy Programming, and Robust Optimization. (Cui et al., 2024) advanced the SP field by developing a probabilistic MILP for dynamic capacity allocation in barge-truck intermodal systems, utilizing discrete probability distributions to manage sequential booking. In contexts where historical data is sparse, (Feng et al., 2025) utilized fuzzy logic to model ambiguous transit times and prices. Conversely, to safeguard against worst-case disruptions, (Lu, 2025) introduced a multi-objective robust optimization framework that balances operational efficiency with network resilience, ensuring system stability even under severe volatility.

As the scale and complexity of these logistical networks grow, Heuristics and Metaheuristics have become indispensable for solving NP-hard combinatorial problems where exact solvers fail. Swarm intelligence and evolutionary algorithms are frequently cited as robust alternatives. A prominent example is the work of (Feng et al., 2025), who designed an Improved Whale Optimization Algorithm (IWOA) incorporating Sobol sequence initialization and variable neighborhood search to navigate hybrid uncertainties in low-carbon container transport. Furthermore, research by (Haris & Nam, 2024) and others continues to validate the efficacy of Particle Swarm Optimization (PSO) in identifying near-optimal solutions for complex routing and capacity constraints in multi-modal environments.

A significant shift in the literature is observed with the introduction of the PI paradigm, which proposes a move from rigid, proprietary supply chains to an open, global, and interconnected logistics system. Unlike traditional approaches that treat shipments as monolithic units, the PI paradigm emphasizes the modularization of freight. (Shahedi et al., 2023) illustrate this by disassembling large shipments into PI-containers, allowing for autonomous routing that exploits the residual capacity of existing vehicle flows across rail and road. This modularity facilitates horizontal collaboration; as demonstrated by (Naganawa et al., 2024) , competing firms can share centralized hubs and transport assets. By combining Fuzzy c-means clustering for hub location with Genetic Algorithms for routing, their research reported a substantial 54% reduction in CO_2 emissions, highlighting the sustainability gains of the PI.

Finally, the PI paradigm challenges the traditional reliance on centralized management through Decentralized and Agent-Based Modeling. While traditional SND models often assume a central coordinator, the PI thrives on autonomous interactions. (Sun et al., 2018) pioneered this transition by utilizing Multi-agent Simulation (MAS), where active agents—such as

autonomous PI-containers and tractors—interact under decentralized protocols. In this environment, containers identify their own transport segments and negotiate with available tractors in real-time. This departure from dedicated infrastructure and fixed scheduling has proven to drastically eliminate "deadhead" (empty) trips, achieving up to 20% fuel savings.
While PI optimization has addressed road routing and rail-road intermodality extensively, inland waterway-centred PI-hub networks integrating future canal infrastructure with joint cost-externality objectives remain unexplored — this paper addresses that gap.

3 GA-Multimodal Itinerary Optimization for PI-Container

The optimization problem is modeled as a multimodal service network design where the primary objective is to minimize a total system cost function that balances internal operational expenses and environmental externalities. The system considers a set of freight demands *D*, where each demand $d_i \in D$ must be routed from an origin to a destination within a specific time window. The IWT network is composed of a set of scheduled barge services *J*, each characterized by a fixed departure time $TD_j \in J$, arrival time $TA_j \in J$, and a capacity $Cap_j \in J$, This section describes the proposed Mixed-Encoding Genetic Algorithm (MEGA) designed to solve the intermodal container routing and modal split problem, as described in Figure 1.
The algorithm simultaneously determines the processing order of freight demands and the share of each demand volume to be assigned to inland waterway transport, with the remainder fulfilled by road.

3.1 Chromosome Encoding

Let $D = \{d_1, d_2, \dots, d_n\}$ be the set of n freight demands to be routed. Each chromosome c is a vector of n real-valued genes: $c = \{g_1, g_2, \dots, g_n\}$, $g_i \in [0,1)$. Each gene g_i encodes two decisions simultaneously through a mixed representation:

- $\sigma_i = \lfloor g_i \rfloor \in \{0,1, \dots n-1\}$, the integer part, represents the index of the demand processed at position i
- $\rho_i = g_i - \lfloor g_i \rfloor \in [0,1)$, the decimal par, represents the fluvial modal split ratio for demand d_{σ_i}.

The integer parts form a permutation of $\{0,1, \dots n-1\}$, ensuring that each demand is processed exactly once. The volume assigned to barge transport for demand d_{σ_i} is $v^f_{\sigma_i} = \lceil \rho_i * V_{\sigma_i} \rceil$ where V_{σ_i} is the total TEU volume of demand d_{σ_i}, and $v^r_{\sigma_i} = V_{\sigma_i} - v^f_{\sigma_i}$ is the residual volume fulfilled by road.

3.2 Itinerary Construction

As depicted in Figure 1, the allocation strategy prioritizes fluvial transport through a hierarchical approach: (1) allocate the target percentage of client demand to barge-only itineraries using a Dijkstra-based search within the fluvial graph; (2) if the target is not met, generate multimodal itineraries combining road and barge (road-then-barge or barge-then-road); (3) route any remaining demand entirely by road. To limit transportation time and transshipment costs, all fluvial itineraries are restricted to a single transshipment port. For each demand ($d_{\sigma i}$) feasible itineraries are constructed following this hierarchy. Barge-only paths are enumerated using a modified Dijkstra's algorithm, while road-barge combinations are explored via breadth-first search (Han et al., 2024). Service capacity constraints, departure time

feasibility, and a maximum of two transshipments are strictly enforced. Any volume that cannot be accommodated within these constraints is fulfilled by direct road service.

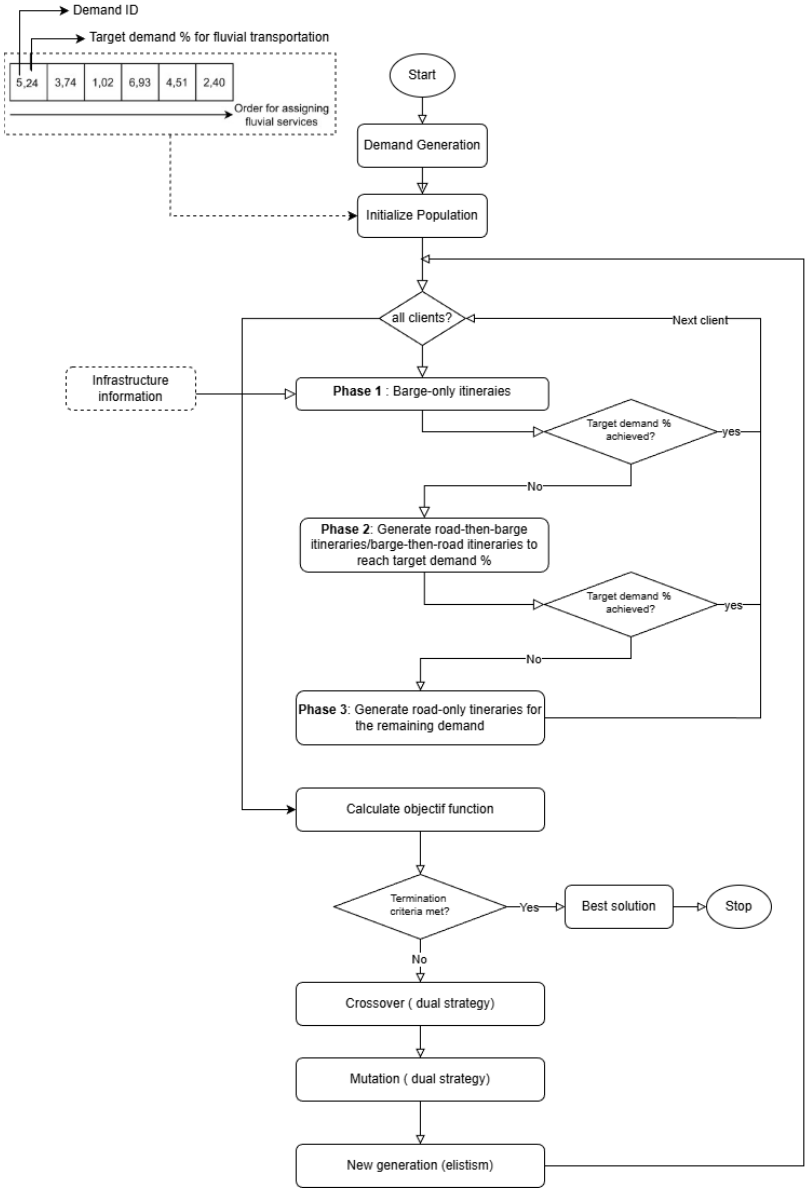


Figure 1: Flow diagram of the proposed MEGA

3.3 Fitness function

Let T_i denote the delivery tardiness and E_i the earliness for demand d_i. Four scalar fitness measures are computed for each chromosome, based on the results of itinerary construction of all demands, as follows:

- $f_1 = {}^{1}/_{\varphi * \sum T_i}$, tardiness minimisation, penalised by a cost φ
- $f_2 = {}^{1}/_{\omega * \sum E_i}$, earliness minimisation, penalised by a cost ω
- $f_3 = {}^{1}/_{(\varepsilon_1 * \varphi * \sum T_i + \varepsilon_2 * \omega * \sum E_i)}$, weighted composite earliness and tardiness costs

Transportation cost per itinerary is computed as the sum of internal costs and externalities (Hofbauer & Putz, 2020) computed with $cost_{road} = 0.4715\ €/TEU \cdot Km$ and $cost_{fluvial} = 0.2143\ €/TEU \cdot Km$ according to VNF[1].

- $f_4 = {}^{1}/_{(\sum_{road} vol_{road} * dist_{road} * cost_{road} + \sum_{fluvial} vol_{fluvial} * dist_{fluvial} * cost_{fluvial})}$

3.4 Genetic Operators

The mixed encoding requires dedicated operators that independently handle the integer permutation and the continuous decimal components

Crossover. Given parents c^1, c^2, integer parts are recombined using Partially Mapped Crossover (PMX). A segment $[cx_1, cx_2]$ is exchanged and conflicts are resolved via the induced mapping. Decimal parts are recombined using BLX-α blend crossover (Takahashi, 2016). For each position i, offspring values are drawn uniformly from $\left[\min(\rho_i^1, \rho_i^2) - \alpha * \left|\rho_i^1, \rho_i^2\right|, \max(\rho_i^1, \rho_i^2) + \alpha * \left|\rho_i^1, \rho_i^2\right|\right] \in [0,1)$. The two offspring are then reassembled as $g_i = \sigma_i + \rho_i$.

Mutation. With probability p_m, one of two operators is applied at random. With probability 0.5, a reversal mutation inverts a randomly selected sub-sequence in both integer and decimal components simultaneously, preserving their correspondence. Otherwise, a complement mutation modifies a single decimal gene as $\rho_i \leftarrow 1 - \rho_i$, thereby shifting the modal split toward the complementary mode.

Selection. Parents are selected using Stochastic Universal Sampling (SUS), which distributes N evenly spaced pointers over the cumulative fitness distribution, reducing selection variance compared to roulette-wheel sampling. An elitism mechanism preserves the top $k = 2$ individuals unchanged between generations.

4 Implementation and Experimental Setup

The proposed MEGA was implemented using Python, guided by the case study of the bimodal infrastructure of the Hauts-de-France region, representing a network of 11 PI-hubs. This includes the strategic integration of the future Canal Seine–Nord Europe (CSNE) infrastructure, which is served by 14 scheduled barge services within a one-week planning horizon.
To evaluate the impact of PI-enabled flexibility, two operational scenarios are compared:

- **Simple Scenario:** represents a conventional multimodal approach where transshipment is restricted. Consequently, if a container is dispatched by a transport mode, it must remain in the same transport mode all the way to its destination.
- **Complex Scenario:** leverages PI principles of hyperconnectivity and modularity, allowing containers to be dynamically transshipped to maximize flow massification and exploit residual barge capacity.

To generate realistic shipment patterns, client-level container demands are produced using a triangular probability distribution, where 2, 10, and 20 TEUs represent the most common

[1] https://www.vnf.fr/vnf/brochure-et-lettress/offre-logistique-fluviale-sur-le-reseau-npdc/

shipment sizes. For each client, the generated demand ranges between 1 and a maximum of 20 TEUs per shipment request. At the origin–destination (OD) level, a minimum aggregated demand of 120 TEUs is imposed to ensure sufficient flow for multimodal consolidation and fair barge-capacity evaluation. These scenarios represent a progression from fragmented, low-density shipments (many small clients) to more concentrated, high-volume clients. This variation enables us to assess the algorithm's ability to allocate barge capacity fairly across heterogeneous shippers while maintaining efficient network utilization. Furthermore, each demand request was assigned a random due-date of 1, 2, or 5 days, representing a spectrum from urgent deliveries to more flexible schedules that enable environmentally sustainable consolidation into inland waterway services.

Prior to the final comparative analysis, a systematic parameterization was conducted using the 10-TEU demand instance. This calibration phase identified the configuration of genetic operators that offered the most stable performance across varying network conditions as summarized in Table 1. Following the tuning results, the best performing configuration was selected for the production runs. This choice avoids over-specialization to specific demand patterns while maintaining high computational efficiency.

Category	*Parameter*	*Value*
GA Configuration	Population	20 individuals
	Generations	200 generations
GA Operators	Crossover	$p_c = 0.70$
	Mutation	$P_m = 0.05$ (with $\alpha = 0.50$)
Time Penalties	Tardiness (φ)	0.1
	Earliness (ω)	0.01
Weight Cost	Tardiness(ε_1)	0.5
	Earliness (ε_2)	0.5
Infrastructure	PI-Hubs	11 Ports
	Barge Services	14 Services
	Barge Capacity	60 TEUs
	Truck Capacity	1 TEUs
Demand Metrics	Scales	{2, 10, 20} TEUs
	Time Windows	{1, 2, 5} Days

Table 1: Genetic Algorithm and Network Parameters

The performance was evaluated through a comprehensive set of Key Performance Indicators (KPIs). The primary KPIs monitored in this study are defined as follows: (1) Road Cost Metrics (Internal, External & Total), which represent the operational expenses and environmental impact of trucking; (2) Barge Cost Metrics (Internal, External & Total) that measure the financial performance of the fluvial mode; (3) The Unitary Cost representing the cost per container per mode of transport; (4) The Total Tardiness measured as a financial penalty for failing to meet the delivery window; (5) The Total Earliness representing the storage cost of arriving before the scheduled window; (6) Barge Utilization (%), which measures the percentage of the total transport task handled by water; and (7) The Total System Cost.

5 Results and Discussion

This section presents the quantitative findings derived from the multi-objective optimization of the PI network. To ensure statistical validity, all experiments were executed on a workstation equipped with an Intel Core i7-11800H CPU @ 2.30GHz and 16GB of RAM, performing 30 independent runs for each scenario. The following analysis evaluates the Genetic Algorithm's performance across varying demand scales and time windows, utilizing the previously defined KPIs.

The results presented in Table 2 reveal an economic and environmental advantage for the PI-enabled "Complex" scenario over the traditional "Simple" multimodal approach. Across all

demand scales (2, 10, and 20 TEUs), the integration of PI principles leads to a reduction in unitary costs for both modes. For instance, in the 20 TEU scale, the unitary cost for road transport remains stable at approximately 29.58 € in the complex scenario, whereas it nearly doubles to 54.92 € in the simple configuration. A similar trend is observed in the barge services, where the complex scenario consistently maintains lower internal and external costs. This efficiency is a direct result of the model's ability to minimize external costs by shifting a larger portion of the freight volume toward the fluvial network, which is inherently more sustainable.

		Road			Barge		
Demand scales	Scenarios	Internal Cost	External Cost	Unitary Cost	Internal Cost	External Cost	Unitary Cost
2 TEUs	Complex	110503,02	28123,78	29,60	14410,27	2353,52	7,63
	Simple	162575,44	41376,57	54,82	20419,39	3334,94	21,42
10 TEUs	Complex	120989,53	30792,67	29,62	8351,07	1363,92	4,14
	Simple	181652,49	46231,81	55,31	22535,37	3680,53	23,81
20 TEUs	Complex	158889,98	40438,60	29,58	15939,03	2603,20	6,04
	Simple	241839,79	61549,89	54,92	24356,53	3977,97	20,55

Table 2: Comparative analysis of road and barge cost metrics across demand scales and operational scenarios.

The systemic impact of these routing strategies is further detailed in Table 3, which highlights the superior scalability of the PI-hub network. A significant finding is the correlation between demand volume and barge utilization. While the "Simple" scenario struggles to consolidate flows, the "Complex" scenario achieves a peak barge utilization of 69.85% at the 20 TEU scale, compared to the 51.22% reached by traditional methods. This confirms that allowing transshipments between different barge services is essential for effective flow massification. Consequently, the Total System Cost is consisted lower in the complex scenario.

Beyond overall cost reduction, the Barge Utilization (%) metrics provided in Table 3 offer critical insights into fleet dimensioning for the Hauts-de-France case study. Currently, the model operates with a standard barge capacity of 60 TEUs. However, the utilization rates—ranging from 50.21% to 69.85% in the complex scenario—suggest that the network could be further optimized by adjusting the fleet composition. The data indicates that while 60-TEU barges are suitable for high-massification corridors, a shift toward smaller, more agile vessels could be more efficient for lower-density PI-hubs. Implementing smaller barges would likely increase the load factor per departure and reduce the high Earliness Total Cost identified in Table 3, as more frequent services with smaller vessels could better align with diverse customer due dates. Equally important, the remaining capacity can be allocated to spot clients (partial or full), who typically pay higher rates because their shipments are arranged at short notice (Bilegan et al., 2022).

Demand scales	Scenarios	Total cost Road	Total Barge cost	Tardiness Total Cost	Earliness Total Cost	Barge Utilization (%)	Total System cost
2 TEUs	Complex	138626,80	16763,79	14561,28	535603,62	50,21	705555,50
	Simple	203952,01	23754,33	23745,50	525366,21	40,85	776818,05
10 TEUs	Complex	151782,20	9714,98	18337,10	559294,69	54,21	739128,98
	Simple	227884,31	26215,90	23878,52	551158,85	44,85	829137,57
20 TEUs	Complex	199328,58	18542,23	39288,93	770982,94	69,85	1028142,69
	Simple	303389,68	28334,49	43519,85	759700,96	51,22	1134944,98

Table 3: Global system performance and PI KPIs for bimodal itinerary optimization.

Finally, the disproportionately high storage-related costs shown in Table 3 point toward a notable rigidity in the current model. These figures represent the financial burden of containers arriving at hubs well before their scheduled windows, likely due to the fixed nature of the 14 scheduled barge services and the absence of granular transshipment processing times. This indicates that while the Genetic Algorithm successfully identifies cost-efficient itineraries, the model serves primarily as a preventive strategy requiring further adjustments.

6 Conclusions and Roadmap for Optimizing Inland Waterway Container Transport

This paper addressed the itinerary optimization of containers in a bi-modal system combining inland waterway transport (IWT) and road transport, where road serves as a complementary mode. The prioritization of IWT is motivated by its well-documented advantages in terms of cost efficiency and environmental impact compared to road-only alternatives. Physical Internet principles were partially incorporated, particularly to limit unnecessary transshipment operations, with broader integration planned in subsequent phases. The proposed genetic algorithm adopts a real-valued chromosome encoding that ensures equitable access to waterway services across all clients. Several scenarios involving partial demand, where freight flows are split between IWT and road services, were analyzed to assess the robustness of the approach under the configuration of the HDF fluvial network. Results on the proposed case study demonstrate that IWT-based itineraries can deliver competitive service quality while significantly reducing costs and external impacts relative to pure road solutions. Nonetheless, this work represents the first phase of a three-stage research roadmap of increasing complexity, with future phases progressively integrating more realistic operational constraints and greater autonomy in routing decisions, as follows:

- Phase 1. *Static Optimization with Enhanced Operational Constraints*: The first phase will continue the development of the static optimization model that generates feasible multimodal itineraries while integrating key operational constraints. Beyond the core routing logic, this phase will incorporate minimum barge-load constraints, barge capacity sharing, and a multi-criteria objective (economic, environmental, service). The output is a set of optimized container paths under deterministic conditions.
- Phase 2. *Simulation-Based Optimization for Realistic Performance Assessment*: The second phase couples the optimization model with a discrete-event simulation to evaluate itineraries under realistic network dynamics. The simulation accounts for water-level variability, lock schedules and delays, truck fleet allocation per port, and drivers' hours-of-service constraints. By connecting optimization and simulation, this phase produces realistic arrival times and KPIs, enabling iterative refinement of routing plans.
- Phase 3. *Decentralized Reactive Decision-Making*: The third phase introduces decentralized mechanisms that react to operational disturbances, such as late barge or truck arrivals. When a connection is at risk, the system evaluates alternatives, delaying departures, rerouting, reallocating resources, to ensure delivery while balancing economic, service-level, environmental, and societal criteria. This phase enables real-time PI-compatible decision-making under uncertainty.

7 References

Alfandari, L., Davidović, T., Furini, F., Ljubić, I., Maraš, V., & Martin, S. (2019). Tighter MIP models for Barge Container Ship Routing. *Omega*, *82*, 38–54. https://doi.org/10.1016/j.omega.2017.12.002

Alias, C., Felde, J. Z., Severin, S., & Alarcon Olalla, F. E. (2022). Measuring the Performance Impact of a Decentralized Waterborne Container Transportation Service on Inland Waterway Hubs in Western Germany. *2022 IEEE International Conference on Industrial Engineering and Engineering Management (IEEM)*, 0965–0969. https://doi.org/10.1109/IEEM55944.2022.9989632

Barros, B. R. C. D., Carvalho, E. B. D., & Brasil Junior, A. C. P. (2022). Inland waterway transport and the 2030 agenda: Taxonomy of sustainability issues. *Cleaner Engineering and Technology*, *8*, 100462. https://doi.org/10.1016/j.clet.2022.100462

Bilegan, I. C., Crainic, T. G., & Wang, Y. (2022). Scheduled service network design with revenue management considerations and an intermodal barge transportation illustration. *European Journal of Operational Research*, *300*(1), 164–177. https://doi.org/10.1016/j.ejor.2021.07.032

Binsfeld, T., Hamdan, S., Jouini, O., & Gast, J. (2025). On the optimization of green multimodal transportation: A case study of the West German canal system. *Annals of Operations Research*, *351*(1), 667–726. https://doi.org/10.1007/s10479-024-06075-5

Calderón-Rivera, N., Bartusevičienė, I., & Ballini, F. (2024). Sustainable development of inland waterways transport: A review. *Journal of Shipping and Trade*, *9*(1), 3. https://doi.org/10.1186/s41072-023-00162-9

Chen, D., Zhang, Y., Gao, L., & Thompson, R. G. (2019). Optimizing Multimodal Transportation Routes Considering Container Use. *Sustainability*, *11*(19). https://doi.org/10.3390/su11195320

Cui, Y., Bilegan, I. C., Duchenne, E., & Duvivier, D. (2024). Demand rerouting mechanisms with revenue management for intermodal barge transportation networks. *Transportmetrica B: Transport Dynamics*, *12*(1), 2416182. https://doi.org/10.1080/21680566.2024.2416182

European Environment Agency. (2024, October 10). *Sustainability of Europe's mobility systems 2024, Web report no. 01/2024*. https://www.eea.europa.eu/en/analysis/publications/sustainability-of-europes-mobility-systems/freight-transport-activity

Eurostat. (2025, March). *Freight transport statistics—Modal split*. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Freight_transport_statistics_-_modal_split

Feng, F., Zheng, F., Zhang, Z., & Wang, L. (2025). A Framework for Low-Carbon Container Multimodal Transport Route Optimization Under Hybrid Uncertainty: Model and Case Study. *Applied Sciences*, *15*(12). https://doi.org/10.3390/app15126894

Gallo, F., Shahedi, A., Di Febbraro, A., Saeednia, M., & Sacco, N. (2025). Robustness evaluation of a physical internet-based intermodal logistic network. *Transportation Engineering*, *20*, 100333. https://doi.org/10.1016/j.treng.2025.100333

Han, Z., Sun, H., Huang, J., Xu, J., Tang, Y., & Liu, X. (2024). Path Planning Algorithms for Smart Parking: Review and Prospects. *World Electric Vehicle Journal*, *15*(7). https://doi.org/10.3390/wevj15070322

Hao, C., & Yue, Y. (2016). Optimization on Combination of Transport Routes and Modes on Dynamic Programming for a Container Multimodal Transport System. *Procedia Engineering, Green Intelligent Transportation System and Safety*, *137*, 382–390. https://doi.org/10.1016/j.proeng.2016.01.272

Haris, M., & Nam, H. (2024). Path Planning Optimization of Smart Vehicle With Fast Converging Distance-Dependent PSO Algorithm. *IEEE Open Journal of Intelligent Transportation Systems*, *5*, 726–739. https://doi.org/10.1109/OJITS.2024.3486155
Hofbauer, F., & Putz, L.-M. (2020). External Costs in Inland Waterway Transport: An Analysis of External Cost Categories and Calculation Methods. *Sustainability*, *12*(14). https://doi.org/10.3390/su12145874
Li, D.-C., & Yang, H.-L. (2023). Ship routing in inland waterway liner transportation with foldable and standard empty containers repositioning. *Ocean Engineering*, *285*, 115391. https://doi.org/10.1016/j.oceaneng.2023.115391
Lu, Y. (2025). A multimodal deep reinforcement learning approach for IoT-driven adaptive scheduling and robustness optimization in global logistics networks. *Scientific Reports*, *15*(1), 25195. https://doi.org/10.1038/s41598-025-10512-1
Montreuil, B., Meller, R. D., & Ballot, E. (2010). *Towards a Physical Internet: The Impact on Logistics Facilities and Material Handling Systems Design and Innovation*.
Naganawa, H., Hirata, E., Firdausiyah, N., & Thompson, R. G. (2024). Logistics Hub and Route Optimization in the Physical Internet Paradigm. *Logistics*, *8*(2). https://doi.org/10.3390/logistics8020037
Oguntola, I. O., Ülkü, M. A., Saif, A., & Engau, A. (2024). On the value of shipment consolidation and machine learning techniques for the optimal design of a multimodal logistics network. *INFOR: Information Systems and Operational Research*, *62*(1), 1–52. https://doi.org/10.1080/03155986.2023.2202079
Pan, S., Trentesaux, D., McFarlane, D., Montreuil, B., Ballot, E., & Huang, G. Q. (2021). Digital interoperability in logistics and supply chain management: State-of-the-art and research avenues towards Physical Internet. *Computers in Industry*, *128*, 103435. https://doi.org/10.1016/j.compind.2021.103435
Shahedi, A., Gallo, F., Saeednia, M., & Sacco, N. (2023). Lead-time-based freight routing in multi-modal networks considering the Physical Internet. *Journal of Supply Chain Management Science*, *4*(3–4), 61–80. https://doi.org/10.59490/jscms.2023.7183
Shobayo, P., Yang, C., Bedoya-Maya, F., Nicolet, A., van Hassel, E., Atasoy, B., & Vanelslander, T. (2026). Improving a cargo consolidation strategy for port-hinterland transport via inland waterways. *Journal of Transport Geography*, *130*, 104477. https://doi.org/10.1016/j.jtrangeo.2025.104477
Siswanto, N., Kusumawati, M., Baihaqy, A. R., Wiratno, S. E., & Sarker, R. (2023). New Approach for Evaluating Berth Allocation Procedures Using Discrete Event Simulation to Reduce Total Port Handling Costs. *Operations and Supply Chain Management: An International Journal*, 535–553. https://doi.org/10.31387/oscm0550410
SteadieSeifi, M., Dellaert, N. P., Nuijten, W., Van Woensel, T., & Raoufi, R. (2014). Multimodal freight transportation planning: A literature review. *European Journal of Operational Research*, *233*(1), 1–15. https://doi.org/10.1016/j.ejor.2013.06.055
Sun, Y., Zhang, C., Dong, K., & Lang, M. (2018). Multiagent Modelling and Simulation of a Physical Internet Enabled Rail-Road Intermodal Transport System. *Urban Rail Transit*, *4*(3), 141–154. https://doi.org/10.1007/s40864-018-0086-4
Takahashi, R. (2016). Empirical evaluation of changing crossover operators to solve function optimization problems. *2016 IEEE Symposium Series on Computational Intelligence (SSCI)*, 1–10. https://doi.org/10.1109/SSCI.2016.7850141
Ypsilantis, P., & Zuidwijk, R. (2019). Collaborative Fleet Deployment and Routing for Sustainable Transport. *Sustainability*, *11*(20). https://doi.org/10.3390/su11205666

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE CargoBeamer

From Digital to Physical Internet: CBoXX-enabled LTL-on-Rail Demonstrators with Automated PI Nodes and Slot-based Booking

Hans-Jürgen Weidemann
CargoBeamer AG, Kreuzstr. 12, 04103 Leipzig, Germany
Corresponding author: hjweidemann@cargobeamer.com

Abstract: This paper presents an industrially grounded Physical Internet (PI) system design to enable less-than-truckload (LTL) unit cargo on rail. The approach couples (i) an interoperable rail load unit optimized for pallets, cage boxes and other unit cargo (CBoXX), (ii) automated, electrified PI nodes using GateModules for fast road/rail transhipment, and (iii) a digital logistics layer providing slot-based booking, unit-level visibility and governance. We describe the end-to-end operating concept and the planned real-world demonstrator on the Kaldenkirchen (DE)–Domodossola (IT) corridor (June/July 2027) within HOLOGISTICS. The demonstrator is designed as a repeatable daily circulation with fixed terminal windows, standard operating procedures, scan-based event capture and exception playbooks. We outline the research methodology, success criteria and expected contributions to PI adoption pathways for rail-based unit cargo networks.

Keywords: Physical Internet; Unit cargo; Less-than-truckload (LTL); Rail freight; Interoperability; PI nodes; Digital logistics platform; CBoXX; GateModule; Slot-based booking

Acknowledgments: Research for this paper was carried out within the framework of the EU project 'HOLOGISTICS - Holistically Optimised Logistics Operations and Green Innovative Solutions for Transport Cargo Sustainably' (GA – 101202556).

1. Introduction

1.1. Motivation and problem statement

Unit cargo (e.g., EU pallets, cage boxes, standardized load carriers and mixed parcels) dominates European merchandise flows, yet LTL volumes remain largely road-based. Rail freight is often unattractive for fragmented shipments due to system-level incompatibilities: specialized loading units, slow terminal operations, and limited flexibility compared to road groupage networks, finally leading to higher transport cost and/or unattractive transport timing per load unit. This paper addresses the question: how can rail become operationally simple, scalable and competitive for LTL unit cargo while aligning with PI principles?

1.2. Contribution and scope

We contribute (i) a PI-aligned system architecture for LTL-on-rail combining modular equipment, automated PI nodes and a digital booking/visibility layer; (ii) an operational concept for repeatable & fast takt-based rail services; and (iii) a demonstrator design for the Kaldenkirchen–Domodossola corridor, including governance, event chains and success criteria.

1.3. Paper structure

Section 2 reviews related work and gaps, summarizing PI core elements and deficits of current rail cargo with respect to unit cargo transport. Section 3 describes the proposed PI system

resulting from a structured design using evaluation metrics to minimize gaps described in Section 2. Section 4 details the demonstrator design. Section 5 concludes and outlines future work.

2. Related Work and Research Gap

2.1. Physical Internet principles for freight networks

The Physical Internet ("PI","π") vision emphasizes modular, interoperable logistics units, open networks of PI nodes, and a digital layer enabling routing, booking and visibility. For rail freight, PI adoption requires both physical interoperability (standardized load units and handling interfaces) and digital interoperability (shared data models, event capture and governance). From ALICE (2022) and Landschuetze (2015), eight "PI" key components are extracted:

1. **Standardized, modular load units "MLU" and "π-containers":** A family of interoperable, modular physical "packets" (from large container scale down to smaller boxes) designed to be combined/split, handled by many parties, and moved seamlessly across modes and networks.
2. **Open, shared logistics networks ("network of networks"):** Logistics resources (capacity, assets, facilities) are pooled across companies so shipments can traverse multiple operators' networks as one connected system, rather than staying inside isolated silos.
3. **PI nodes and PI hubs (π-nodes / π-hubs) as standardized transfer points:** Physical locations (terminals, hubs, DCs, ports, etc.) evolve into Physical Internet nodes where operations are standardized, and where modular units can be efficiently transferred between incoming/outgoing transport services (including multimodal transfers).
4. **Standardized interfaces (physical + informational):** Common physical handling/locking/stacking compatibility and common information interfaces so units can be identified, tracked, and processed consistently across equipment, facilities, and IT systems.
5. **PI protocols (handover + routing rules):** Agreed "protocols" that govern how units are handed over node-to-node and how they are routed through networks while maintaining network performance/integrity, analogous in spirit to Internet routing/transport protocols.
6. **Digitally visible, usable node services (plan–book–execute):** Services at PI nodes are visible and digitally accessible for planning, booking/transactions, execution, and information sharing (think: capacity, slots, handling services as "APIs").
7. **Governance & rules for participation (data, access, neutrality, clearing):** A governance layer that defines roles, access rights, data-sharing boundaries, and decision rights so the PI can function as a "network of networks" without collapsing under trust/competition issues.
8. **Performance & sustainability targets with measurement (KPIs):** PI is goal-driven (efficiency, resilience, emissions/energy). It assumes systematic measurement and improvement using shared KPIs.

2.2. LTL barriers to rail – a customer's perspective and resulting design targets

Road-based groupage networks achieve flexibility through frequent departures, consolidation, and standardized handling of pallets and cages. Shifting LTL to rail typically increases process steps and dependencies (pre-/post-carriage to terminals, crane operations, specialized equipment), reducing speed and reliability. Existing rail solutions often optimize either the long-haul rail leg or terminal handling, but rarely integrate interoperable unit-cargo equipment, high-throughput automated nodes, and a booking/visibility layer into a coherent PI system.
Table 1 summarizes today's deficit of rail cargo with respect to the 8 π components described above from a customer's perspective; Table 2 highlights the areas and shortcomings that generally undermine the competitiveness of rail freight transport today:

Table 1: Deficit of rail cargo with respect to the 8 core-π-elements

#	PI element	Deficits today (LTL using rail / intermodal) from a customers perspective	Resulting design targets for PI-enabled LTL-on-rail
1	Standardized modular load units (π-containers / MLUs)	Pallets are not a true "intermodal packet." A few EU pallets typically require ad-hoc consolidation into a container/swap body; packaging/stacking constraints vary; higher damage risk due to re-handling and re-palletizing.	Small-shipper-friendly modular unitization: pallet-compatible modular boxes/cages that can be mechanically & robototized consolidated into rail-capable units (swap body/container) with minimal repacking. Design targets: fewer touches, lower damage, better cube/weight utilization, easy reverse logistics for empties.
2	Open, shared networks (network of networks / pooling)	Customers are "stuck" with one forwarder's rail option set. If a corridor is full/disrupted, alternatives are limited; pooling across forwarders & operators is weak; small volumes get low priority.	Capacity pooling across forwarders, slot-based booking system and corridors; high priority on rail quality & frequency; tendering of linehaul services.
3	PI nodes & π-hubs (standardized transfer points)	Terminals + consolidation hubs behave differently by country/provider: cut-offs, appointment rules, handling quality, operating hours. For small shipments, every hub adds delay variance and "lost in hub" risk.	"Rail LTL routers" with standardized service levels: designated PI-ready consolidation hubs and terminals with defined processes (check-in/out, consolidation windows, damage prevention, exception handling).
4	Standardized interfaces (physical + informational)	Physical compatibility issues (equipment, securing, load carriers) and IT fragmentation (different labels, portals, EDI/API). Tracking milestones differ; cross-border partner switches cause visibility gaps.	Plug-and-play interfaces: one label/ID scheme, one milestone model, standard event semantics (picked up, inducted, hub scan, terminal gate-in/out, train departed/arrived, out for delivery, delivered).
5	PI protocols (handover + routing rules)	Handover is contract- and email-driven; routing is opaque; customers can't easily trade cost vs. speed vs. CO_2; re-routing during disruption is slow and manual.	Machine-executable handover & routing rules: standardized handover protocol between customer - truck - terminal - rail - terminal - truck - customer.
6	Digitally visible node services (plan–book–execute)	Booking rail capacity for a few pallets is not "retail-like." Customers face multiple systems and unclear feasibility until late; appointment booking is painful; ETA precision is weak.	Single service catalog: instant price + feasibility check for 1–10 pallets, Digital slot/appointment orchestration across hubs/terminals; exception workflows (rebook slot, change delivery window) built into the product.
7	Governance (data rights, neutrality, clearing, liability)	Multi-actor chains create blame-shifting and slow claims. Data sharing is limited; customers don't know who is responsible at each leg; cross-border liability boundaries feel opaque.	One accountable contract + neutral governance: single point of responsibility for the customer (even if execution is multi-operator). Clear liability/claims SLA, standardized evidence package (photos, scans, timestamps).
8	KPIs & sustainability measurement	Customers get inconsistent service KPIs and hard-to-audit CO_2 numbers; reliability is not expressed as a probability; performance varies by lane but is not transparent.	Customer-relevant KPIs by lane/product tier: on-time probability, lead-time distribution (not just average), damage rate, claim cycle time, and shipment-level CO_2e with consistent allocation logic.

Table 2: Areas and shortcomings that generally undermine the competitiveness of rail freight

#	General	General deficits today (LTL using rail / intermodal) from a customers perspective	Resulting design targets for PI-enabled LTL-on-rail
9	Cost	Rail transport cost plus terminal hub cost plus first&last mile often exceed cost of pure road transport	Competitive pricing per modular box (one transparent price from shipper to receiver)
10	Complexity	Rail transport needs synchronization with different service providers (first&last mile, rail operator, terminal timings)	One-stop-shop for information, booking, tracing, clearing. Transport service from the shipper's loading dock to the final destination
11	Quality	Rail freight transport often has a punctuality rate of less than 70%	Very high quality on punctuality & route frequency
12	Timing	Cross-border rail freight transport in the EU often has only one departure per day and route, and an unacceptable total transit time of 2–4 days from the shipper to the destination	Train speed 120km/h, several departures per day & route, fast first & last mile services synchronized with train timing

3. Research Methodology and Demonstrator Design

3.1. Methodological framework

The research follows a design-and-validate approach: (i) system synthesis based on design targets resulting from today's gaps to PI principles and general rail cargo shortcomings (see Tables 1,2) ; (ii) process modeling of current vs. target LTL chains (see Fig. 1); (iii) KPI

evaluation of potential LTL chains vs. road; and (iv) validation through a real-world LTL-on-rail demonstrator with defined success criteria and KPI verification measurement.

3.2. System Design based on Design Targets

3.2.1. Design target #1 “Shipper-friendly modular unitization”

The proposed LTL-on-rail design focuses on industrial boxes in standardized sizes 1.2 m x 1.0 m and 0.8 m x 1.2 m (see Fig 1). Customers do not need any equipment except a forklift. A key objective of the research project is to provide end customers with complete flexibility regarding shipment size, while ensuring cost-effectiveness and a sustainable logistics chain.

3.2.2. Design target #2 “Open, shared networks”

CargoBeamer AG designed a rail container with capacity for 104 MLU (1.2 m x 0.8 m base) such as cage-boxes, the “CBoXX”, which will serve as a “micro-π-hub” allowing pooling of shipments from several customers.



Figure 1: The “CBoXX” rail container results from the design targets listed in tables 1,2

3.2.3. Design targets #3 and #4: Standardized transfer points & physical interfaces

The proposed LTL on rail solution focuses on standardized interface consisting of direct cross-docking of small modular loading units between first&last mile eTrucks and the CBoXX. At present, standard forklifts are used. The transfer points between road & rail are standardized GateModules, which automate the CBoXX unloading/loading onto/from railcars, a fully automated and energy optimized process.

3.2.4. Design targets #5 and #6 Protocols, Digital Services

The analysis, synthesis, and implementation of machine-executable handover and routing rules—specifically, a standardized handover protocol between customer–truck–terminal–rail–terminal–truck–customer—are planned for the demonstrator 2027. In the 2nd half of 2026 we will analyze and utilize results from completed and ongoing EU projects: a) Alignment and synergies with Shift2Rail’s Flagship Area 5 DAC and Seamless Freight workstreams, which are expected to further enhance yard and hub operations; b) HOLOGISTICS Work Package 4 on standardized data sharing between intermodal operators and dynamic optimization of

internal routes, reducing travel times and fuel/electricity consumption. Results will be applied to the demonstrator 2027.

3.2.5. Design target #8: Sustainability & Customer KPI measurement

Table 2 lists the key shortcomings of today's rail freight transport, particularly for time-sensitive, small-lot commercial goods. Both in the analysis system design as well as in the booking and logistics management software foreseen for the demonstrator, we will calculate and measure KPIs from the customer's perspective for the entire transport chain, specifically (i) full transport costs per MLU from shipper to destination, (ii) emissions per MLU, (iii) transit time from shipper to end customer (compared to road transport), (iv) energy consumption, (v) quality in terms of schedule adherence (delays > 2 hours as a percentage of shipments)

3.2.6. Design targets #9-#12: Cost, Complexity, Quality, Timing

To minimize transportation and handling costs in the MLU-LTL transport chain, particular emphasis should be placed on leveraging the system advantages of rail over road, namely (i) high load capacity and (ii) greater volume capacity.

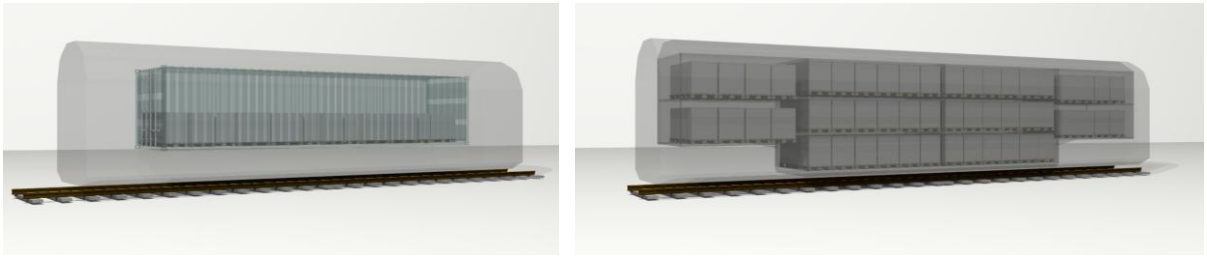


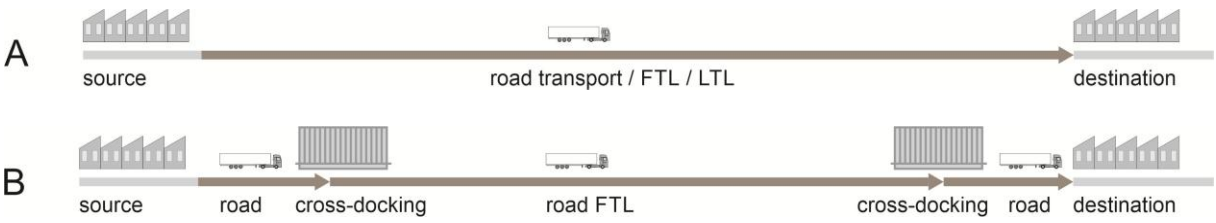
Figure 2: Volume efficiency given by net MLU volumes vs. maximum allowed geometrical rail profile. Left: ISO-container with 24 non-stackable EU-pallets. Right: CBoXX with 104 EU-pallets (Table 3)

Table 3: CBoXX KPI vs. standard 40ft-ISO-Container: Significantly increased capacity

	40ft. ISO-container	CBoXX
Volume efficiency*	20%	70%
Load capacity, EU-pallet, non stackable	26	104
Load capacity, Industry Box 1.2 x 1 x 1.1 m (Width/Length/Height)	52	84
Max. load weight [t]	26,7	33 (2026 version)
		44 (2028 planned)
**Portion of loaded EU-pallet volumes vs. theoretical max. railcar gabarit profile x 20m*		

3.3. Process modeling: Today's vs. Future FTL/LTL transport chains

Starting with traditional road transport using diesel-powered semi-trucks, we model and analyze six different scenarios (A) through (F) for the transition from full-truckload (FTL) road transport to flexible, low-emission, and competitive less-than-truckload (LTL) transport:



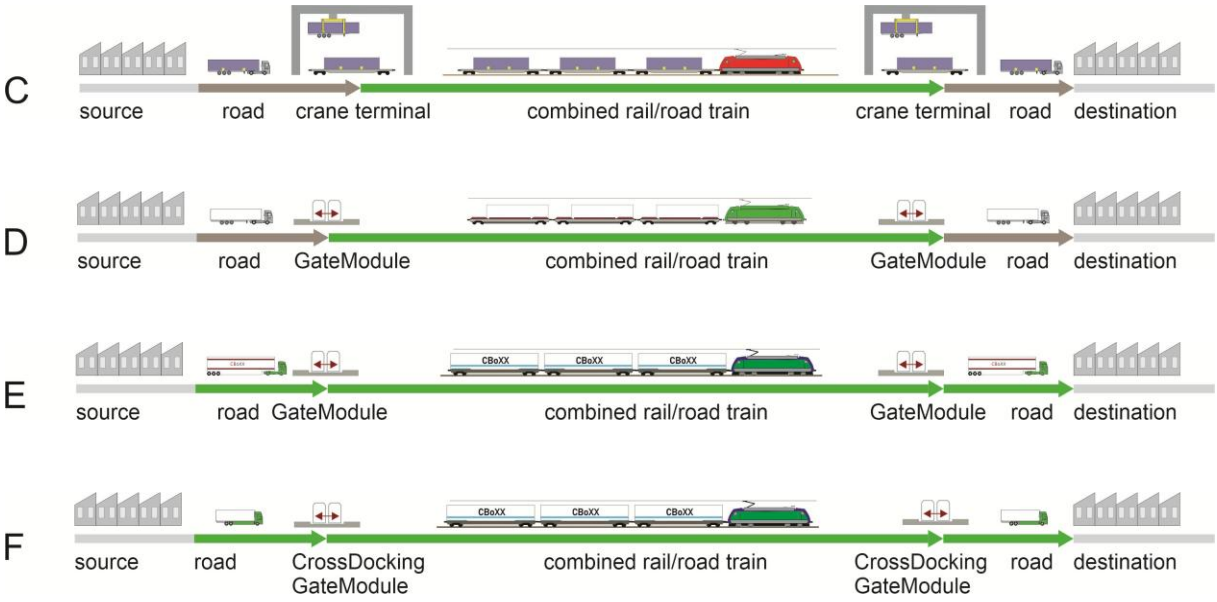


Figure 3: Today's conventional road & Combined Transport chains (A,B,C) and potential future LTL on rail transport chains (D,E,F). Electrified segments marked in green color.

3.4. KPI evaluation of potential LTL chains vs. Road

Given the transport chain scenario's A – F, we evaluated KPI's for LTL shipments such as timing, cost, energy & emissions per multimodal logistics unit (MLU) transport.

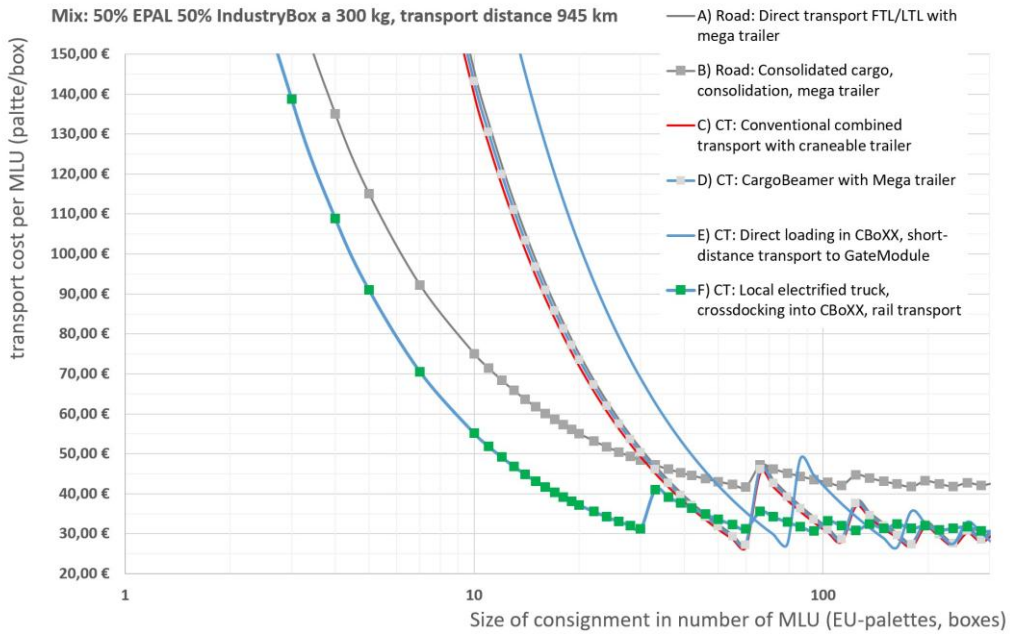


Figure 4: Specific cost per modular load unit (MLU). Note: The curves for scenarios A, C, and D are very close together. They reach their minima at 60, 120, 180, ... MLU, since exactly one Mega-trailer is filled for each of these values when, as in this parameter set assumed, 50% industrial boxes and 50% palletized goods are allowed to be transported stacked in two layers. The CBoXX scenario F reaches minima at multiples of 30 MLU, which corresponds precisely to the capacity of 12-ton eTrucks for this MLU mix. Key scenario parameters: Road transport

distance 945 km, rail distance plus first/last mile 990 km. FTL road cost €1.50/km, rail traction cost €24/km.

This study focuses on LTL shipments, which account for a large but not precisely defined share of road freight transport in Europe. All producers / shippers with irregular quantities of MLU over long distances in Europe (for example, 5 EU pallets on one day, 43 on the next, and 17 on the third day, and so on), need to rely on freight forwarders that can consolidate these loads with shipments from other shippers and stack them into FTL ("full truck load") semi-trailers that are as full as possible for the main route. Figure 4 shows the key bottom-up calculated transport costs excluding overheads for varying quantities of MLUs given a specific scenario.
It becomes apparent that, if we exclude the currently unavailable "F" scenario involving e-trucks and CBoXX, scenario *"B: Road transport with consolidation in warehouses"* represents the cost optimum for smaller shipment volumes. This result is logical, since, running a long-distance heavy vehicle with less MLU significantly increases cost per MLU.
Scenarios A,C and D, where one Mega-Trailer is filled at shipper's site, result with cost minima depending on stackability and size of the MLU, in the scenario depicted in Fig. 4 we assumed a mix of EU-palletized goods and Industry Boxes with a max. stacking level of 2, resulting in max. 60 MLU per Mega-Trailer. Correspondingly, these scenario's show minima of specific cost at 60, 120, 180, ... MLU.
Given the CBoXX's high volume and weight capacity and based on the analyzed route the scenario *"F: First and last mile with eTruck, consolidation into CBoXX, automated loading of CBoXX using GateModule"* is significantly more cost-effective than other scenarios for these LTL shipment volumes up to approximately 30 MLU.
Depending on the type and mix of MLUs, the route, and key cost assumptions, the specific costs curves per MLU & consignment size vary; however, Scenario F shows increasingly attractive specific costs for MLU volumes of up to approximately 30–40 units starting at certain transport distances. We therefore anticipate that this solution will be attractive for the EU transport market in the future.

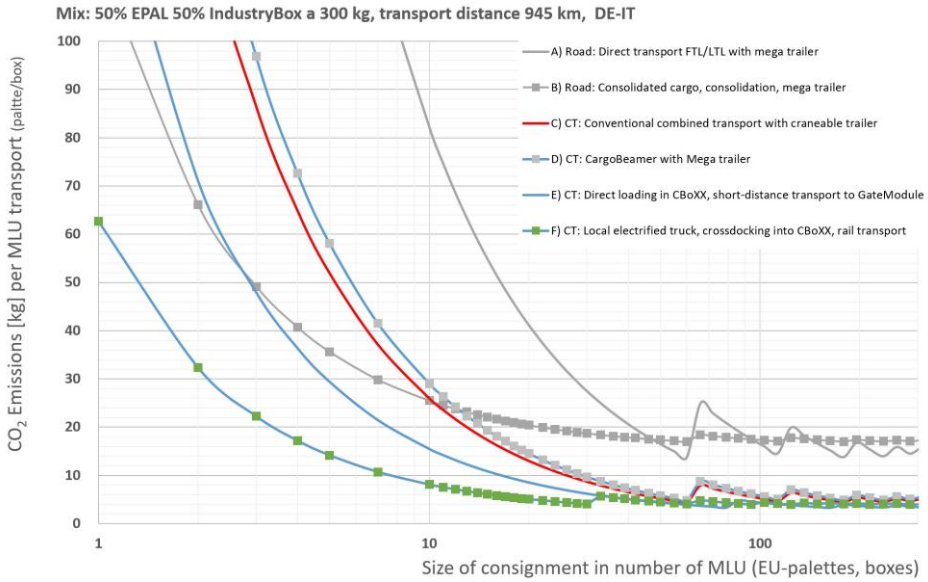


Figure 5: Scenario-based specific CO_2-emissions per MLU (example, route-specific)

Figure 5 shows the direct CO_2 emissions from transporting MLU's according to the scenarios presented on a sample route (western Germany to northern Italy). Here, too, it is evident that,

due to the environmental and energy efficiency advantages of rail freight transport combined with the high capacity of the CBoXX, the proposed scenario “F” has the lowest emissions per MLU transported.

3.5. Proposed PI System: CBoXX-enabled LTL-on-Rail

3.5.1. System overview

The proposed system accordingly follows scenario “F” (Figure 3) and couples three layers: (1) modular transportation equipment for unit cargo (electrified first/last mile using available e-Trucks e.g. of 12t load size, the CBoXX and compatible railcars), (2) automated PI nodes (GateModules) enabling fast, electrified road/rail transfer, and (3) a digital logistics platform enabling slot-based booking, unit-level IDs, event capture and KPI monitoring.

3.5.2. CBoXX as interoperable rail load unit for unit cargo

CBoXX (Fig. 1) is designed to carry MLU unit cargo such as EU pallets and cage boxes with defined packaging and handling rules. It acts as a micro PI hub by enabling pallet-/box-level consolidation on rail. The design leverages compartment-level identification and supports condition/status monitoring (e.g., temperature, humidity, vibration) via sensors and identifiers linked to a database.

3.5.3. Automated PI nodes: GateModules for high-throughput transhipment

GateModules are automated road/rail loading stations enabling parallelized handling of complete trainsets by simultaneously loading/unloading all railcars of a cargo train. This reduces dwell time and handling energy and supports predictable takt operations through repeatable handling times and full process automation.

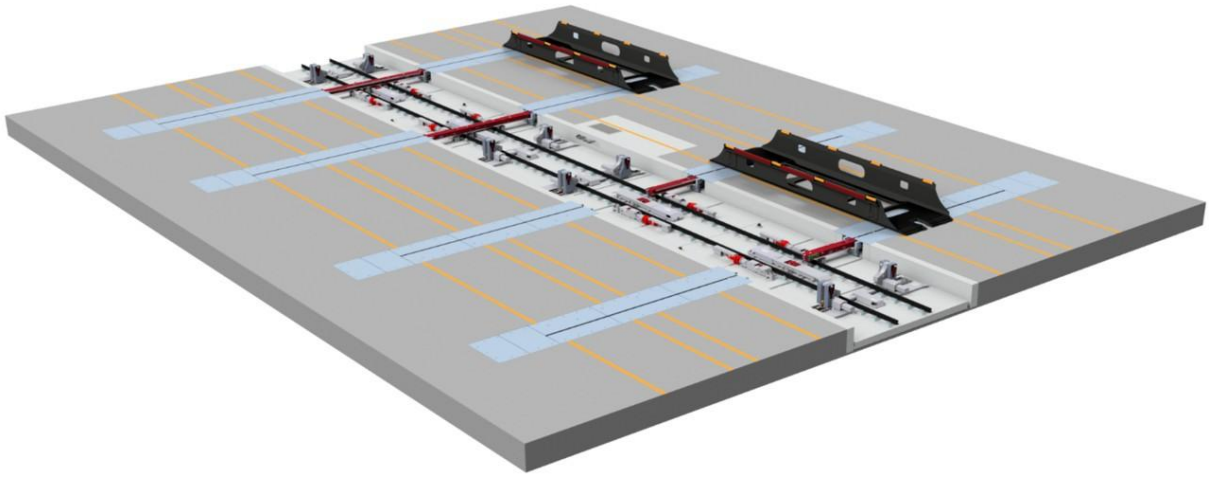


Figure 6: GateModule – Fully automated & electrified road/rail interface for FTL/LTL loads

3.5.4. Digital layer: slot-based booking, event chain, and governance

A central PI enabler is coupling the physical system with a digital internet of bookings, data and visibility. Slot-based booking targets conflict-free reservation of load-space areas for piece goods. Unit-level IDs and scan events create a minimal but robust event chain (gate-in, loaded, departed, arrived, unloaded, gate-out). Governance is implemented via a rulebook covering packaging, acceptance, scanning, exceptions and claims.

4. Real-world LTL on Rail Demonstrator

4.1. Demonstrator context: HOLOGISTICS and corridor selection

Within HOLOGISTICS (Horizon Europe Innovation Action), the demonstrator targets European LTL/unit cargo on the Netherlands and West Germany to Northern Italy corridor and v.v. The corridor is aligned with new-build/upgrade CargoBeamer terminal projects in Kaldenkirchen (DE) and Domodossola (IT) at both ends designed to operate with GateModules and standardized processes, providing the physical foundation for PI-node performance and scalable throughput.

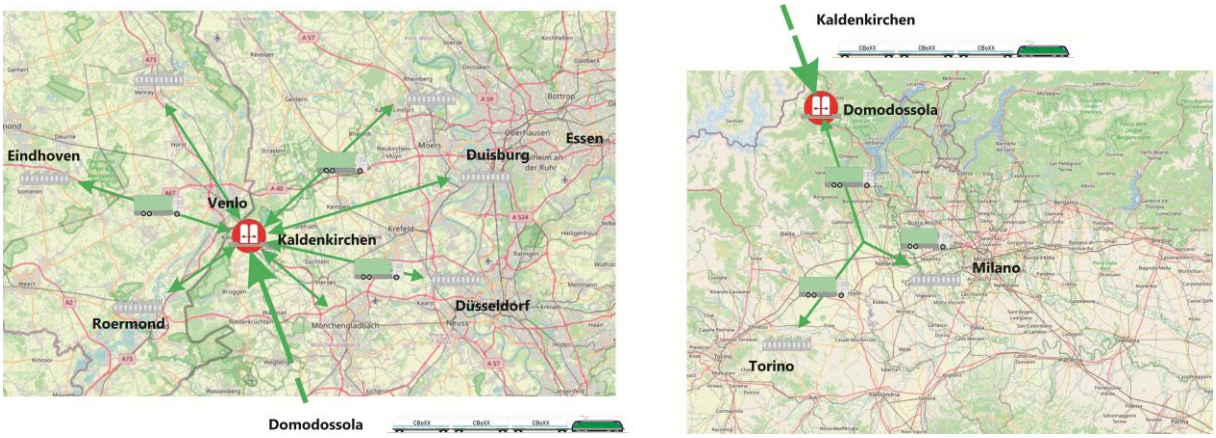


Figure 7: Real-world demonstrator route planned for mid 2027: Kaldenkirchen (DE) near Venlo (NL) to Domodossola / Milano (IT) and v.v., connecting major industrial regions in Europe

4.2. Operating model (takt / circulation plan)

The demonstrator is operated as a repeatable daily circulation with fixed terminal windows and a defined train schedule coordinated with GateModule-driven loading/unloading. Foreseen baseline timing anchors are: Kaldenkirchen loading window 23:00–01:00; departure 02:00; arrival Domodossola 20:00; return departure 02:00. Each cycle includes inbound staging (gate-in, unit ID scan, basic QC), automated transhipment within the fixed timebox, event capture at each handover, and exception playbooks (missed cut-off, non-conforming packaging, damage/claims, documentation issues).

4.3. First/last mile concept and partner operations

Regional operating partners will provide pickup and delivery capacity in defined catchment areas in the north and south. Pickup/delivery runs are planned to meet terminal cut-offs and fixed loading windows. Electrical mid-size trucks will be used for first&last mile transport to minimize emissions. Scan compliance, exception reporting and slot adherence are treated as core PI enablers. The outcome is a scalable pattern for low-emission first/last mile that can be replicated as additional corridors and terminals come online.

4.4. Deliverables and success criteria

Key deliverables include: (i) a working unit-cargo rail LTL process on the corridor integrated with GateModule-based automated terminal handling; (ii) a minimal but robust digital event chain for visibility and KPI reporting; (iii) a rulebook v1 (packaging, acceptance, scanning, exceptions, claims); and (iv) a KPI cockpit for continuous improvement.

4.5. KPIs and measurement

Evaluation focuses on PI-node performance and end-to-end service quality and resulting cost, referring to design criteria #8 and KPI listed in chapter 3.2.5. Where possible, comparisons are made against baseline road-based LTL processes.

4.6. Expected environmental effects and PI adoption pathways

The demonstrator supports decarbonization through electrified energy-efficient terminal handling, low-emission first/last mile, energy-optimized electrified long-distance rail transport. Beyond corridor performance, the main contribution is a replicable PI pattern: Standardized modular units, SOP (Standard Operating Procedure) - driven PI nodes, a digital event chain, and governance that can inform broader PI adoption and policy discussions for rail-based unit cargo networks.

5. Conclusion and Future Work

This paper outlined a PI-aligned system design for competitive and environmentally-friendly LTL-on-rail transport system combining electrified trucks, CBoXX rail container equipment, GateModule automated PI nodes, and a digital booking/visibility layer. We detailed the planned 2027 EU Innovation Action "HOLOGISTICS" demonstrator on the Kaldenkirchen - Domodossola rail corridor and its operating model, governance and success criteria.

Future work will

- report digitalization & software structure foreseen for the designed CBoXX-based LTL-on-rail system following π-criteria,
- report demonstrator results based on measurement vs. calculated KPI,
- refine interoperability and data standards,
- report results of additional research on automation within the π-node (e.g. autonomous forklifts executing cross-docking from e-Trucks into and oout of the CBoXX in the GateModules), and
- extend the approach to additional corridors and PI nodes, and
- report on mixed FTL / LTL rail cargo both using GateModules as road/rail nodes, a procedure which promises further enhancement of competitiveness of rail cargo.

References

- ALICE (2022): Alliance for Logistics Innovation through Collaboration in Europe (ETP-Logistics).
  Roadmap to the Physical Internet — Executive Version (PDF). November 2022 (as hosted on ETP-Logistics).https://www.etp-logistics.eu/wp-content/uploads/2022/11/Roadmap-to-Physical-Intenet-Executive-Version_Final-web.pdf
- Landschuetzer, C (2015): *Containers for the Physical Internet: requirements and engineering design related to FMCG logistics*. Open Springer Press, Containers for the Physical Internet: requirements and engineering design related to FMCG logistics | Logistics Research | Springer Nature Link
- HOLOGISTICS (2025): EU Research Project Grant Agreement 101202556, CORDIS fact sheet.
- MODULUSHCA (2015): EU Research Project Homepage**.** www.MODULUSHCA.com
- MultiRELOAD (2025): EU Research Project Homepage. www.multireload.eu

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE
TNO innovation for life

Investable Interoperability for the Physical Internet

Participant-Centricity and a Viable Market for Common Services

Inge Lucassen[1], Huib van Nieuwenhuijze[1] and Harrie Bastiaansen[1]
TNO, the Hague, the Netherlands
Corresponding author: inge.lucassen@tno.nl

Abstract: *This paper addresses the architecture of the digital infrastructure for the emerging Physical Internet (PI) with its diverse and dynamic data sharing requirements. It reflects on the various European initiatives that are currently driving the development of data sharing solutions for the logistics domain, specifically the European Mobility Data Space (EMDS) and the FEDeRATED architecture. Building and extending upon these EC data sharing initiatives, and recent lessons learned, this paper proposes two complementary approaches to serve the needs of the PI: (1) PI-participant-centricity that puts individual PI participants in the lead for establishing (their prioritized) data needs and pain points, and (2) stimulating a viable market for common services and their providers to support data sharing that can be leveraged across domains. This two-pronged approach supports converging of data space developments and leads to enablement of cross-cutting use cases at large-scale, thus providing the interoperable ecosystem as required for the PI.*

Keywords: *Physical Internet, Data Sharing Needs, Data Spaces, Trust, Services Market, Customer-Driven Design, Roadmap, Adoption*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☐ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☒ *Logistics/Commercial Data Platform,* ☒ *Access and Adoption,* ☒ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

1 Introduction

Growing demand for transport and societal expectations around sustainability, climate neutrality and the responsible use of scarce resources are exceeding the boundaries of existing optimization measures. As Ballot et al (2020) states, a shift to more sustainable transport modes alone is insufficient; genuinely meeting climate objectives requires a substantial reduction in the total deployment of logistics assets. The Physical Internet (PI) aims for greater efficiency by next level optimalization of the logistics system, and sets an ambitious target on its roadmap: a future in which a 300% increase in transport demand can be handled with only 50% growth in assets, enabled by radically higher utilization rates. Alongside this, the PI seeks to reduce emissions, energy consumption and congestion, while strengthening the resilience of logistics networks. In the PI, companies will no longer work in static linear supply chains. Instead, today's fragmented landscape of logistics stakeholders must evolve towards a hyperconnected PI network (a "system-of-systems"). More dynamic routing and greater operational flexibility reduces dependencies on specific actors, modalities or locations, whilst supporting faster recovery from disruptions.

Large-scale (digital) interoperability provides the basis of this logistics PI network. It must be seamless, flexible and resilient, with complete visibility of services, capacities and available resources. Both transport demand and assets are pooled and they can be efficiently matched, supported by secure exchange of operational, commercial and financial information. Ballot et al. (2020) outlines a phased roadmap leading to full PI deployment by 2040, covering changes in organizational practices, digital infrastructure, physical operations and governance, while Michiels et al. (2025) provide a blueprint for its digital infrastructure.

Digital interoperability must provide the stakeholders with a comprehensive digital representation of the entire PI network, referred to as the "*network state*". Combined with real-time "*track and trace*" information, this enables "*network dynamics*", a state in which digital agents autonomously determine optimal routing strategies and continuously adjust shipments to achieve system-wide efficiency The relationships between these concepts are visualized in Figure 1.



Figure 1: Relationship between Network state, Track & Trace and Network Dynamics

Based on the network state, PI network stakeholders may engage into logistics business activities by means of smart contracts. Smart contracting refers to the digital encoding of contractual terms in a way that allows them to be executed automatically, while digital negotiation concerns the joint, digital agreement and coding of those terms. Although these concepts are not deeply explored in existing PI literature, such as Ballot et al (2020) and Michiels et al (2025), they are highly relevant to the ambition of Network Dynamics and the PI roadmap, which envisions a fully autonomous logistics system with minimal human intervention. For such autonomy to function reliably, contractual agreements must be available in machine-readable form so that execution aligns with agreed conditions, making smart contracting a prerequisite. As the system evolves, autonomous agents may not only execute existing agreements but also negotiate new transport contracts in real time. Digital negotiation therefore represents an optional but potentially powerful enabler of Network Dynamics, enabling the system to identify and activate alternatives rapidly, thereby increasing resilience and supporting the PI's aim of seamless, data-driven coordination across the entire network.

In such a digitalized approach for the PI network, large-scale sharing of data between stakeholders will be key. This will be the case both for building the digital representation of the PI network (the network state) and for engaging into logistics business activities (the smart contracting). A complex challenge on data sharing arises as (1) many stakeholders are involved requiring dynamic data sharing and digital relationships to be established, (2) the data to be shared may be (commercially or privacy) sensitive, requiring adequate and legally binding data sharing policies to be made in a real-time and automated manner, and (3) the types of data to be shared may range widely, e.g. including logistics events, operational status and logistics assets availability and smart contract negotiation information.

At the same time, the European Commission (EC) is taking the lead in shaping Europe's data economy by developing and deploying both regulations and technical implementations (architectures, building blocks, open-source solutions) for interoperable data sharing. Building and extending upon these current EC data sharing initiatives for the logistics sector, this paper proposes two complementary approaches for developing a data sharing infrastructure that

serves the needs of the PI: (1) participant-centricity that puts individual PI participants in the lead for establishing (their prioritized) data needs and pain points, and (2) stimulating a viable market for common services and their providers to support data sharing that can be leveraged across domains. These are depicted in Figure 2.



Figure 2: Participant-centricity and a focus on a viable market enable the Physical Internet

To address these approaches, this paper has the following structure: Section 2 presents the diverse and dynamic data sharing requirements and their operationalization challenges for the emerging PI, after which section 3 describes the current EU logistics data sharing initiatives with the lessons learned on their applicability for the PI. Subsequently, section 4 addresses the need for a new two-pronged approach. These approaches are further motivated and elaborated in section 5 and section 6, respectively. Finally, section 7 presents the overarching conclusions.

2 PI data sharing needs and their operationalization challenges

In the PI network, data from many sources need to be combined to achieve the concept of the "*network state*". Routing decisions, for instance, require order data for planned orders as well as those in transit, plus information on available services at nodes and the capacity and speed of movers between them. Increasing standardization of logistics services should simplify bookings and reduce administrative burdens, while deeper digital integration will enable fully automated processing.

Therefore, to achieve an operational PI, reliable data need to be shared, mostly in real-time. This data is often commercially sensitive, because the optimization needs not only order information. It also needs more sensitive service details such as capacity availability, service levels and pricing to optimize for various outcomes and KPIs. As Figure 1 shows, the complexity and sensitivity of the data increases when moving to the right and approaching Network Dynamics. Sharing such sensitive data is counter-intuitive for the logistics sector that is known for its competitiveness and its pragmatic 'can-do' attitude. This attitude leads to less appetite for fundamental changes but is also unfavourable towards describing and thereby limiting the offered services upfront. To approach a network optimum, many PI participants need to onboard. The sector is fragmented and shows very different levels of digital maturity. This poses additional requirements for onboarding, ability to invest and willingness to pay.

To operationalize the PI concept, advanced mechanisms for sharing (commercially) sensitive data between many PI stakeholders and participants are required. The PI requires a decentralized data-sharing architecture in which trust, data sovereignty and granular control over who can access what information are essential. Ballot et al. (2020) already highlights this need by referring to the FEDeRATED initiative, described as "a federated Network of Platforms consistent with the PI vision and its decentralized architecture". The FEDeRATED architecture is based on data spaces technology.

The PILL project explored how data spaces concepts can support PI implementation. According to Michiels et al. (2025), data spaces offer an ideal architectural foundation for the PI because they prioritize data sovereignty, mirror the PI's decentralized nature, support discoverability and transparency through event-based information exchange, and rely on open standards and

modular components that can flexibly integrate with diverse logistics processes. While this assessment is positive, it also identifies important challenges. Publishing and discovering logistics services still requires a substantial degree of standardization, which is equally critical for enabling smart contracting and achieving end-to-end automation. Furthermore, constructing a reliable and actionable overview of the network state demands consistent modelling and standardization of complex logistics processes; without such harmonization, orchestrating and monitoring execution across a federated PI environment becomes significantly more difficult.

3 EU Data Strategy and its EU logistics data sharing initiatives

The European Commission (EC) is driving the development of data sharing solutions to strengthen the data economy and ensure Europe's competitiveness in general, with the EU Strategy for Data, the Data Union Strategy and a focus on creating Common European Data Spaces. The EC Directorate-General for Mobility and Transport (DG MOVE) is leading the European Mobility Data Space (EMDS), encompassing both personal mobility and logistics, to support and interconnect the plethora of existing data sharing initiatives in the mobility and logistics sector. Simultaneously, it promotes the FEDeRATED architecture for event-driven data sharing in logistics. Both initiatives will be harmonized under the EU's new EDIC M&L governance initiative (Kuiper, 2025) and are set to provide the foundation for future digital infrastructures, enabling the PI network, as was also highlighted in Michiels et al. (2025).

Data spaces are organizationally siloed because they are fundamentally positioned as “community-driven” ecosystems with their own governance authorities, membership rules, onboarding processes, accreditation models and legal agreements. This works well within a single domain and region but almost by definition collapses when participants need access across multiple sectors or regions. The first wave of data spaces focused on building trust, legal certainty and operational rules inside small, like-minded communities, rather than across them. Emerging from this first stage, it is clear the technology can deliver in this restricted environment, and the focus moves to cross-domain applications, e.g., combinations of energy + logistics + mobility.

4 Lessons learned: need for a new approach to support the PI

There is already an extensive landscape of existing data sharing initiatives in the logistics sector but it is technically fragmented. An inventory for the mobility and logistics sector, as done by a European Mobility Data Space initiative (EU PrepDSpace4Mobility project, 2023), identifies hundreds of such operational, community-driven, data sharing initiatives and data spaces. In the Netherlands, the Dutch Ministry of Infrastructure and Water Management has invested in piloting of data sharing concepts in several individual communities. A comparative assessment (Geonovum, 2025) has shown that it resulted in: different interpretations of “trusted data sharing” across ecosystems, divergent approaches to trust, identity, interoperability, metadata, and data models, overlapping functionalities, built independently (e.g., identity management, catalogues, trust frameworks) and misalignments in component maturity levels (some initiatives were strong in semantics, others in sovereignty & trust, but no ecosystem is mature on both). Overall, there is a lack in use of common components and consequently insufficient reuse, despite many components being functionally similar across domains. As a result, Geonovum (2025) concludes, there is no common technical layer that enables seamless participation in multiple of these ecosystems, and natural convergence has proven unrealistic without deliberate coordination.

Moreover, these initiatives have been funded with a particular use case and stakeholders in mind. This results in individual data sharing initiatives and data spaces being economically

fragile because most lack sustainable revenue models, rely heavily on public subsidies, and struggle to articulate clear, direct value for participants. An analysis of 155 European data spaces (Moonen et al., 2025) reveals that only 15% have a clear revenue model. Many initiatives remain stuck in pilots without adoption or network effects to support recurring business models.

This fragility has emerged because early data space efforts were driven by public policy goals, such as sovereignty, openness, and compliance, rather than market needs. As a result, governance frameworks, technical infrastructure and community processes were built without considering how value is created, captured, and monetized. Moonen et al. (2025) conclude that initiatives struggle to find a sustainable revenue model. Usage-based pricing does not work well for a federated solution such as data spaces. Because there is no central platform or oversight on data flows, the organization operationalizing the data space has no visibility on data accessed or used and can therefore not measure usage and bill transparently or fairly. This is a key difference with centralized data platforms that often rely on (a combination of) usage fees, subscriptions and/or licensing. Limited understanding of how value is created, limited adoption at scale, and unclear cost–benefit trade-offs for participants and solution providers, limits willingness to invest, for both participants and solution providers (Moonen et al, 2025).

Nevertheless, for the PI, a bottom-up approach was suggested in the Roadmap to the Physical Internet (Ballot et al., 2020). The expectation was that smaller networks develop their own governance and over time these smaller communities mature organically to a central, interoperable network. That might work for some areas of development for the PI. But from the lessons learned and the observations on the need and challenges on interoperability as addressed in this paper, it can be concluded that the plethora of data sharing initiatives and data spaces is far from converging, and that without structured action this will also not happen organically. A bottom-up approach will not drive adoption. Focus on developments for specific use cases results in solutions that are siloed and not interoperable. This quickly results in a poor PI participant experience as businesses will need to invest in and onboard to multiple communities with a multitude in variations of technology, legal conditions and rules. A bottom-up approach will also not drive necessary investments from solution providers. With siloed development, the services that need to be offered cannot scale sufficiently to make a positive investment case while the revenue model remains unclear. Instead, deliberate coordination, shared development of common building blocks, and structured governance can drive convergence.

With on the one hand the challenging requirements for the operationalization of the PI concept, and on the other hand the EC driven initiatives for the data economy, the question automatically arises whether continuing the current approach will naturally lead to a harmonized and perfect match. This paper focuses on the development of a two-pronged approach for the next phase of establishing such a harmonized digital foundation for the PI. These two approaches are further motivated and elaborated in the following sections, respectively.

5 Putting participants front and center: participant-centricity

Participant-centricity adds the perspective of PI participants with varying digital maturity levels to research and development initiatives. This suggested approach builds onto the common concept of customer centricity; a growing field of study, often discussed together with product centricity, customer experience and customer loyalty (Da Silva Muniz Sobrinho et al., 2025). Research on customer centricity is prominent in marketing, information systems and especially in financial services, where digitization and service innovation are key.

Customer centricity drives improvements in customer experience, and is foundational for successful innovation (Athaide et al., 2024). There is strong empirical evidence of the relation between Customer Experience Management (CXM) and innovation success. CXM is the

strategic and systematic management of how customers interact with a company with the aim of delivering a seamless, valuable, and consistent experience. Companies that excel in innovation typically have a customer-centric culture, intentionally design touchpoints, and allocate resources in ways that maximize impact on experience improvements. This mindset ensures that the entire customer journey is understood, and that innovation is not technology-push but anchored in real customer needs. In dynamic and complex environments, an intentional and structured focus on journeys is particularly beneficial. Athaide et al. advise that the most straightforward first step is customer journey mapping: a visual representation of the steps a person takes when pursuing a goal. The next step is identifying the most critical customer segments and pain points.

CXM typically includes customer journeys of pre-purchase, purchase, and post-purchase and the concept of these journeys can also be applied to the adoption of data sharing services that enable businesses to participate in the PI. Additionally, the journeys should include pain points and learnings of earlier work on data spaces technology to help focus future efforts and maximize impact. Figure 2 visualizes a set of customer journeys and highlights the detailed steps of the "Onboard to Ecosystem" journey as example. An important difference between the typical use of CXM within businesses and our use of it here, is that in our case, there is typically not one organization that offers all these touchpoints to PI participants, but instead multiple service providers jointly deliver these journeys. This makes a structured approach of journeys even more valuable as the experience is complex. It is difficult for PI participants to navigate and achieve their goals.

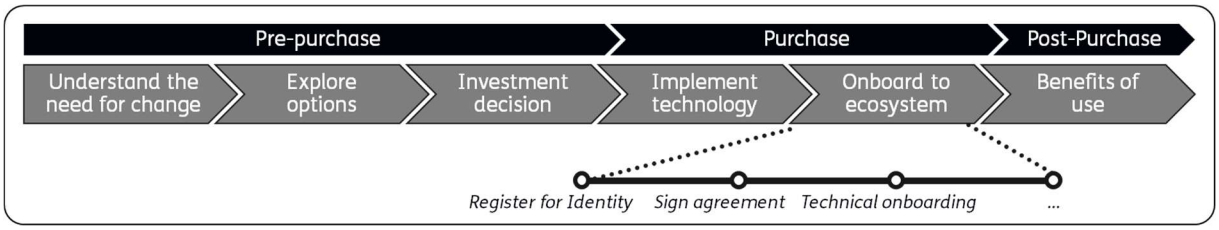


Figure 2: Illustration of Customer Journeys and touchpoints for data sharing in the Physical Internet

The "Onboard to Ecosystem" journey connects to challenges that have been identified in the previous section. It was concluded there that data spaces are organizationally siloed and consequently technically fragmented. This is reflected in the onboarding experience for participants. They will need to onboard to an ecosystem before realizing benefits of use. They might require multiple identities for themselves and go through different onboarding processes before being accepted in relevant ecosystems, including multiple and varying community contracts to be signed. This is a poor experience and limits cross-community or cross-sector data sharing. Ensuring a customer or PI participant-centric mindset and having a structured customer journey view, ensures that the early needs of the targeted community of adopters are identified. Only with early understanding and focus, will it lead to adequate consideration of the complexities of their running business and the need for cross-community interoperability instead of focussing on greenfield implementations.

Customer journeys are typically not the same for all customers. A customer segment is a clearly defined group of customers who share relevant characteristics that lead them to have similar needs, expectations, behaviours, and experiences across a customer journey, and who therefore require a coherent and tailored experience strategy. Meaningful dimensions (perspectives) for segmentation of customers to adopt PI data sharing solutions could be:

- *Needs*: Role in the Physical Internet, e.g. Mover, Node, Owner of the Goods, …
- *Behaviour*, e.g. self-service oriented or support/consultancy dependency, …

- *Attitude and Motivations*, e.g. price-sensitive, risk-averse, innovative mindset, …
- *Context of Interaction*, e.g. digital maturity (routines in implementation of new technology), time-criticality or urgency of adoption, ….
- *Lifecycle stage*, e.g. first time onboarding, experience with PI pilots, ….

Overall, the approach using customer journeys, combined with relevant customer segmentation ensures that needs of participants are reflected throughout all stages of research, innovation adoption and participant lifecycle.

The other key lesson from the previous section is that data spaces are financially fragile and part of the solution is that participants need to pay for data sharing services. Understanding what goes into this decision for businesses with different characteristics can improve by researching the "investment decision" journey for various segments. An important factor in the competitive logistics industry is clearly whether the benefits of the solution outweigh the costs. Buyers need to see the value of multiple use cases and the possibility to leverage common components. Businesses might consider a growth path in adopting new technology and think about how they prepare their data, now siloed within vendor-owned technology stacks, and turn it into high quality data products that can be shared. They might think about where they could buy software and value that leverages the newly available data. Moreover, they might think about how this impacts business continuity, competitiveness, sovereignty and performance. Researching these business needs and addressing them will improve the business case for PI participants, reduce risks and support adoption.

6 Developing a viable market for service providers

Many of the foundational components for data spaces are common, mature, and not specific to the PI or even the logistics or mobility sector. A viable market for common data sharing services is interoperable and provides technology solution and service providers with the ability to scale and grow their market. These service providers need the opportunity to focus on their own value-add and Return-on-Investment (RoI). Especially infrastructure service providers need to be able to create economies-of-scale. Providers that are offering solutions for highly use case or domain specific value-adds should be prevented from having to spend (too) much effort and budget on developing common infrastructural services.

Recent research (CoE DSC, 2026; Berkers et al., 2026) shows that it is possible to achieve this standardization of services by the market. A Data Spaces Services Model (DSSM) was developed that structures all necessary services for enabling data spaces. It describes an adequate set of data space building blocks as re-usable, intermediary services. This study focuses explicitly on services as "value-creating offerings delivered to customers", rather than technical building blocks or implementation activities of data spaces. Services are thereby directly connected to the value they provide, independent of how these services are technically realized. It uses the classification by the DSSC of four service categories: Participant Agent Services, Federation Services, Support Services, and Value Creation Services. During market dialogues, it was reviewed what the development strategy for these services currently is and should be in future, distinguishing: Unique (specific and highly dynamic), Adaptive (low differentiation, responsive), Common (low differentiation, standard), and Specialized (stable but specific). The development (from current to future) was assessed for all services in the DSSM, as shown in Figure 3. The figure shows that none of the services is expected to be Unique. Instead, most services are expected to be or become Common. Many services are currently Adaptive to specific applications or use cases, but this number is expected to decrease by 50% according to service providers.

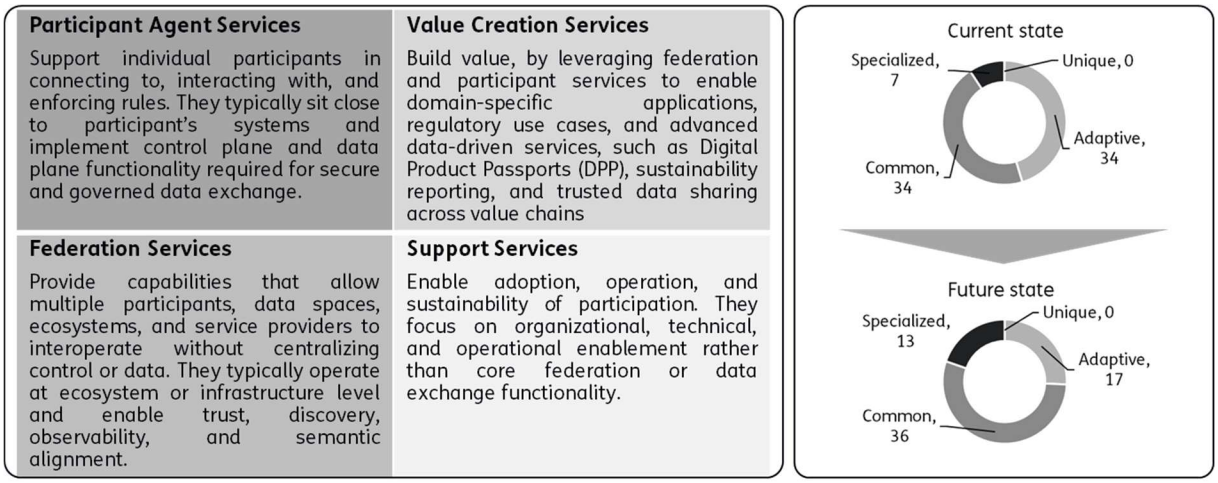


*Figure **3**: DSSM service categories (l) and evolution of Development strategies (r)(CoE DSC, 2026; Berkers et al., 2026)*

Due to its complexity and fundamental enablement of data spaces, trust frameworks require special attention for a viable market. For sharing (potentially sensitive) data in the PI network across data spaces and community borders, the development, deployment and adoption of aligned community-driven and participant-driven cross-cutting trust frameworks will need to be pursued (Bastiaansen et al, 2025). The participant-driven trust framework will be dominant in the PI network due to its complexity in number of stakeholders, variety in processes, exceptions and sensitivity of the data. The participant-driven trust framework matches very well with the participant-centric model (as argued for in the previous section) by putting control over data sharing conditions firmly under the control of individual participants. This can be enabled by the Common Carrier Layer (CCL), whose onboarding provides participants with a single point-of-entry to access data across data sharing initiatives and data space borders (Bastiaansen et al., 2025; Pretsch et al., 2026). A community-driven trust framework may serve the specific needs of smaller, well-defined but confined use cases. With one common onboarding, participants should be able to participate in larger and smaller ecosystems. To help the market develop common services that can be leveraged in both frameworks, strong EU-wide governance is needed on this topic.

Clear guidance is needed for common services to be compliant to the EU Data Strategy, specifically the protocols and standards as being developed by the IDSA and the DSSC to support the EU (common) data spaces. Open-source building blocks thereof are rapidly maturing. Two main open-source software stacks are currently being developed that are becoming the per-default preferred solutions:

- the EDC software stack, and
- the SIMPL software stack.

As interoperability of technical components is key for infrastructure harmonization, designated open standards and open-source components are essential. In this emerging domain and with the complexity of the sector and the needs of the PI, technical support for compliance and certification is essential to provide service providers with the confidence to offer new services. Governance needs to ensure that services and solutions remain compatible.

Next to established trust frameworks and (technical) compliance assurances, viable business models are key for service providers to step into this emerging market, with confidence to

invest. Development of these business models requires additional research and development in close collaboration with the service provider market itself.

7 Conclusions & Recommendations

To enable the PI and to allow cross-domain applications with a good experience for PI participants and a viable market for service and solution providers, this paper has argued that a different approach is needed. Requirements for the operationalization of the PI concept are challenging. And the current approach for developing (community-driven) data spaces for specific use cases is not yet converging into an interoperable ecosystem. Therefore, there are two paths that need to be executed on; a two-pronged approach. A summary of recommendations for these two approaches is presented in Table 1.

Table 1: Summary of recommendations for the 2-pronged approach

Participant-Centricity	***Viable Market***
• Develop a set of participant journeys to ensure participant-centricity. • Research participant needs, pain points and segments according to the Customer Experience Management approach. • Put special focus on the "investment decision" of participants to inform development of new business models. • Develop a harmonized trust-framework that fits participant needs.	• Put in place strong EU-wide governance on architecture and common services, with a special focus on Trust Frameworks. Such as was done for eIDAS for Identities. • Designate preferred technical, open-source, open standard components. • Develop large labs with leading providers to validate the common service approach. • Put special focus on research of business and revenue models.

Common principles of Customer Experience Management need to be leveraged to the benefit of innovation adoption. This brings more focus and attention to how individual businesses, that want to participate in the PI, prepare for their onboarding, purchase necessary solutions, and execute operations in the new ecosystem. A participant-centric mindset brings a new perspective, leads to different questions being asked on pain points and onboarding paths and thereby gives insights on what needs to be solved to deliver a compelling (business) case for adoption.

For a viable market for data sharing services, strong governance is urgently needed. Given the EU-wide developments described in section 3, plus the cross-border nature of the industry, an EU-level governance body is best positioned to drive this. It needs to be a prerequisite that any trust and identity service is by default fitting the participant-driven trust framework and can also be applied to more closed communities. This ensures that these services will not only work for specific use cases but allow PI participants to leverage these services across multiple use cases and domains and bring added benefits. The governance structure needs to organize collaboration, align priorities, and ensure that architectural decisions are consistent across initiatives. Support for compliance with wider EU regulations is a prerequisite to reduce risks. For the next phase of data spaces development for the PI, the common services and solutions need to be developed with a group of preferred providers and be pressure tested in large scale labs that go beyond current piloting. Along with these technical developments, the business, revenue and cost model of these providers need special consideration.

Acknowledgements

The work presented in this paper was done under the umbrella of the TNO project "Future Proof Smart Logistics", with valuable contributions from the Dutch Centre of Excellence for Data Sharing and Cloud (CoE DSC) and the EU-funded project SM4RTENANCE with the aim to

facilitate a neutral cross-sectorial data space for the manufacturing industry. We would like to thank these initiatives for providing the opportunity to do this highly relevant and interesting research within a stimulating and cooperative setting.

References

Athaide G., Zhang J., Klink R. (2024) : Customer Experience Management (CXM) and new product success : an empirical investigation ; European Journal of Innovation Management.

Ballot E., Barbarino S., van Bree B., Liesa F., Franklin R., 't Hooft D., Nettsträter A., Paganelli P., Tavasszy L. (2020) : Roadmap to the Physical Internet, Executive version ; by ALICE ETP.

Bastiaansen H., CoE-DSC. "Towards Scalable Data Space Interoperability and Federation - A Combined Community-Driven and Participation-Driven Approach based on a Common Carrier Layer - CoE-DSC Discussion Paper". https://coe-dsc.nl/wp-content/uploads/2025/11/CoE-DSC-Towards-Scalable-Data-Space-Interoperability-and-Federation.pdf.

Bastiaansen H., Farias Lóscio B., Bouman E., Haijenga F. (2025) : Cross-cutting Trust Frameworks to support Data Sharing across Communities ; by Centre of Excellence Data Sharing & Cloud.

Berkers, F., Punter, M., Stolwijk, C., van Houwelingen, G., Stoter, A. (2026): The Influence of the Emerging Service Provider Market on the Data Space Business Model. Accepted as DSSC Book Chapter, to be published Q2 2026.

CoE DSC / Center of Data Sharing & Cloud (2026) : Market Dialogue in the context of SCSN.

Common European data spaces: https://digital-strategy.ec.europa.eu/en/policies/data-spaces, last accessed 20/02/2026

Da Silva Muniz Sobrinho, J., de Oliviera Dias, M., Vieira Marcolino, G., Bastos da Cruz P. (2025) : Systematic Literature Review on Product and Customer Centricity, Customer Experience , and Loyalty : Open Journal of Business and Management.

Data union strategy - Unlocking data for AI: https://digital-strategy.ec.europa.eu/en/library/data-union-strategy-unlocking-data-ai, published 19/11/2025

European Commission : Communication "Creation of a Common European Mobility Data Space", https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52023DC0751, published 29/11/2023.

EU PrepDSpace4Mobility project (2023). Report on the Inventory of Existing Data Ecosystems https://mobilitydataspace-csa.eu/wp-content/uploads/2024/03/ 2024-03-19-deliverable-d2.1-inventory-of-data-ecosystems-v3.pdf, last accessed 2026/03/19.

European Strategy for Data: https://digital-strategy.ec.europa.eu/en/policies/strategy-data, last accessed 20/02/2026

FEDeRATED Network of Platforms project website, https://www.federatedplatforms.eu/index.php/library/category/architecture, last accessed 24/03/2026.

Geonovum (2025) ; Inventarisatie Architectuur Data Ecosystemen. Commissioned by the Dutch Ministry of Infrastructure and Water Management.

Kuiper J. (2025) : European digital infrastructure Consortium for Mobility and Logistics, ALICE webinar 26/03/2025.

Michiels P., van Bever D. (2025) : A blueprint for the Physical Internet. By the PILL project.

Moonen N., Mollee N., Wentzel V., van den Born A., Vossen A. (2025): Sustainable Revenue Models for Data Sharing Initiatives. Commissioned by the Dutch Ministry of Economic Affairs. By Jheronimus Academy of Data Science.

Pretzsch S., Bastiaansen H., Farias Lóscio B., Chawla K., Mulder P., Udayan S. (2026): "Towards a Common Carrier Layer - A First Step for a Participant-Driven Data Sharing Approach for Data Spaces". Accepted as DSSC Book Chapter, to be published Q2 2026.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE
KEDGE BUSINESS SCHOOL

Improving Transport Utilization via Proactive Consolidation and Coordinated Purchasing

Najat Bara[1*], Mohamed Nour Kitri[2], Onur Ozturk [3]
, Frédéric Gautier[4], and Junkai He[1]

1. The Centre of Excellence for Supply Chain Innovation & Transportation (CESIT), Kedge Business School, Bordeaux, France
2. EMINES- School of Industrial Management, Mohammed VI Polytechnic University, Ben Guerir, Morocco
3. Telfer School of Management, University of Ottawa, 55 Laurier Avenue East, Ottawa, Ontario K1N 6N5, Canada
4. IAE Paris I Panthéon-Sorbonne (Sorbonne Business School), Paris, France.

*Corresponding author: najat.bara@kedgebs.com

Abstract: *This study addresses inefficiencies in decentralized procurement systems where independent project-level decisions fragment shipment volumes, limiting consolidation opportunities and increasing logistics costs. It proposes a proactive, buyer-driven consolidation framework that integrates procurement and transportation decisions to improve overall supply chain performance. Unlike traditional reactive approaches, the framework anticipates consolidation opportunities at the procurement stage, enabling coordinated purchasing across projects. Decision-makers are encouraged to balance trade-offs between purchasing, transportation, and inventory costs, even adjusting suppliers or delivery schedules to enhance truckload utilization. The model incorporates both new and ongoing orders, revealing additional consolidation potential. An integrated optimization model, supported by an Iterated Local Search algorithm, ensures feasible and scalable solutions. Numerical experiments using industrial data show cost reductions and improved resource utilization. The results highlight the managerial and operational value of proactive consolidation in achieving more efficient, cost-effective, and sustainable logistics systems.*

Keywords: Proactive freight consolidation; Heterogeneous product procurement; Global optimization; Transportation

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☒ *Transportation Equipment,* ☐ *PI Networks,* ☒ *System of Logistics Networks,* ☒ *Vertical Supply Consolidation,* ☒ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

1 Introduction

The transportation cost is often a piecewise linear function as the unit cost decreases with volume. One solution to improve transport resource utilization is shipment consolidation. It is the process of grouping different shipments into a larger one to gain economies of scales. This

work is inspired from a real company that deals with the electric works in large construction projects (malls, train stations…). It procures most of the material from suppliers abroad. Depending on the agreed Incoterm (INternational COmmercial TERM), either the selected supplier deals with transportation or one of the company carriers does. The company does local optimization to select the cheapest offer (the supplier and the Incoterm) for each order. This policy leads to the dispersion of orders in space over several regions and over spaced delivery dates in time which decreases the transportation efficiency.

Typically, shipment consolidation occurs reactively; at the time of shipment, the transportation carrier bundles orders whose delivery dates and destinations coincide. Moreover, shipment consolidation is straightforward for conventional, standardized products when order volumes closely align with truck capacity. In such instances, the decision-making process can be conceptualized as a lot-sizing problem that determines the optimal order quantities from each supplier. This paper presents a proactive strategy for shipment consolidation on time and space. We have developed a mixed integer linear model which selects the suppliers based on the trade-off between the purchasing costs, transportation costs and the storage costs. The model distinguishes between already launched orders for which only the truck is determined, and new orders to be launched for which the region (supplier), the picking time (week) and the truck must be determined.

To operationalize this framework, the problem is formulated as an integrated decision model capturing the interdependencies between procurement and transportation planning. The model considers heterogeneous orders, supplier-region constraints, delivery windows, and truck allocation decisions, while explicitly including existing orders with predetermined suppliers and shipment dates. Given the computational complexity of large-scale instances, a tailored Iterated Local Search (ILS) is employed to generate high-quality solutions within practical planning timeframes. The algorithm is designed to respect delivery windows, balance loads across trucks, and iteratively explore alternative supplier assignments for new orders. While the methodological component ensures feasibility and scalability, it is presented as a supporting element; the primary contribution lies in demonstrating the operational and managerial value of proactive consolidation and coordinated purchasing.

The rest of this paper is organized as follows. The second section reviews the literature, reports on the main consolidation configurations. The third section presents the flowchart of the mathematical model. Section four presents the first results of the study. Section five contains a conclusion.

2 Literature review

Freight consolidation, given its significant economic and environmental implications, has been widely explored in the literature. It is generally defined as a deliberate effort to improve the utilization of transportation and storage capacities (Crainic et al., 2009). Such initiatives may be undertaken by different actors within the supply chain, including suppliers, customers, or third-party logistics providers (3PLs).

Several studies adopt the customer's perspective. For instance, works by Ertogral (2008), Aguezzoul and Rezg (2011), and Glock and Kim (2014) focus on standard products and rely on the Economic Order Quantity as a key decision variable. Their models primarily aim to minimize costs related to transportation, ordering, and inventory holding. In addition, Aguezzoul and Rezg (2011) incorporate purchasing costs by accounting for economies of scale. These contributions are generally formulated as lot-sizing problems, which differ from the approach developed in this paper.

Other contributions investigate consolidation from the supplier's standpoint. Zegordi and Nia (2009), for example, analyze a setting where geographically proximate suppliers cooperate by sharing transportation resources, such as trucks, in order to reduce logistics costs. Similarly, Gao et al. (2022) consider a vendor managing multiple shipments of varying sizes alongside a set of containers. Their model assigns shipments to containers while minimizing both transportation and container usage costs, and is formulated as a bin-packing problem. Brown and Ronen (1997) propose an optimization-based decision support system designed for a U.S. manufacturer, where order consolidation is integrated within a vehicle routing framework to minimize total travel distance and support sales operations.

A large portion of the literature focuses on consolidation strategies implemented by 3PL providers. Hanbazazah et al. (2020) develop a mathematical model in which orders, already linked to suppliers, are consolidated from multiple geographically dispersed origins. Shipments are first collected and routed to consolidation terminals before being delivered to a single customer. The objective function captures both inbound and outbound transportation costs, modeled as piecewise functions of shipment volume. Chabot et al. (2018) examine a collaborative setting in which multiple shippers pool their flows, with transportation handled by a 3PL. Their problem is formulated as a vehicle routing problem that integrates both economic and external costs, with consolidation driven by piecewise linear transportation cost structures.

In the context of e-commerce, Zhang et al. (2019) propose a time-based consolidation approach, where orders destined for the same customer are grouped at a delivery station. Assuming prior assignment of orders to warehouses, the 3PL determines consolidation decisions by balancing transportation and inventory holding costs. Nasiri et al. (2018) study a three-echelon supply chain involving suppliers, a distributor acting as a 3PL, and retailers. Their model integrates supplier selection for heterogeneous products, allowing for price variability across suppliers. Upon receiving retailer orders specifying quantities and delivery time windows, the 3PL optimizes both procurement and distribution decisions, minimizing total costs including purchasing, transportation, cross-docking, inventory holding, and penalty costs for early or late deliveries.

This latter work is closely related to our contribution. However, our approach differs in two key aspects: first, consolidation decisions are analyzed from the buyer's perspective; second, we explicitly incorporate piecewise linear transportation cost structures into the modeling framework.

3 Mathematical model for proactive shipment consolidation

In this section we first present an overview of the problem. Second, we present the mathematical model flowchart. Third, we describe the resolution method in case of large instances of the problem.

Overview of the problem

We study a procurement environment in which several purchasing agents, all belonging to the same firm, generate orders for a wide range of products supplied by vendors distributed across multiple geographic areas (see Figure 1). These sourcing decisions play a crucial role in ensuring that materials are available when needed for subsequent production processes. Instead of allowing each buyer to make isolated decisions, the problem is addressed from a system-wide perspective, where all orders are considered jointly in order to enhance coordination and improve overall performance.

In such a setting, both the unit purchase cost and the replenishment lead time vary according to the supplier selected. At the same time, transportation activities are subcontracted to third-party logistics providers, whose pricing schemes depend on the consolidated shipment volume originating from the same region within a given time window. This creates strong interactions between sourcing and transportation choices, making it essential to consider them simultaneously.

To address this, the proposed approach integrates supplier selection, delivery scheduling, and shipment consolidation into a unified decision framework. The objective is to strike a balance between procurement expenses, transportation costs, and inventory-related costs. By aligning decisions across buyers and coordinating shipment schedules, the framework promotes the grouping of orders both geographically and temporally. This, in turn, enhances vehicle fill rates, reduces reliance on fragmented shipments, and leverages scale economies in transport, while still meeting production requirements on time.

An important aspect of the model is its explicit differentiation between two categories of orders. The first category corresponds to ongoing orders that have already been placed with suppliers and are currently in production. These orders are associated with fixed sourcing decisions and predetermined shipping periods; therefore, the model focuses solely on assigning them to appropriate transportation resources. The second category consists of new orders, for which greater flexibility exists. For these orders, the model simultaneously determines the supplier, the shipping period, and the transport allocation. This distinction enables the coordination of new decisions with existing commitments, uncovering additional consolidation opportunities and ultimately leading to more efficient and cost-effective supply chain operations.

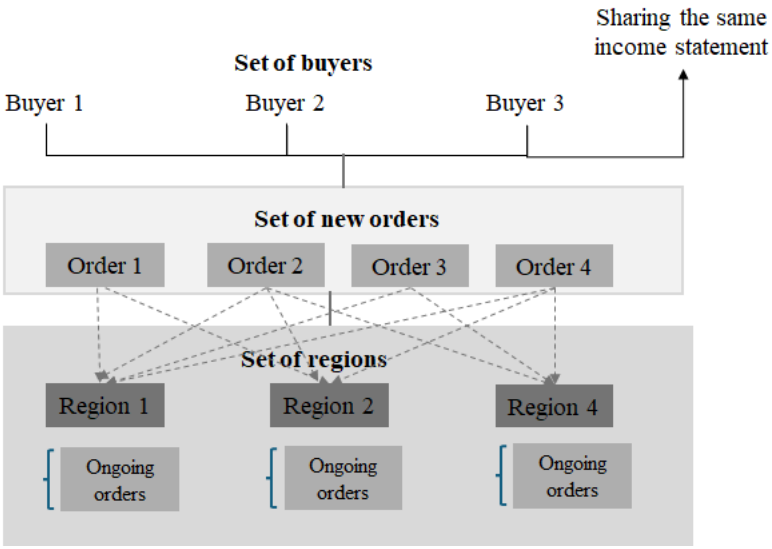


Figure 1: decisional context

The mathematical model flowchart

Figure 3 presents a flowchart of the model. It describes the main parameters associated with new and ongoing orders, structed data and the considered decision criteria. The right side of the figure lists the main decision variables.

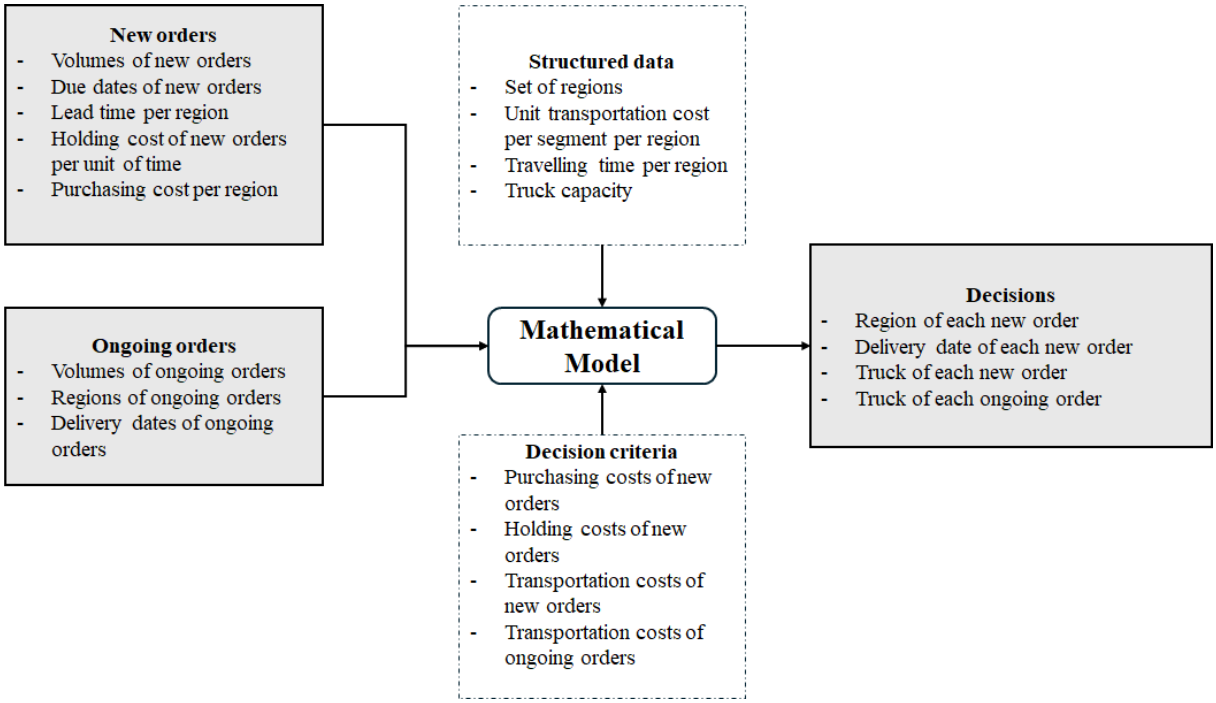


Figure 2: mathematical model flowchart

Iterated Local Search for large instances of the problem

This work tackles the limits of exact optimization when applied to large-scale freight consolidation problems. Although a mixed-integer linear programming model can identify optimal solutions for small instances, its computational burden grows rapidly with problem size, often preventing timely or even feasible solutions. To address this issue, a customized metaheuristic grounded in the Iterated Local Search (ILS) paradigm is developed.

The solution process starts by constructing a feasible plan through a heuristic that distributes loads while satisfying operational constraints. This initial solution is then progressively improved using two complementary neighborhood search procedures: one focuses on optimizing vehicle loading decisions within the same region, while the other explores alternative regional assignments for orders.

To enhance exploration and avoid stagnation, the method introduces controlled modifications to the current solution, with increasing intensity when improvements become scarce. The search continues until predefined stopping conditions are met, based either on runtime limits or a lack of progress. Overall, this approach enables the efficient generation of high-quality solutions for complex and large problem instances where exact methods become impractical.

4 First results of the study Performance Assessment

The effectiveness of both the MILP formulation and the proposed ILS metaheuristic was assessed through a series of test instances derived from data reflecting the historical operations of the studied company. For each category of instances, ten scenarios were generated. It is important to note that the medium- and large-scale instances are the most representative of real-life situations encountered by the company.

For small-scale instances, both the MILP model and the ILS metaheuristic were able to consistently identify optimal solutions. While the MILP achieved optimality in most cases with negligible optimality gaps, it required significantly higher computational time compared to the

ILS, which reached identical objective values within a fraction of a second. This highlights the suitability of both methods for small problems, with a clear advantage for ILS in time-sensitive contexts.

As problem size increases, performance differences become more pronounced. For medium-sized instances, the MILP approach maintained good solution quality but struggled to prove optimality within the time limit. In contrast, the ILS method systematically generated feasible solutions across all instances and frequently matched or outperformed the MILP results, while maintaining substantially lower computation times. This demonstrates the robustness and efficiency of the ILS when addressing moderately complex problems.

The gap between the two approaches widens further for large-scale instances. The MILP formulation encountered significant computational limitations, failing to establish optimality and producing feasible solutions in only a subset of cases. Conversely, the ILS consistently delivered feasible and high-quality solutions for all instances, with notably reduced computational effort. These results underline the scalability of the ILS and its relevance for handling complex, real-world problems.

A closer analysis reveals that instances with a higher proportion of new orders are inherently more difficult to solve, as they introduce greater uncertainty and require additional resource allocation. In this context, the ILS demonstrated strong adaptability, maintaining solution quality even under increased variability. Overall, the aggregated results confirm that while both methods are effective for small instances, the ILS significantly outperforms the MILP in terms of scalability, robustness, and computational efficiency for medium- and large-sized problems, making it a practical alternative for real-world applications.

Table 1: Aggregated results

Size	MiLP Formulation				ILS procedure (algorithm 3)			
	Optimal	Feasible	Avg objective	Avg cpu	Optimal	Feasible	Avg objective	Avg cpu
Small	8	10	465560	1455	8	10	465560	1
Medium	1	9	1737297	5797	1	10	1666634	482
Large	0	5	2947458[1]	7200	0	10	2836419[1]	3676

[1] Average computed for the instances for which the MILP found a feasible solution.

Effect of consolidation effort

The framework is evaluated using instances derived from historical industrial data. Numerical experiments compare decentralized, project-level decision-making with proactive, coordinated procurement. Results show that proactive consolidation consistently improves truck utilization and reduces transportation costs, yielding total cost reductions ranging from 1.76% to 4.05%. Including ongoing orders further expands consolidation opportunities, highlighting the importance of continuous coordination across projects. Beyond economic gains, the proactive and coordinated approach enables faster consolidation by carriers, more predictable delivery patterns, and improved resource utilization, offering both operational and environmental advantages.

In summary, this research highlights the strategic importance of proactive consolidation and coordinated purchasing in project-based supply chains. By embedding consolidation opportunities into procurement decisions, organizations can achieve higher truckload utilization, reduce logistics costs, and improve operational efficiency. The study provides

actionable guidance for managers seeking to align procurement practices with transportation efficiency and resource optimization, showing that anticipating consolidation opportunities and fostering cross-project coordination produces measurable improvements in both cost and operational performance. This proactive, coordinated approach demonstrates a practical pathway toward more efficient and responsible logistics management in complex supply chains.

5 Conclusion

This study presents a proactive shipment consolidation approach that simultaneously optimizes supplier selection, delivery scheduling, and truck assignment for heterogeneous orders. By integrating both new and ongoing orders, the model enables coordinated procurement across regions, revealing consolidation opportunities often missed by decentralized strategies. This approach not only ensures timely material availability but also reduces overall operational costs through upstream coordination.

Computational experiments on industry-inspired instances show that the proposed Iterated Local Search (ILS) method outperforms the MILP formulation, particularly for medium- and large-scale problems, by providing high-quality solutions with significantly lower computation times. The consolidation results highlight the economic benefits of coordinated procurement, primarily through reduced transportation costs, yielding total cost savings between 1.76% and 4.05%.

The current model focuses on a single buyer perspective. Future work could extend it to multi-actor, cross-company collaboration where suppliers, 3PLs, and buyers share PI-enabled logistics infrastructure. This could examine how decentralized actors can jointly consolidate shipments, balancing cost reduction and service reliability. Furthermore, building on the proactive model, research could investigate dynamic, real-time consolidation where new orders are continuously integrated into ongoing flows across a PI network. Using IoT-enabled tracking and smart containers, the system could adapt delivery timing, vehicle allocation, and hub usage on-the-fly to optimize cost and service levels.

References

- Aguezzoul, A., Rezg, N., 2011. The Impact of Transportation Network on the Supplier Selection. IFAC Proceedings Volumes 44, 2162–2166.
- Brown, G.G., Ronen, D., 1997. Consolidation of customer orders into truckloads at a large manufacturer. Journal of the Operational Research Society 48, 779–785. https://doi.org/10.1057/palgrave.jors.2600430
- Chabot, T., Bouchard, F., Legault-Michaud, A., Renaud, J., Coelho, L.C., 2018. Service level, cost and environmental optimization of collaborative transportation. Transportation Research Part E: Logistics and Transportation Review 110, 1–14.
- Crainic, T.G., Marcotte, S., Rei, W., Takouda, M.P.L., 2009. Proactive order consolidation in the retail supply chain. Logistics and Transportation (CIRRELT) and Department of Management and Technology 3, 8.
- Ertogral, K., 2008. Multi-item single source ordering problem with transportation cost: A Lagrangian decomposition approach. European Journal of Operational Research 191, 156–165.
- Gao, Z., Birge, J.R., Chen, R.L.-Y., Cheung, M., 2022. Greedy algorithms for the freight consolidation problem, in: 22nd Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2022). Schloss Dagstuhl-Leibniz-Zentrum für Informatik.

- Hanbazazah, A.S., Castro, L.E., Erkoc, M., Shaikh, N.I., 2020. In-transit freight consolidation of indivisible shipments. Journal of the Operational Research Society 71, 37–54.
- Nasiri, M.M., Rahbari, A., Werner, F., Karimi, R., 2018. Incorporating supplier selection and order allocation into the vehicle routing and multi-cross-dock scheduling problem. International Journal of Production Research 56, 6527–6552.
- Zegordi, S.H., Beheshti Nia, M.A., 2009. Integrating production and transportation scheduling in a two-stage supply chain considering order assignment. The International Journal of Advanced Manufacturing Technology 44, 928–939.
- Zhang, Y., Sun, L., Hu, X., Zhao, C., 2019. Order consolidation for the last-mile split delivery in online retailing. Transportation Research Part E: Logistics and Transportation Review 122, 309–327.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE
KEDGE BUSINESS SCHOOL

Adaptive Capacity Orchestration in Physical Internet Parcel Networks via Resource Pooling

Yaxin Pang[1], Walid Klibi[1], Matthieu Lauras[1], Eva Petitdemange[2], Johan Leveque[3]

[1]KEDGE Business School, CESIT, France
[2]IMT Mines Albi, CGI, France
[3]La Poste, France

Corresponding author: matthieu.lauras@kedgebs.com

Abstract: *Prior studies on pooling and shared as well as research on adaptive rebalancing in mobility systems demonstrate the operational value of coordinated resource flows, particularly in distribution networks. However, they do not provide an integrated multi-period tactical framework combining uncertainty, heterogeneity, environmental constraints, and inter-node mobility. To address this gap, we propose a two-stage stochastic multi-period optimization model designed to support resource pooling in Physical Internet-enabled parcel hub networks. The framework explicitly represents residual capacity as a mobile asset that can be redistributed through dynamic pooling across interconnected logistics nodes under uncertainty. We carried out numerical experiments inspired by a real parcel operator, to illustrate the proposed approach.*

Keywords: *Physical Internet; adaptive capacity pooling; stochastic optimization; parcel logistics; hyperconnected networks; tactical planning; resource rebalancing.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☒ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☒ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

1 Introduction

Parcel distribution systems are entering a new era shaped by strong demand variability, urban logistics constraints, and rising expectations for sustainability and responsiveness. To cope with uncertainty, operators often deploy precautionary levels of resources, and this allocation occurs at a low frequency, leading to systematic residual capacity such as idle vehicles or underutilized infrastructure. While this redundancy improves reliability, it also increases operational costs and environmental impact. Rather than treating unused capacity as an unavoidable inefficiency, the Physical Internet (PI) paradigm (Montreuil, 2011) suggests reinterpreting it as a shared and reconfigurable resource that can circulate within hyperconnected logistics networks. In such environments, dynamic pooling becomes a central coordination mechanism enabling assets to be reassigned across locations and time periods according to system-wide needs.

Despite this conceptual shift, most existing tactical planning approaches still rely on static allocation principles documented in parcel logistics literature (Wang et al., 2022). These approaches typically separate initial deployment decisions from corrective mechanisms such as subcontracting or transfers, which prevents a consistent assessment of how early commitments affect later adaptability. Moreover, current models rarely capture the temporal evolution of internal capacity, limiting understanding of how residual capacity can be mobilized across periods. Prior studies on pooling and shared resources (Jain, 2007; Zhong et al., 2018) as well as research on dynamic rebalancing in mobility systems (Maggioni et al., 2019) demonstrate the operational value of coordinated resource flows, yet they do not provide an integrated multi-period tactical framework combining uncertainty, heterogeneity, environmental constraints, and inter-node mobility. Stochastic optimization theory provides tools to address such variability and avoid rigid planning decisions (Halperin, 2022), but unified formulations capturing all these dimensions remain scarce.

To address this gap, this research utilizes stochastic multi-period optimization approach to support PI-enabled parcel networks. The framework explicitly represents residual capacity as a mobile asset that can be redistributed through dynamic pooling across interconnected logistics nodes. A two-stage stochastic programming structure is adopted, consistent with established approaches to supply chain design under uncertainty (Santoso et al., 2005). In the first stage, decision-makers allocate vehicles of different types to hubs and time periods while optionally reserving anticipatory subcontracting capacity as a hedge against forecast errors and fluctuation in demand over time. In the second stage, once actual demand scenarios are realized, corrective actions can be activated, including additional subcontracting and inter-hub transfers. These transfers are governed by temporal conservation relations that track resource availability over time, allowing capacity to move through the network and respond to emerging imbalances. Environmental requirements related to low-emission vehicles are also incorporated.

The proposed formulation yields three main contributions. First, it reconceptualizes residual capacity as an operational lever that enhances system adaptability when coordinated across a network. Second, it integrates anticipatory and reactive decisions within a unified optimization structure, enabling systematic evaluation of cost–flexibility trade-offs. Third, it embeds fleets allocation, sustainability constraints, and uncertainty within a single tactical planning model aligned with PI principles of interoperability and resource sharing. Figure 1 illustrates our research methodology.

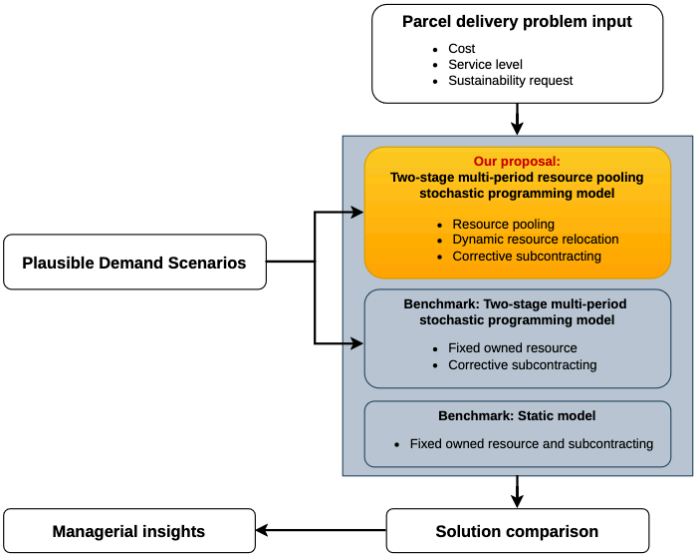


Figure 1: Research methodology.

This paper is organized as follows. Section 2 reviews the literature relevant to the problem and positions the paper within the existing research. Section 3 provides a detailed description of the problem under study. Section 4 discusses illustrative examples and reports the computational results. Finally, Section Conclusion concludes the paper and outlines directions for future research.

2 Literature background

The concept of the PI (Montreuil et al., 2012) has emerged as a promising paradigm for addressing the structural inefficiencies of traditional logistics networks. The PI envisions a system where physical goods flow seamlessly through a network of open, standardized, and interconnected nodes and resources. Modularity, interoperability, and decentralized decision-making enable logistics actors to pool capacity, share infrastructure, and dynamically reconfigure flows in response to evolving conditions.

One of the key operational implications of the PI is hyperconnectivity (Crainic and Montreuil, 2016) which is the ability of distributed logistics resources (vehicles, lockers, staff, robots, and hubs) to interact across organizational boundaries. In urban parcel delivery, hyperconnected networks facilitate dynamic pooling strategies, where capacity in one node can be leveraged to serve demand at another (Faugère et al., 2022).

Several studies have shown that embedding pooling mechanisms within hyperconnected networks can yield efficiency gains. For instance, coordinated reallocation of heterogeneous resources across facilities can lead to substantial cost reductions and service improvements, particularly during high-demand periods (Kim et al., 2021; Pang et al. 2025). Such flexibility is enabled by the PI's core principle of interconnectivity, whereby nodes are spatially linked through physical infrastructure and digitally integrated through shared information systems.

Building on the PI vision, resource allocation in parcel logistics involves the tactical assignment of delivery resources across a network to meet fluctuating and spatially distributed demand (Malairajan et al., 2013). In urban settings, this challenge is intensified by decentralized network structures, uncertain demand patterns, and the need for dynamic coordination across both time and space (Xu et al., 2024).

Modern resource allocation problems consider different dimensions: centralized versus multi-hub architectures, static versus adaptive decision frameworks, and homogeneous versus hybrid resource fleets (Ghiani et al., 2025), furthermore, multi-criteria objectives balancing economic efficiency, environmental sustainability, and service quality. Prior research has demonstrated the value of cross-node or multi-period coordination in reducing resource underutilization, and balancing loads (Dahmen et al., 2020; Saleh et al., 2025).

Despite these advances, most existing models operate under deterministic assumptions or treat residual capacity as a passive by-product rather than an actively managed asset. While the PI provides a conceptual foundation for dynamic resource sharing, formal optimization frameworks that explicitly operationalize hyperconnectivity through residual capacity pooling and corrective actions under uncertainty remain limited.

3 Problem statement

This study aims to explore the impacts of resource pooling, enabled by the PI paradigm, in supporting logistics service providers to achieve greater cost efficiency and sustainability, as well as to optimize resource deployment through the dynamic relocation of shared resources among participants, specifically, parcel delivery hubs in our case.

Network Structure and Problem Description

We adopt the hyperconnected city logistics framework proposed by Montreuil et al. (2018) and further applied by Faugère et al. (2022) and Xu et al. (2026), which discretizes urban agglomerations into three hierarchical levels: unit zones (clusters of customer locations), local cells (clusters of unit zones), and urban areas (clusters of local cells). This results in a three-tier logistics network comprising gateway hubs, local hubs, and access hubs, each designed to manage parcel flows across urban areas, local cells, and unit zones, respectively. The present study concentrates specifically on the local hub tier.

The parcel logistics service provider offers package delivery and pickup services to customers within a region. Each region is composed of multiple local cells. Each local cell exhibits varying logistics service demands across time. The service provider must fulfill all demand through a network of local hubs, utilizing different types of vehicles, including cargo bikes, diesel vans, and electric vans.

In practice, resource levels at each hub are planned annually based on demand forecasts. However, due to demand fluctuations, some hubs frequently experience residual capacity and idle resources during low-demand periods, while others may face capacity shortages that necessitate outsourcing through subcontracting. Negative demand correlations between hubs are quite common in urban logistics networks. For example, during major holiday periods, many people leave their homes to travel, which may temporarily reduce parcel deliveries to residential areas. At the same time, commercial areas in tourist destinations often experience increased activity. Similarly, hubs located near industrial zones may experience demand surges during certain production cycles, whereas residential hubs remain relatively quiet. These kinds of complementary demand patterns create natural opportunities for resource sharing across hubs.

To address the inefficiencies of underutilized resources and better leverage the interconnections among hubs, we propose a tactical resource planning strategy that integrates two key elements: (1) resource pooling among hubs, enabling shared resource utilization based on individual hub demand and system-wide total costs; and (2) dynamic resource relocation on a weekly basis, allowing reactive adjustments to resource levels at each hub in response to uncertain demand realizations.

Resource Pooling Framework

Figure 2 illustrates the proposed resource pooling strategy for parcel delivery local hubs, which involves resource relocation among hubs based on the demand of each local cell. Under this framework, the resources in this network are pooled together and the stakeholder can coordinate resource levels across hubs by transferring resources from hubs with residual capacity to hubs whose demand exceeds their handling capacity.

We consider a set of local hubs N operated by a single stakeholder. Let $G = \{N, A\}$ denote the network of hubs with A as the complete ser of directed arcs between hubs in N. The planning horizon T is one year, divided into 52 periods. Eeach period is one week, i.e., $T = \{1,2,\ldots,52\}$.

Resource levels in the pooled hub network are determined at the beginning of each year, including the initial resource deployment at each hub and the planned subcontracting quantities for each period, based on forecasted demand. Three types of resources are pooled: diesel vans, electric vans, and cargo bikes. All resource types can be used to fulfill logistics services at any local hub. Notably, a minimum portion of demand must be satisfied by green vehicles (cargo bikes or electric vans) due to sustainability requirements.

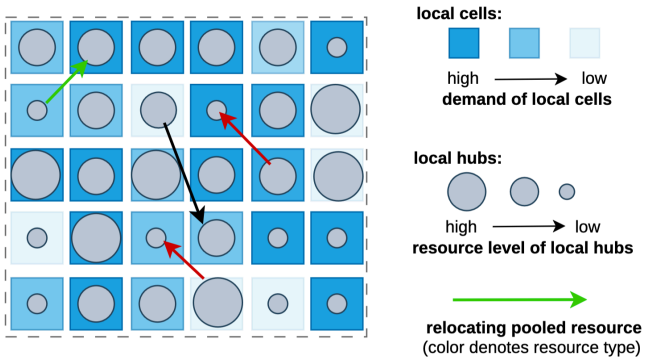


Figure 2: Resource pooling in the network of local hubs considering demand in the local cells.

Dynamic Relocation

While resource levels are initially determined based on annual demand forecasts, actual demand at each hub is uncertain and often deviates from forecasts. Moreover, demand arrival patterns vary across hubs. The probability distribution of demand at each hub is assumed to be known from historical data. Under such uncertainty, dynamic resource deployment becomes an effective strategy to adjust resource levels at each hub to meet realized demand.

With pooled resources, we can relocate resources between hubs at the end of each period (week) to support hubs in need, without immediately resorting to purchasing or outsourcing from external providers. Figure 3 illustrates tactical resource relocations within the hub network over three periods, using one resource type as an example. As shown, the resource level at each hub (denoted by circle size) adjusts in response to demand changes (denoted by square size) at each period, reflecting resource inflows or outflows.

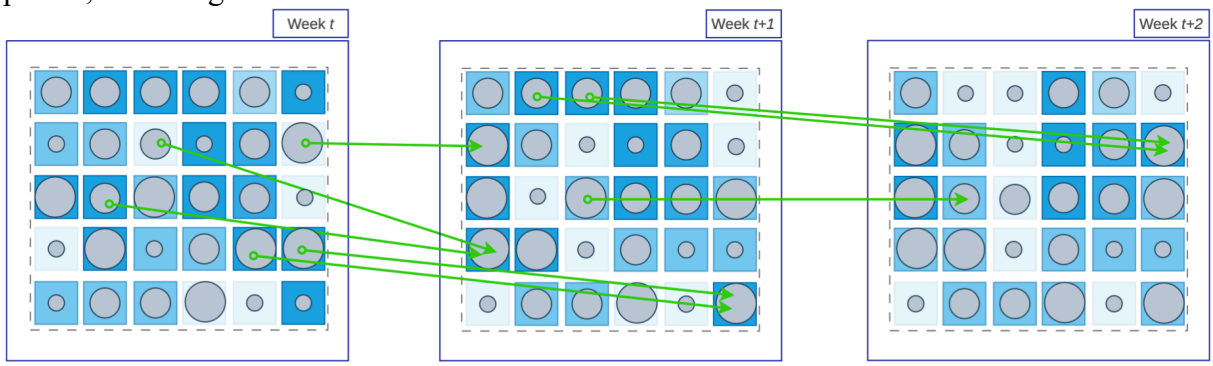


Figure 3: Illustration of dynamic relocation of pooled resources (example of single type resource).

Two-Stage Stochastic Programming Model

To address this tactical planning problem, we propose a two-stage multi-period stochastic programming model that accounts for demand uncertainty with the objective of minimizing total expected costs. First-stage decisions: to determine the initial resource level of each type at each hub, as well as the planned subcontracting quantities for each period of the year, considering purchasing and subcontracting costs. Second-stage decisions: after demand scenarios are realized, to determine resource flows between hubs on a weekly basis, along with any additional subcontracting required to fulfill the observed demand.

Within this tactical resource pooling and relocation framework, costs are incurred in four categories: (1) first-stage: purchasing new resources at the beginning of the planning horizon. (2) first-stage: planned subcontracting from rental services to obtain additional resources. (3) second-stage: corrective operational movements of resources for relocation between hubs under

each demand scenario. (4) second-stage: corrective subcontracting of resources based on updated demand information for the upcoming period.

4 Illustrative examples

To evaluate the effectiveness of the proposed multi-period resource pooling strategy (MRP), we benchmark its performance against an optimal static resource deployment plan (STATIC) and an optimal multi-period resource deployment plan without considering resource pooling (MNP). The benchmark STATIC is obtained by solving a static variant of the model without corrective subcontracting and resource relocation. The benchmark MNP is obtained by solving a variant of the model that considering multi-period corrective subcontracting but no resource pooling. Both benchmarks operate under the constraints of fully satisfying all realized demand while adhering to the prescribed sustainability requirements, thereby serving as a baseline.

Representative numerical experiments were designed based on a multi-hub logistics network. The example captures the key trade-offs between resource investment (first stage) and operational flexibility under demand uncertainty (second stage) with different resource deployment strategies: resource pooling and adaptive resource relocation, or fixed resource deployment level.

The hub network for parcel delivery consists of three hubs. The planning horizon spans 52 weeks, allowing for analysis of long-term fleet deployment under stochastic demand. Three types of resources for delivery are considered, reflecting a transition toward sustainable urban logistics: Type 1: Cargo bikes with low cost, low emissions, and low capacity. Type 2: Electric vans with higher capacity, low environmental impact but higher cost. Type 3: Diesel vans with higher capacity but higher cost and environmental impact. This case allows a maximum purchase of M_v units for each resource type, calculated as follows:

$$M_v = \left\lceil \frac{R_static}{3} \right\rceil$$

where R_static is the total number of vehicles of all types used in the optimal plan of STATIC strategy. In our instances, M_v is set as. 617. The cost structure is listed in Table 1.

Table 1: Parameter setting.

Resource Type	Capacity (Units)	Purchase Cost (€)	Movement Cost (€ per move among hubs)	Subcontracting Cost (€ per week)
Cargo bike	650	44,000	10	1,950
Electric van	910	49,000	50	2,730
Diesel van	910	51,000	50	2,730

Mean value of the demand (in parcels) at each hub take value randomly from interval [300000, 400000] for each hub respectively. Demand scenarios are generated for each week each hub, with a total of 100 scenarios of equal probability ($p^s = \frac{1}{|S|}, \forall s \in S$). Each scenario is generated by adding a noise following gauss distribution $\mathcal{N}(0,1)$ to the forecasted demand $d_{i,t}$. The sustainability baseline is set as $g_i = 0.3$, representing a target share of demand met by sustainable resources (e.g., cargo bikes) in the fleet group.

The proposed model was solved using Gurobi Optimizer version 13.0.1, conducted on a system equipped with an Apple M1 processor, featuring 8 physical cores and 8 logical processors. A

total of 10 random instances were generated following the demand and scenario generation rules described above. The solving time is 17 seconds.

Figure 4 compares the total costs of the 3 strategies for the example instance, where the MRP and MNP models report per-scenario expected costs across 100 scenarios, while the Static model provides a deterministic total cost. The MRP model achieves a mean per-scenario cost of 68.92 million euros, compared to 71.46 million euros for MNP and 113.85 million euros for the STATIC benchmark. Both stochastic models substantially outperform the static approach, with MRP delivering a 39.5% cost reduction over Static. Moreover, MRP outperforms MNP by approximately 3.6% (2.55 million euros on average), demonstrating that resource pooling and dynamic relocation yield measurable efficiency gains beyond stochastic planning alone. Furthermore, both MRP and MNP present acceptable standard deviations (MRP: 0.72 million euros; MNP: 0.79 million euros). This indicates consistent performance across demand scenarios, with MRP exhibiting slightly greater robustness.

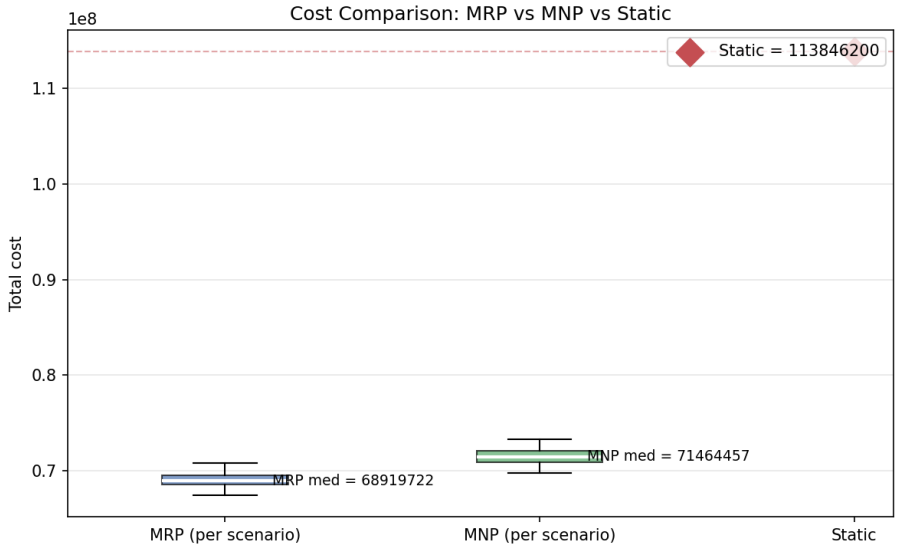


Figure 4: Costs comparison between MRP, MNP, and STATIC strategy.

Figure 5 demonstrates the deployment resource level (without counting resource from subcontracting) at each hub with periods. Figure 6 illustrates the level of subcontracting with time at each hub. It suggests that MRP enables active resource reallocation based on demand fluctuations and avoids residual resources while meeting service requirements. Furthermore, it reveals that the MRP model can reduce reliance on subcontracting by dynamically transferring resources among hubs with negatively correlated demands.

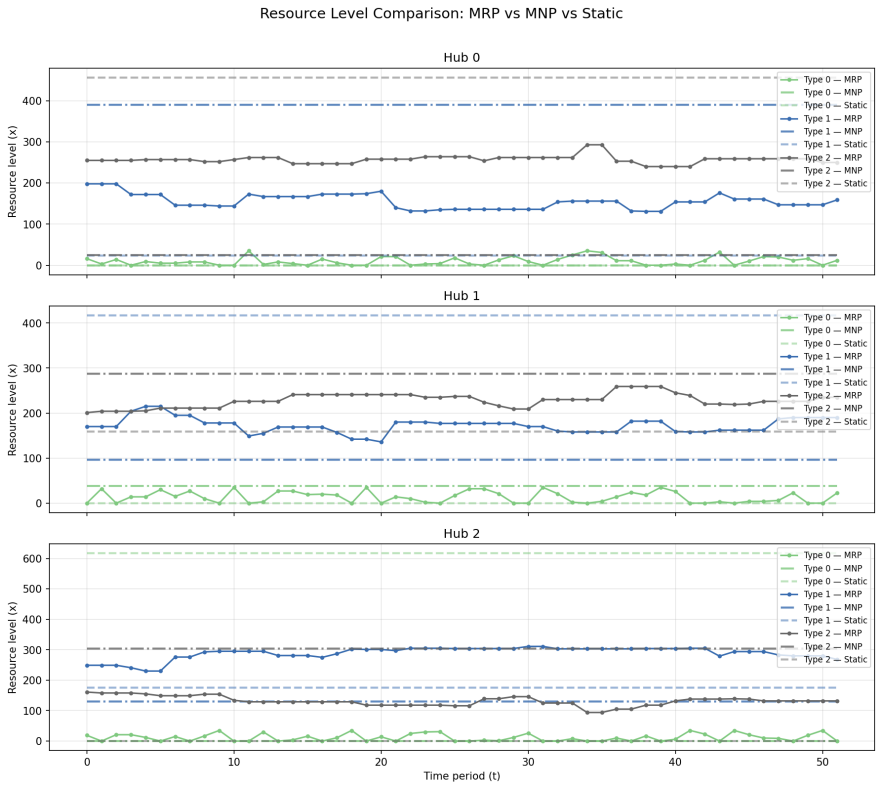


Figure 5: Resource level in each hub.

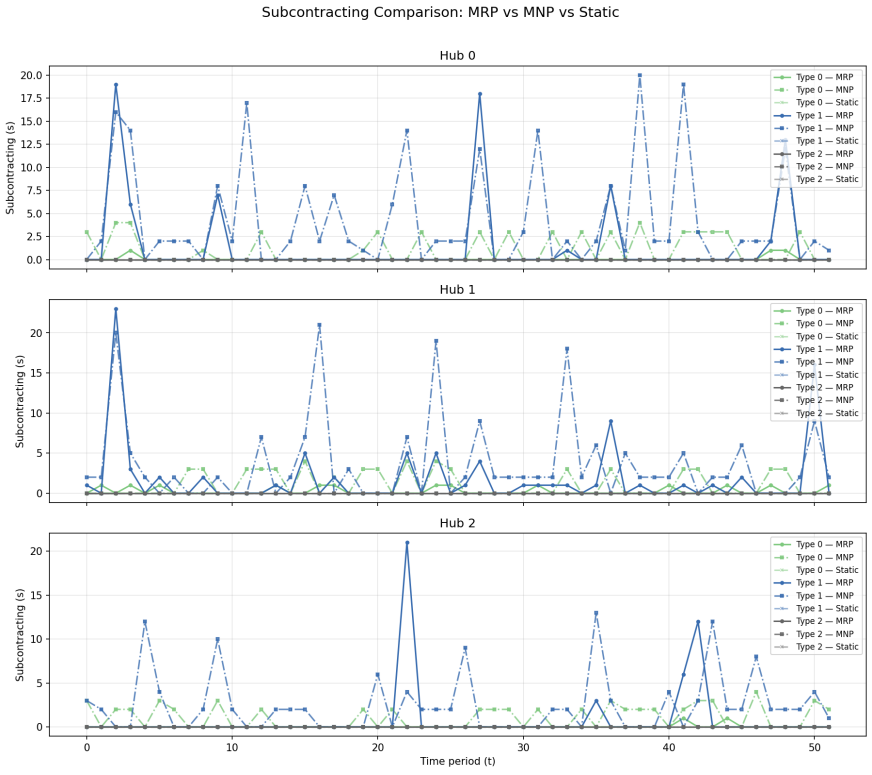


Figure 6: Subcontracting resource level in each hub.

The results also reveal that model MRP can promotes greener operations. It maximizes utilization of owned green vehicles. The dynamic reallocation of owned resources, particularly electric vans and cargo bikes, ensures that green vehicles are utilized more efficiently across the network rather than sitting idle at low-demand hubs. Moreove, it helps meet the

sustainability baseline target by ensuring a minimum portion of demand is served by environmentally friendly resources.

Overall, the results showcase that the proposed resource pooling strategy and two-stage dynamic resource relocation stochastic programming model deliver significant cost savings, enhanced sustainability, and stable performance, offering a practical strategy for logistics providers to enhance efficiency under demand uncertainty.

Conclusion

This study proposed a two-stage multi-period stochastic programming model with resource pooling and dynamic relocation (MRP) to address tactical resource allocation in parcel logistics under demand uncertainty. Within the hyperconnected logistics framework, the model optimizes resource level and deployment across local hubs, which balances first-stage tactical decisions with second-stage operational adjustments.

A stylized application inspired by a real parcel operator illustrates the approach. The experimental setting includes multiple hubs, several planning periods, and different vehicle types subject to fluctuating demand patterns. The results show that MRP enables active resource reallocation across hubs with negatively correlated demands. When dynamic pooling is enabled through controlled transfers, a small number of reallocations significantly reduces external capacity needs, demonstrating how residual capacity circulation can improve performance without increasing total assets.

Beyond numerical gains, the findings emphasize a key principle of the PI: logistics performance depends less on static resource ownership than on the ability to orchestrate shared assets across a hyperconnected network. By quantifying the operational value of dynamic pooling and residual capacity mobility, this study provides analytical foundations for designing parcel distribution systems that are more resilient, efficient, and sustainable. It thus contributes to the development of next-generation logistics networks in which coordination, interoperability, and adaptive resource sharing form the core operating logic. Several directions remain for future research. First, the model could be validated using real-world operational data to confirm its practical applicability. Second, the framework can be extended to multi-stakeholder settings to investigate the potential of resource pooling across competing or collaborating logistics providers. Third, workforce allocation including driver scheduling and shift planning could be integrated into the model to capture additional operational constraints.

References

Crainic, T. G., & Montreuil, B. (2016). Physical internet enabled hyperconnected city logistics. Transportation Research Procedia, 12, 383-398.

Dahmen, S., Rekik, M., Soumis, F., & Desaulniers, G. (2020). A two-stage solution approach for personalized multi-department multi-day shift scheduling. European Journal of Operational Research, 280(3), 1051-1063.

Faugère, L., Klibi, W., White III, C., & Montreuil, B. (2022). Dynamic pooled capacity deployment for urban parcel logistics. European Journal of Operational Research, 303(2), 650-667.

Ghiani, G., Guerriero, E., Manni, E., & Pareo, D. (2025). Combining autonomous delivery robots and traditional vehicles with public transportation infrastructure in last-mile distribution. Computers & Industrial Engineering, 203, 111001.

Halperin, I. (2022). *Reinforcement Learning and Stochastic Optimization: A Unified Framework for Sequential Decisions: by Warren B. Powell (ed.), Wiley (2022). Hardback. ISBN 9781119815051* (Vol. 22, No. 12, pp. 2151-2154). Routledge.

Kim, N., Montreuil, B., Klibi, W., & Kholgade, N. (2021). Hyperconnected urban fulfillment and delivery. Transportation Research Part E: Logistics and Transportation Review, 145, 102104.

Jain, A. (2007). Value of capacity pooling in supply chains with heterogeneous customers. *European journal of operational research*, *177*(1), 239-260.

Maggioni, F., Cagnolari, M., Bertazzi, L., & Wallace, S. W. (2019). Stochastic optimization models for a bike-sharing problem with transshipment. *European Journal of Operational Research*, *276*(1), 272-283.

Malairajan, R. A., Ganesh, K., Muhos, M., & Anbuudayasankar, S. P. (2013). Class of resource allocation problems in supply chain–a review. International Journal of Business Innovation and Research, 7(1), 113-139.

Montreuil, B. (2011). Toward a Physical Internet: meeting the global logistics sustainability grand challenge. *Logistics Research*, *3*(2), 71-87.

Montreuil, B., Meller, R. D., & Ballot, E. (2012). Physical internet foundations. IFAC Proceedings Volumes, 45(6), 26-30.

Montreuil, B., Buckley, S., Faugere, L., Khir, R., & Derhami, S. (2018). Urban parcel logistics hub and network design: The impact of modularity and hyperconnectivity. 15th IMHRC Proceedings (Savannah, Georgia. USA – 2018)

Pang, Y., Klibi, W., & Labarthe, O. (2025). Resource Pooling in Hyperconnected Warehouse Networks. 11th International Physical Internet Conference (IPIC), Hong Kong. https://doi.org/10.35090/gatech/13245

Saleh, M., Lauras, M., Klibi, W., & Leveque, J. (2025). Pooling Residual Delivery Capacity in Urban Parcel Logistics. IFAC-PapersOnLine, 59(10), 2153-2158.

Santoso, T., Ahmed, S., Goetschalckx, M., & Shapiro, A. (2005). A stochastic programming approach for supply chain network design under uncertainty. *European Journal of Operational Research*, *167*(1), 96-115.

Wang, Y., Zhao, L., Savelsbergh, M., & Wu, S. (2022). Multi-period workload balancing in last-mile urban delivery. *Transportation Science*, *56*(5), 1348-1368.

Xu, Y., Klibi, W., & Montreuil, B. (2026). Dynamic deployment of pooled human-robot resources in urban parcel logistics. Transportation Research Part B: Methodological, 205, 103409.

Zhong, Y., Zheng, Z., Chou, M. C., & Teo, C. P. (2018). Resource pooling and allocation policies to deliver differentiated service. *Management Science*, *64*(4), 1555-1573.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE
cetic Your Connection to ICT Research

Towards Federated Routing Optimization for Road Transport

Renaud De Landtsheer and Quentin Meurisse
CETIC - Center of Excellence in ICT (Belgium)
Department of Combinatorial Algorithmics (COAL)
Corresponding author: renaud.delandtsheer@cetic.be

Abstract: *Road transport faces a major efficiency problem, evidenced by the high rate of empty truck movements in Europe (21% in 2024). Widespread horizontal collaboration, which offers significant profit gains, is typically hindered by the challenge of consolidating routes among many small and medium-sized enterprises (SMEs) without compromising sensitive competitive data. This position paper proposes an algorithm and protocol to implement a federated optimization system for road transport that resolves this data confidentiality barrier. Using this system, transport companies can simultaneously optimize vehicle routes and trade Transport Orders (TOs) with partners and marketplaces. This enables them to automatically and optimally accept or subcontract TOs to consolidate routes and improve global efficiency, addressing fuel, societal, and economic concerns. Crucially, the system ensures confidentiality by not requiring the sharing of private information, such as internal fleet data or confidential TOs, with any third parties, which is a key requirement for its adoption. Besides outlining the proposal, the paper identifies and discusses related challenges, notably regarding the fair pricing policy of subcontracted TOs and the global impact of the system on average TO prices. Our proposal could serve as a basis to build a Physical Internet, focusing on road transport at first.*

Keywords: *Routing optimization, Federated optimization, Local Search optimization, Road Transport, Auction Systems, Physical Internet, Game Theory*

Physical Internet (PI) Roadmap Fitness: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:☐ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),☐ Transportation Equipment, ☐ PI Networks, ☒ System of Logistics Networks, ☐ Vertical Supply Consolidation, ☐ Horizontal Supply Chain Alignment, ☒ Logistics/Commercial Data Platform, ☒ Access and Adoption, ☐ Governance.

Targeted Delivery Mode-s: ☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation

1 Introduction

Empty truck movements in Europe reached 21% in 2024 (Eurostats, 2025), largely due to the many SMEs in the logistics sector (80% of EU transport companies have <10 trucks), hindering consolidation. Horizontal collaboration can increase transport companies' profits by 20–30% by reducing empty trips (Cruijssen *et al.*, 2007; Montoya-Torres *et al.*, 2016; Soysal *et al.*, 2018).

There are services of request exchange among trucking carriers, like WTransnet (https://www.wtransnet.com/) and Trans.eu (https://tfc.trans.eu/en/), which provide secured platforms for putting carriers in contact with each other. These platforms can be used for peer-to-peer exchange of requests; however, they do not provide tools for organizing vehicle routes of their members or selecting TOs to best complete partial routes.

Inspired by the Physical Internet concept, this position paper proposes a technical approach to automatically consolidate truck routes among collaborating transport companies using optimization tools while maintaining data privacy. This is achieved through a "social network" and an optimizing communication protocol. The protocol enables the route

optimizer in each company's Transportation Management System (TMS) to identify the best subcontractors (via pre-defined "partnership" links) and the most profitable Transport Orders (TOs) to take.

The core is an iterative auction protocol, similar to a local search optimization for vehicle routing. Transporters bid on visible TOs based on marginal cost, leading to temporary assignments. A TO is reassigned only if another transporter has a lower marginal cost for it in a subsequent auction round. This system functions as a federated routing optimization algorithm, where the TMS of each transporter optimizes their routes and trades TOs without sharing sensitive data as illustrated in Figure 1. Transporters can also keep some TOs confidential.

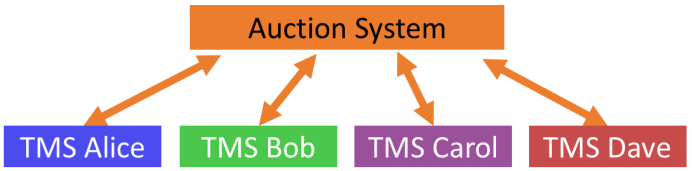


Figure 1: Optimization happens through a federated optimization algorithm

Like all local search heuristics, the iterative auction can get stuck in sub-optimal local minima. To address this, similar to how ejection chains improve simple TO assignment, we propose considering the use of more complex moves like the Very Large Scale Neighborhood heuristic, which is known to handle strong constraints more effectively (Ahuja, 2002).

The objective of this paper is to delineate the concept of federated routing optimization, focusing primarily on enumerating the essential requirements for its enforcement and the associated research challenges. Furthermore, we have developed a preliminary Proof of Concept (PoC) demonstrating a potential implementation of such a system. This PoC will facilitate the study of our proposal's efficiency, optimality, and scalability, and enable experimentation with other research challenges, such as the price agreement mechanism.

The structure of the remainder of this paper is as follows: Section 2 provides a review of the state of the art concerning various approaches enabling similar collaboration among transporters. Section 3 introduces a simplified example used to illustrate the concept. Section 4 presents the mandatory requirements that the protocol must satisfy for practical application. Section 5 formalizes the optimization and trading protocol. Section 6 describes an illustrative PoC developed to experiment with the concept. Section 7 concludes the paper.

2 Literature Review

The studies by Gansterer and Hartl (2018, 2020) provide two comprehensive surveys on collaborative routing. Their primary focus is on horizontal collaboration, defined as cooperation where enterprises operating at the same level of a supply chain share resources. They delineate two categories of collaboration: centralized and decentralized.

In a centralized collaboration framework, a central decision maker possesses access to all pertinent information. Centralized collaboration aspects can be readily integrated into optimization problems. These aspects can function as a new objective or constraint, thereby allowing the utilization of established methodologies such as local search, branch-and-cut algorithms, and similar techniques.

(Defryn and Sörensen, 2018) demonstrate that horizontal collaborations are advantageous even when the carriers maintain different or even conflicting objectives. Generally, research indicates that centralized collaborative approaches augment profits (the discrepancy between revenues and costs) by 20-30% compared to non-collaborative methods (Cruijssen et al., 2007; Montoya-Torres et al., 2016; Soysal et al., 2018).

Conversely, in decentralized collaboration, the decision maker holds only partial information. (Gansterer and Hartl, 2018; 2020) emphasize the substantial potential of auction systems for facilitating decentralized collaboration.

It has been established that aggregating requests into a bundle enhances the efficiency of the auction mechanism. These bundles are subsequently placed into single-shot auctions, meaning they are definitively allocated to the most economical transporter. (Gansterer *et al.*, 2020), in particular, developed an approach where TOs are initially bundled based on their geographical location. This methodology yields significant improvements in both processing speed and solution quality compared to a single-shot auction conducted without bundling. This approach is only viable if the auctioned TOs are published and accessible to all participating transporters. Moreover, while this method is suitable for last-mile delivery, where a high volume of TOs is available for aggregation, it may not be appropriate for other contexts, such as Full Truck Load (FTL) transports across Europe. This limitation stems from the European cabotage rule, which, broadly speaking, restricts trucks to a maximum of six stops per trip.

Iterative auctions appear to refine the exchange mechanism. In each iteration, carriers submit a bid based on feedback derived from preceding iterations. Such a protocol has the potential to increase the collaboration profit (Lyu *et al.,* 2019).

Analogous to centralized collaboration, decentralized approaches have been demonstrated to increase profitability and mitigate CO2 emissions (Nataraj 2019; Yao 2019).

Recently, Los et al. (2025) experimented with an auction system comprising 1000 carriers, based on real-world problems. Their findings demonstrated that travel cost gains of 77% can be achieved. Bundling can further enhance the results by 13%.

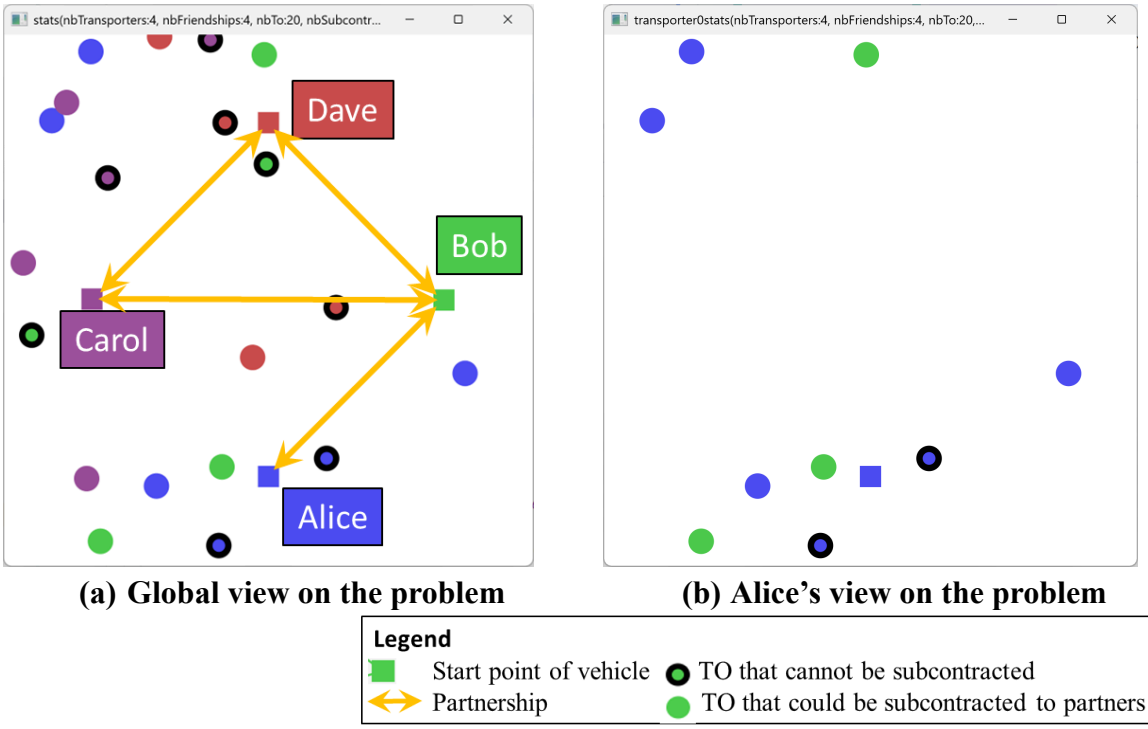
To summarize, the primary limitations associated with public auction systems are:

- **Loss of Confidentiality**: Auction system compromises confidentiality because TOs are disclosed to all participants. This pervasive exposure enables competing organizations to potentially monitor each other's transport operations through continuous observation of the auctions. Furthermore, the expanded scope of data sharing exacerbates the challenges associated with maintaining compliance with the General Data Protection Regulation (GDPR).
- **Disruption of the Business Ecosystem**: The shift from directly sending TOs to transporters to entering them into auction systems disrupts established customer-transporter business relationships. This change may be met with reluctance from transporters.
- **Loss of Control**: Carriers relinquish control over the specific transporter that will execute the order they have sold, leading to a diminished level of trust in the resulting economic arrangement.

3 Toy Example of Last Mile Delivery optimization

We illustrate the concept with a simple last mile delivery problem shown in Figure 2(a). Alice, Bob, Carol and Dave are transporters on the same territory. They must reach points

spread in the territory that represent TOs. Each TO is owned by a transporter as represented by its color. Furthermore, transporters have partnership agreements. Alice has a partnership agreement with Bob, but not with Carol and Dave. Bob, Carol and Dave have a partnership agreement with each other. These agreements formally capture that the system could subcontract TOs following these partnership links. Transitive subcontracting is not allowed. TOs are either considered private, meaning that the transporter that currently owns the TO does not want to subcontract it to anyone, or public, meaning that the TO can be subcontracted to a partner of its owner. To make it more realistic, the transporters have capacity constraints. There is a maximal number of TOs that they can serve and a maximal distance that they can drive.



(a) Global view on the problem **(b) Alice's view on the problem**

Figure 2 Collaborative routing optimization problem

The key contribution here is to solve it without using a trusted third party; the optimization must be performed through communication established among the TMSs of these transporters. To illustrate this, Figure 2(b) shows the perspective of Alice on the right side. Her TMS and route optimizer only see this part of the routing problem. This includes her vehicle start point, all her TOs and Bob's public TOs since Bob is her sole partner. Actually, the global view is not available because none of the actors possesses a complete view of the global optimization problem. The federated optimization system, where TMSs can talk to each through an auction system is illustrated in Figure 1.

4 Requirements on a Federated Routing optimization System

The key technical research challenge is to develop the proper communication protocol and algorithms so that route optimization engines can connect with each other and perform route optimization while trading transport orders. We have identified the following main requirements to address SME adoption barriers:

- **Confidentiality**: For any transporter, the start points and routes of their trucks, and their private TOs must not be disclosed to any other transporter nor to the auction system.

Their public TOs are only disclosed to their selected partners. The auction system should not know the public TOs either. This is a strong requirement further supported by EU regulations (GDPR, IA-act).

- **Optimality:** The system should deliver a good solution compared to the segregated situation where transporters do not cooperate with each other. The federated optimization should not degrade too much compared to a centralized optimization approach.
- **Scalability**: The system should scale up to the number of transporters and TOs of the considered market segment, which could be based on a territory and on a type of transport. Logistics is split into different and incompatible market segments like FTL first/last mile, parcel, pallets etc.; each of these segments can exploit a specific instance of the system.
- **Efficiency:** The federated optimization algorithm should converge fast enough to be usable in practice.
- **Integration of domain constraints:** The system should take constraints into account, such as capacity constraints of trucks, or maximum availability of trucks or the EU cabotage rules that limits an international trip to a maximum of six stops within the EU. These constraints must be handled at the level of the TMSs because they are related to confidential aspects like the capacity of trucks. However, the value of these constraints associated with TOs must be expressed and communicated among transporters through a standard exchange format. It should also be able to accommodate structural constraints like FTL problems that have an origin and a destination or pick-up and delivery problems where transported parcels can share a portion of the route.
- **Fair and efficient price agreement**: The result of the trading is that TOs are assigned to transporters based on the global efficiency of the system. The value that a transporter will get from the system should be higher than if she were not using the system. This encompasses the TOs that are exchanged and the overall pricing agreement. There is also a more global notion that the average value of a TO should not be driven down by the system, which would be detrimental to all transporters. Similarly, it should not be driven upwards by the system.
- **Market non-distortion**: The system should not make it possible for a transporter to attract TOs and drain them from other transporters. To this end, constraints can be added on the partnership links. For instance, they could specify a maximal number of TOs to be transferred, that TOs can be exchanged against other TOs, or that the value of transferred TOs should be similar.
- **Security:** Beyond market distortion, the system should also be robust against malicious actors, which could, for instance, send over- or under-valued auctions into the protocol either to manipulate the price or the result, or to prevent the convergence of the algorithm. Coalitions of actors should also be considered.

5 A Trading and Optimization Protocol for Transport Orders

We propose formalizing the collaborative routing problem as an auction system where transporters bid on TOs. To achieve optimal outcomes, the system is designed around the concept of iterative auctions: transporters repeatedly submit bids on visible TOs, and the TOs are temporarily assigned to the transporter offering the lowest marginal cost. The auction is then restarted and TOs can be moved to another transporter if their marginal cost changes. The key mechanism is that the marginal cost for a TO is influenced by the presence of other TOs typically because they are geographically close to each other, so that the transporter can consolidate his route.

This iterative auction process closely mirrors classical local search optimization, a technique widely used for handling vehicle routing problems (Glover and Kochenberger, 2003). This should enable transporters to collect TOs that are close to each other throughout the iterations, and massify the transports.

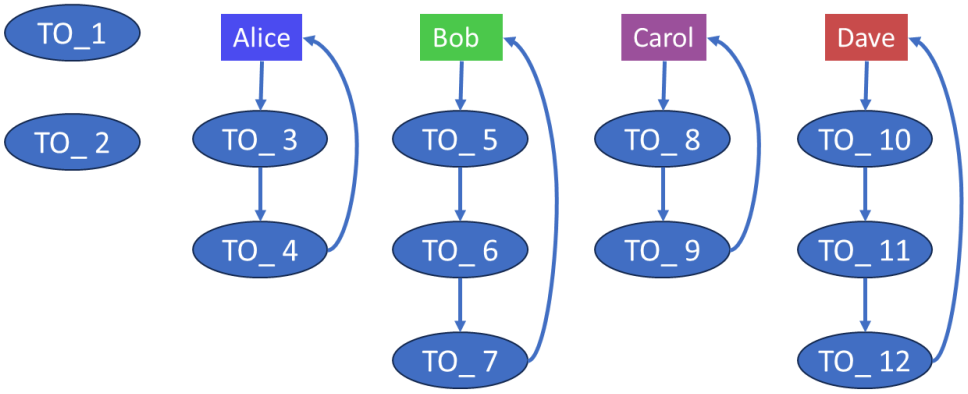


Figure 3 Temporary assignment of TOs to transporter and logical routes

Pure local search however will encounter the issue of local minima, where the current TO assignment is sub-optimal but cannot be improved by transferring a single TO between transporters. This notably arises due to strong constraints, such as the maximum availability of a vehicle over the day or a maximal number of TOs that a vehicle is allowed to serve in a route (EU cabotage rule). To address this issue, we re-use established concepts such as ejection chains and the more efficient Very Large Scale Neighborhood (VLSN) search that have been developed to address such issues (Ahuja et al, 2002).

The VLSN is similar to the iterative bidding described here above with an additional swap bid. A *swap bid* is a bid made by a transporter on a TO provided another TO that is currently assigned to her is reassigned to another transporter or unrouted. The value of this swap bid is the marginal difference of cost when performing this complex move. The VLSN searches for ejection chains made of coherent swap bids and simple bids described in Figure 4.

Similarly to the iterative auction, the VLSN repeatedly performs three main steps:

1. **Exploration of Atomic Moves:** This step explores atomic moves that re-assign a node from one route to another (classical bids) and moves that re-assign a node from one route to another under the assumption that a node from the target route has been reassigned elsewhere (swap bids). This information is summarized in a directed graph, termed the *VLSN graph*, where nodes correspond to those of the routing problem, and bids are represented by annotated directed edges. The annotation includes the marginal cost of the bid and information about the reassignments.
2. **Cycle Detection:** The process identifies cycles within this VLSN graph where the summation of the weights of the edges is negative. This is critical because swap bids are inherently infeasible in isolation; the ejected TOs must be assigned to another transporter who submitted a corresponding simple or swap bid. Global feasibility requires identifying a chain of bids that are collectively coherent, which translates to finding a cycle in the VLSN graph. Although finding such a cycle is an NP-complete problem, efficient heuristics with acceptable complexity have been proposed.
3. **Assignment Execution:** The assignments represented by the cycles are executed, reassigning the TOs according to the selected bids/edges of the VLSN graph

By reformulating the VLSN approach as an auction system, we split the VLSN into two distinct, communicating algorithms as shown in Figure 1. A central node communicates with the Transportation Management System (TMS) of each participating transporter via a defined protocol. The protocol is about TMSs sending bids to the central auction node which in return sends temporary assignments to the TMSs and coordinates the auction cycles described here above.

Figure 4 illustrates typical cycles identified by the VLSN on the temporary assignment in Figure 3. Figure4(a) depicts a simple cycle representing the assignment of an unrouted node to a transporter. Figure4(b) shows a more complex move where TO_3, assigned to Alice, is reassigned to Bob, which permits Alice to accommodate the unrouted TO_2. Figure4(c) details an ejection chain involving three swap moves among Bob, Carol, and Dave. The black node conceptually represents the set of all unrouted nodes. Given that the objective function is the cost function, the value of the edge emanating from the black node is proportional to the monetary value the TO giver is willing to pay for the order's execution. The black node is a mathematical artifact that enables us to represent the action of adding a previously unrouted TO into a transporter's route. The edge leaving the black node captures that the TO give will put money into the system, hence greatly improves over the objective function.

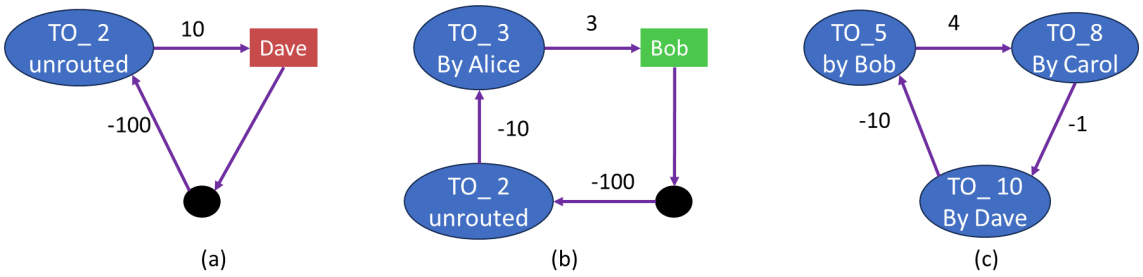


Figure 4 Example of ejection chains identified by the VLSN

Despite the decentralized nature of the system, the auction mechanism requires a central node to coordinate the auction. While the individual TMSs protect sensitive data such as truck routes and fleet capacity, the marginal costs submitted as bids to the central node introduce the risk of indirect data leakage. These marginal costs are the only information that each transporter has to share with the auction system. The central node must receive and process these marginal costs to identify cycles in the VLSN graph that have a negative total cost. If a transporter's marginal cost for a TO is exceptionally low, it might allow a malicious central node to infer the existence of nearby destinations or a high consolidation potential, thereby potentially compromising the transporter's confidential cost function. This information is known to the central node, and potentially to all transporters that can bid on the concerned TO. However, it is not clear if this is to be considered as a confidentiality breach or as a price transparency; the acceptable trade-off between confidentiality and transparency might be application-dependent.

6 Illustrative Proof of Concept

An illustrative PoC was implemented to solve the example shown in Figure 2. It is based on OscaR.cbls (https://github.com/cetic/oscar-cbls) and uses Pekko (https://pekko.apache.org/) to implement the communication protocol among actors.

The resulting solution, illustrated in Figure 5, details the optimized vehicle routes for each transporter and the subcontractions of TO. The figure is divided into the perspectives of the

four transporters. Each perspective displays the optimized route for the transporter and the TOs she could bid on. Figure 5 also outlines the identified subcontractions with yellow dashed arrows, showing the flow from the initial TO owner to the selected subcontractor.The auction system is recalled at the center to remind that the optimization was performed in a federated way.

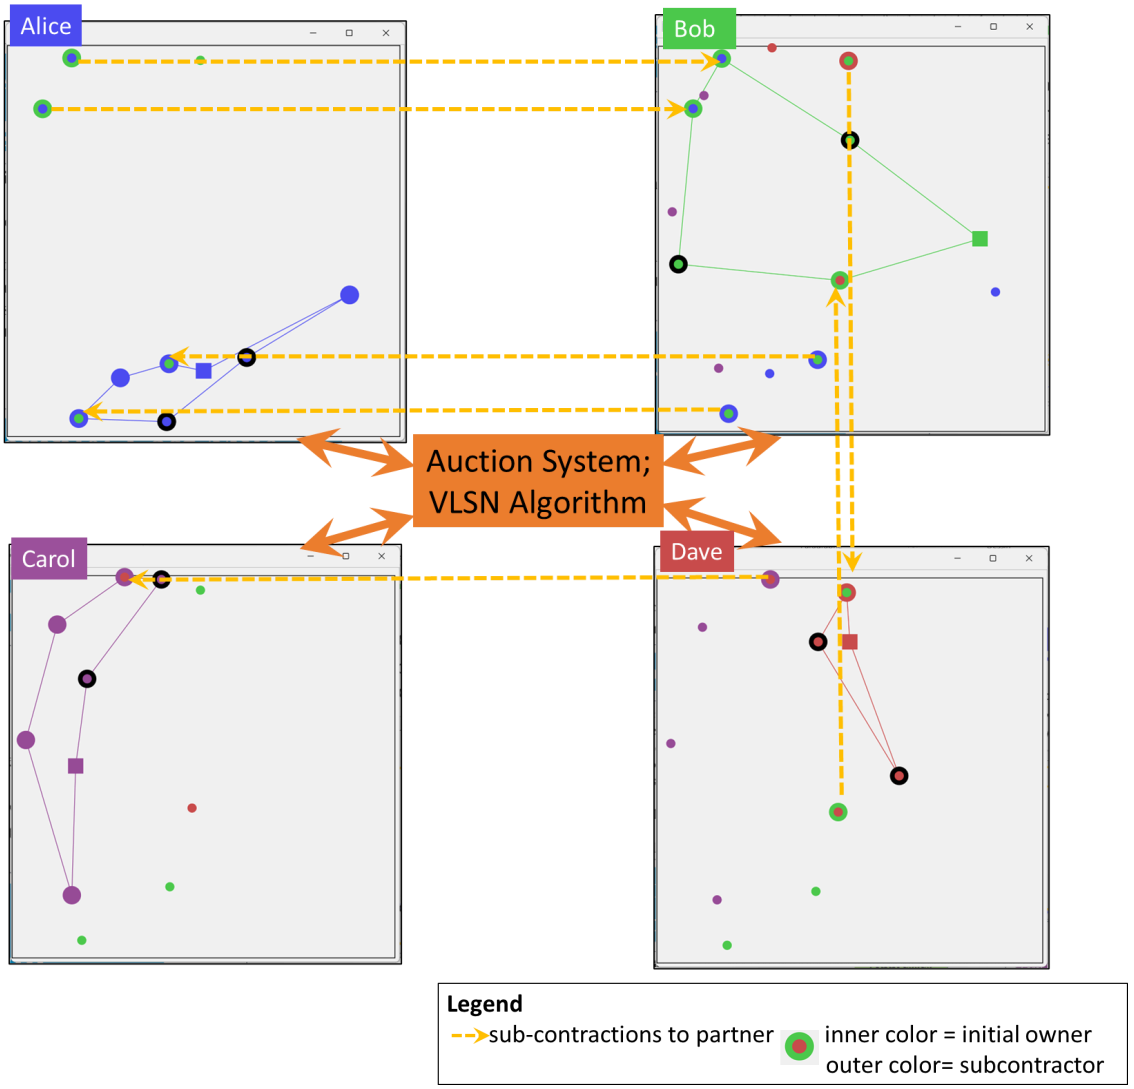


Figure 5 : the resulting optimization for the simple example

A key outcome in Figure 5 is the bilateral exchange between Alice and Bob: Alice subcontracted two TOs to Bob (geographically positioned in the northern sector) and, in return, accepted two TOs from Bob, thus facilitating a net reduction in her routing costs. Additionally, the optimized solution includes the reassignment of one further TO between Bob (who maintains partnership links with all other transporters in this example) and Dave. Dave also subcontracted a TO to Carol. These reassignments involve public TO and are consistent with the pre-established partnership agreements shown in Figure 2.

Based on this example and our prototype, we formulate the following observations:

- The distributed algorithm converges rapidly, within 2 seconds on this toy example. The majority of this execution time is attributed to message logging and graphical

rendering.

- The number of completed auction rounds is non-deterministic, ranging from 7 to 12.
- The central node has flexibility in its cycle detection strategy for temporary TO reassignment, such as prioritizing the smallest cycle or searching for the best cycles first. These algorithmic choices may influence the final solution quality and the potential for subcontracting.
- The current implementation did not account for the overhead associated with subcontracting. Incorporating a price overhead to account for administrative work related to subcontracting is necessary for a realistic assessment.
- The auction protocol can be further enhanced by incorporating other mechanisms, such as the bundling mechanisms proposed by Gansterer *et al* (2020). For instance, the two TOs that Bob subcontracted to Alice in Figure 5 could be bundled before being put into the system as a single item. However, these bundles require adaptation, as each transporter has a unique view of the global optimization problem, so that the relevant bundles may differ for each transporter.

7 Conclusion

This position paper proposes a federated optimization system using an iterative auction protocol to implement horizontal collaboration among road transport companies. The goal is to automatically and optimally consolidate truck routes and trade Transport Orders (TOs) based on marginal cost, significantly reducing the high rate of empty trips (21% in Europe in 2024) and improving global technical efficiency.

The system ensures confidentiality by operating decentrally, preventing the disclosure of sensitive data like private TOs and routes to competitors or central parties It only requires declaring a marginal cost to the auction system. The approach incorporates advanced heuristics like ejection chains and Very Large Scale Neighborhood (VLSN) search to overcome local minima. Such optimization could be applied to numerous sectors from logistics (such as FTL, specific transport), transportation (local on-demand transport) or services (home nurse).

The proposed PoC is still in a preliminary stage and should be benchmarked with respect to quality and scalability. This proposal exclusively focuses on the optimal subcontractions and route optimization. In some contexts, a fair and efficient price agreement mechanism should be integrated in the auction system. Such price agreement can typically be based on Fisher Market model where the price is determined after the auction based on some constraint system or on the simpler Vickrey auction where the value of a TO is the second smallest bid that was placed on the TO (Nissan *et al.* 2007).

Acknowledgement: This research was financed by the service public de Wallonie Recherche under grant n°2010235 – ARIAC and grant n° 8936 – DACS and grant n° 9156 – LogiCourts

References

Ahuja et al (2002) '*A survey of very large-scale neighborhood search techniques*'. Discrete Applied Mathematics, 123(1-3), 75-102, https://doi.org/10.1016/S0166-218X(01)00340-4

Cruijssen, F. *et al.* (2007) '*Joint route planning under varying market conditions*', *International Journal of Physical Distribution & Logistics Management*, 37(4), pp. 287–304. Available at: https://doi.org/10.1108/09600030710752514.

Defryn and Söresen (2018), '*Multi-objective optimisation models for the travelling salesman problem with horizontal cooperation*'. European Journal of Operational Research 267(3), 891–903. Available at: https://doi.org/10.1016/j.ejor.2017.12.028

De Landtsheer *et al.* (2021), '*A Very Large Scale Neighborhood Approach to Pickup and Delivery Problems with Time Windows*', in Proceedings of the 10th International Conference on Operations Research and Enterprise Systems - ICORES, 273-280. Available at: https://www.scitepress.org/PublishedPapers/2021/103159/103159.pdf

De Landtsheer *et al.* (2025), '*Enabling Decentralized Collaboration Among Transporters Through an Optimizing Trading Network for Transport Orders*'. In: Dassisti, M., Madani, K., Panetto, H. (eds) IN4PL 2024. Communications in Computer and Information Science, vol 2373. Springer, Cham, https://doi.org/10.1007/978-3-031-80775-6_8

Eurostat (2026) *Key figures on European transport – 2025 edition*. Eurostat. Available at: https://ec.europa.eu/eurostat/web/products-eurostat-news/w/wdn-20260120-1.

Eurostats (2025) *Road freight transport statistics*. Eurostats. Available at: https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Road_freight_transport_statistics.

Gansterer and Hartl (2018) '*Collaborative vehicle routing: A survey.*' European Journal of Operational Research 268(1), 1–12, https://doi.org/10.1016/j.ejor.2017.10.023

Gansterer and Hartl (2020) '*Shared resources in collaborative vehicle routing*'. Top 28, 1–20 Available at: https://doi.org/10.1007/s11750-020-00541-6

Gansterer *et al.* (2020) '*The value of information in auction-based carrier collaborations*'. International Journal of Production Economics 221, 107485. Available at: https://doi.org/10.1016/j.ijpe.2019.09.006

Glover and Kochenberger (2003) '*Handbook of Metaheuristics*'. Springer Science & Business Media. ISBN: 978-1-4020-7263-5

Los *et al.* (2025) '*Large-scale collaborative vehicle routing*'. Annals of Operations Research, *350*(1), 201-233. Available at: https://doi.org/10.1007/s10479-021-04504-3

Lyu *et al.* (2019) '*A multi-round exchange mechanism for carrier collaboration in less than truckload transportatio*n'. Transportation Research Part E: Logistics and Transportation Review 129, pp 38-59. Available at: https://doi.org/10.1016/j.tre.2019.07.004

Montoya-Torres, J.R *et al* (2016) '*On the impact of collaborative strategies for goods delivery in city logistics*', Production Planning & Control, 27(6), pp. 443–455. Available at: https://doi.org/10.1080/09537287.2016.1147092.

Montreuil, B. (2011) '*Toward a Physical Internet: meeting the global logistics sustainability grand challenge*', Logistics Research, 3(2–3), pp. 71–87. Available at: https://doi.org/10.1007/s12159-011-0045-x.

Nataraj *et al.* (2019) '*Consolidation centers in city logistics: A cooperative approach based on the location routing problem*'. International Journal of Industrial Engineering Computations 10(3), 393–404. Available at: https://hdl.handle.net/10609/99649

Nissan *et al.* (2007) '*Algorithmic Game Theory*'. Cambridge University Press. Available at: https://doi.org/10.1017/CBO9780511800481

Yao *et al.* (2019) '*Assessment of collaboration in city logistics: From the aspects of profit and CO2 emission*'. International Journal of Logistics Research and Applications 22(6), 576–591. Available at: https://doi.org/10.1080/13675567.2019.1594733

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE
PILOTS

Policy-based and Process-Aware Interoperability in the Physical Internet

Philippe Michiels[1], Julián Rojas[2], and Birger Schrevens[1]

1. imec, Leuven, Belgium
2. Ghent University, Ghent, Belgium

Corresponding author: philippe.michiels@imec.be

Abstract: *The Physical Internet (PI) requires coordinated execution of logistics processes across independent actors, but existing interoperability approaches remain data-centric or rely on tightly coupled integrations. This paper proposes a process-aware interoperability model in which coordination is achieved through policy-governed event exchange representing the evolution of shared process state. The approach integrates event-driven interaction with ODRL-based governance and semantic mediation, enabling loosely coupled interoperability across heterogeneous systems. A federated architecture extends data space principles from data sharing to process coordination. The model is validated through a proof-of-concept in container weighing, demonstrating feasibility and reduced integration effort.*

Keywords: *Physical Internet; Federated Interoperability; Logistics Data Platform; Governance; Access and Adoption; Event-Driven Architecture; Semantic Interoperability.*

1 Introduction

The Physical Internet (PI) envisions logistics networks in which independent actors dynamically coordinate transport and handling processes across organizational boundaries (Garbers et al., 2025). In practice, such coordination remains limited by interoperability challenges, as existing approaches rely on bilateral integrations or platform-centric solutions that introduce tight coupling and high integration effort, particularly for SMEs operating with heterogeneous systems (Michiels & Schrevens, 2025). While data space technologies address key challenges related to trust and controlled data sharing, they do not provide a generic mechanism for coordinating the evolution of shared process state (Niemann et al., 2025).

This paper proposes a process-aware interoperability model in which cross-organizational coordination is achieved through policy-governed event exchange representing the evolution of shared process state. The approach integrates event-driven interaction with policy-based governance, supported by semantic mediation and controlled data disclosure mechanisms, enabling loosely coupled interoperability without centralized control. The remainder of this paper is structured as follows. Section 2 discusses related work. Section 3 introduces the conceptual model. Section 4 presents the event-oriented interoperability model. Section 5 describes the policy framework. Section 6 outlines the architecture. Section 7 demonstrates the approach through a proof-of-concept, and Section 8 concludes.

2 Background and Related Work

A substantial body of Physical Internet (PI) research focuses on network design, routing, and large-scale optimization, including data-driven and hyperconnected logistics systems (Garbers et al., 2025). Such work highlights the operational potential of PI principles but typically assumes the availability of timely and sufficiently granular data in highly dynamic logistics ecosystems.

As highlighted in earlier work, these assumptions are difficult to satisfy in practice due to limited visibility across organizational boundaries (Cassan et al., 2023), the absence of sufficient trust between actors (Sun et al., 2024), and the absence of standards for frequent exchange of fine-grained process data to support coordination and planning.

Foundational work in enterprise interoperability has long distinguished technical, semantic, and organizational interoperability concerns in collaborative enterprise networks (Charalabidis et al., 2008). Existing approaches, primarily based on EDI, APIs, or platform-centric integration, nevertheless exhibit structural limitations that hinder broad adoption in heterogeneous logistics ecosystems (Michiels & Schrevens, 2025). More recent data space architectures address key challenges related to federation, trust, and governance, but remain largely focused on data exchange rather than coordinated process execution (Kovach et al., 2026). At their core, such architectures rely on semantic vocabularies, including the W3C Open Digital Rights Language (ODRL), to enable interoperable and policy-governed information exchange (Bader et al., 2020; Slabbinck et al., 2025). Collaborative supply-chain platforms such as C2NET (Andres et al., 2016) have further explored federated coordination and information sharing across enterprise networks, primarily focusing on collaborative planning and optimization rather than process-aware interoperability and policy-governed process execution.

Event-driven architecture has long explored asynchronous coordination through distributed event exchange (Etzion, 2010). Attempts to introduce process-aware interaction models include collaboration protocols and state transition models (Van Gessel & Hofman, 2023), which provide the necessary granularity for representing process evolution but do not fully address complementary concerns such as governance, discoverability, or integration with existing standards. Similarly, semantic approaches to process modelling have been explored in other domains (Harth et al., 2024), though without integrating policy-based enforcement. What remains missing is a unified approach that combines choreography-based process interaction, federated trust and governance, policy-governed admissibility of process evolution, and practical integration with existing enterprise systems and standards.

3 Anatomy of a Logistics Service Process

In this section we introduce a logistics service use case that is sufficiently involved to serve as a reference example and proof-of-concept, yet simple enough to be easily understood.

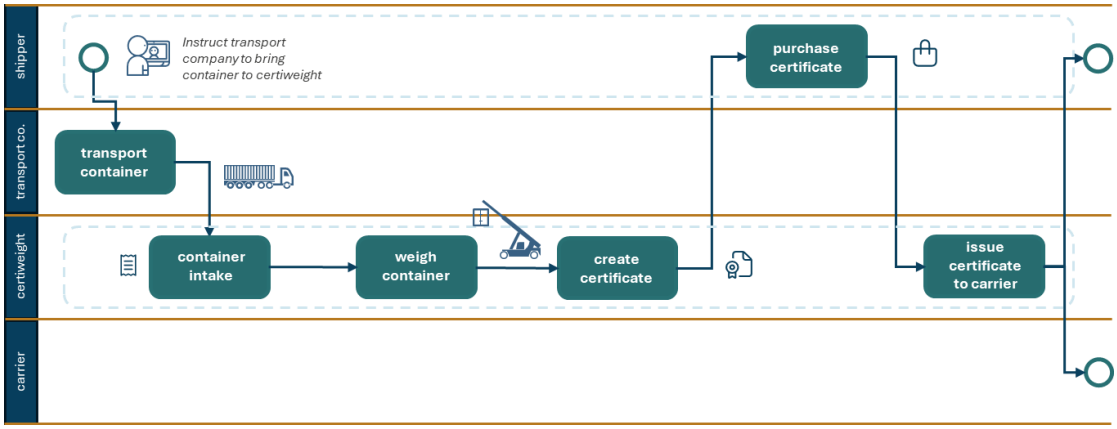


Figure 1: *A BPMN view of the weighing process*

3.1 Certified Container Weighing Process (Use Case Description)

Our use case concerns a certified container weighing process, starting from the arrival of a container at a weighing facility following prior booking and pre-announcement. The weighing service provider registers the container intake and performs the weighing operation, producing

a verified gross mass (VGM). Based on this result, a weighing certificate is created, purchased by the customer, and subsequently issued. The core process reflects a bilateral interaction between a customer (shipper) and a service provider (weighing company). Additional actors may participate either through delegated authority or as notification parties. In this example, the weighing provider is authorized to transmit the verified weight directly to the ocean carrier on behalf of the shipper. The weighing process depicted in **Figure 1** is modeled as a sequence of state transitions driven by discrete activities, which are later mapped to events and governed through policy constraints.

3.2 BPMN as a conceptual reference

BPMN is used as a design-time language to specify roles, activities, transitions, and data objects in inter-organizational service processes. Execution is instead realized through choreography-based, policy-governed event interactions, avoiding centralized orchestration and enabling loosely coupled coordination across autonomous participants. Since BPMN is highly expressive, this work restricts the model to roles, activities, transitions, and data objects, while the definition of a broader executable subset is left for future work.

Actors participate through roles defined in the service definition and bound to participants at service-instance creation. The model distinguishes primary actors responsible for execution, delegated secondary actors, and notification actors with read-only visibility.

Process activities are represented as events associated with a service instance, an authorized actor role, and, where applicable, a state transition. Rather than enforcing execution through a centralized workflow engine, coordination emerges from the exchange of events evaluated against shared service definitions and policies.

Data (objects) exchanged with events is represented through semantically defined parameters and an optional payload containing richer business information in domain-specific formats. A semantic layer maps payload structures to parameters, enabling interoperability while preserving compatibility with existing standards.

4 Event-Oriented Process Interoperability Model

This section defines a generic event-oriented interoperability model for coordinating distributed process execution across organizational boundaries. Inspired by the Eclipse Data Space Protocol (Steinbuss et al., 2025) and event-driven architectures(Etzion, 2010), coordination proceeds through policy-governed exchange of typed events associated with shared service definitions and process state. Unlike traditional event-processing approaches, events represent admissible process actions evaluated against shared policies and local process state.

4.1 Event API Specification

To operationalize the event-oriented interoperability model, this section defines a minimal set of API endpoints that enable participants to: (i) submit events, (ii) retrieve the process state, (iii) access the event history. The API combines a resource-oriented design with asynchronous event propagation, in line with data space protocol interactions. **Table 1** summarizes the different API endpoints and their interaction patterns.

Table 1: *A simple event-driven API specification.*

Endpoint	Method	Description	Body	Response	Key semantics
`/events`	**`POST`**	Submit a new event for a service instance	Event object	202 Accepted	Validates event against service definition and policies; may trigger state transition; propagates to authorized participants

/events/{id}	**GET**	Retrieve a specific event	(none)	Event object	Returns full or filtered event details depending on policies
/serviceInstance/{id}	**GET**	Retrieve current state of a service instance	(none)	Last Event Object	Returns latest locally derived state; response may be filtered by policy
/serviceInstance/{id}/ events	**GET**	Retrieve event history for a service instance	(none)	Array of Event objects	Provides audit trail of process execution; events may be filtered or redacted based on access rights

4.2 Generic Event Model

In the proposed model, an event is a structured message representing a process-relevant activity within a service instance. Events refer to a service instance. They are interpreted against a governing service definition specifying admissible activities, actor roles, and state transitions, and may result in an update of shared process state if accepted under the applicable policies. To support interoperability across heterogeneous systems, events follow a canonical structure that separates control information from data content, while remaining expressive enough to capture both lightweight signals and richer business information. Although inspired by ODRL evaluation concepts and the notion of a State of the World (Bonatti et al., 2025; Slabbinck et al., 2025), events are not treated as self-contained policy evaluation requests. Instead, upon receipt, an event is interpreted in the context of the receiver's local state and the governing service definition, from which a policy evaluation request can be derived.

As shown in **Table 2** and **Listing 1**, each event includes identifiers for the service instance and service definition, the event type, the actor role, and a timestamp. Where applicable, it may also refer to a state transition, allowing it to signal progression in the process lifecycle.

__Table 2__: Event model description.

Field	Semantic definition	Required	Description
serviceInstanceId	odrl:permission / odrl:target	Yes	Identifier of the service instance, mapped to ODRL target during evaluation
serviceDefinitionId	pilots:service	Yes	Reference to governing service definition
eventType	rdf:type	Yes	Type of event (activity identifier)
actor	pilots:actor/rdf:type	Yes	Role under which the event is emitted
timestamp	dct:issued	Yes	Time of event creation
stateTransition	odrl:permission/odrl:action + sotw:context	No	Transition triggered by the event, used to derive action and constraints
payload	pilots:payload	No	Optional business information (e.g. EDI, JSON)
parameters	pilots:parameters	No	A dictionary containing data points (mapped from the payload) that can be indicated as required by the policy

Through this structure, events act as the operational mechanism for coordinating distributed process execution. A sequence of validated events forms a consistent and auditable representation of process progression, effectively realizing a distributed state machine in which state emerges from the shared interpretation of event exchanges.

```
{
  "@context": {
    "odrl": "http://www.w3.org/ns/odrl/2/",
```

```
    "pilots": "https://pilots-project.be/ns#",
    "dct": "http://purl.org/dc/terms/",
    "xsd": "http://www.w3.org/2001/XMLSchema#",
    "ex": "http://example.org/"
  },
  "@type": ["pilots:ServiceUpdate", "pilots:completeWeighing" ],
  "@id": "urn:uuid:9a3f3b0e-8f1a-4f5c-b3b4-92bcb8d7c001",
  "pilots:serviceDefinitionId": "pilots:WeighingServiceDescription",
  "pilots:serviceInstanceId": "urn:uuid:12435aaa-1234-1234-1234-123412341234",
  "dct:issued": {
    "@value": "2026-04-02T10:15:00Z",
    "@type": "xsd:dateTime"
  },
  "pilots:actor": {
    "@id": "ex:CertiWeight",
    "@type": "pilots:serviceProvider"
  },
  "pilots:involves": [
    {
      "@id": "ex:VanMoerLogistics",
      "@type": "pilots:serviceUser"
    }
  ],
  "pilots:previousState": "pilots:containerReceived",
  "pilots:newState": "pilots:weighingCompleted",
  "pilots:parameters": {
    "pilots:containerReference": "MSKU1234567",
    "pilots:weighingTimestamp": {
      "@value": "2026-04-02T10:14:30Z",
      "@type": "xsd:dateTime"
    },
    "pilots:measuredGrossMass": {
      "@value": "28750",
      "@type": "xsd:decimal"
    },
    "pilots:massUnit": "KGM",
    "pilots:weighingLocation": "Certiweight Antwerp LO"
  }
}
```


***Listing 1**: An example event object.*

4.3 Event Data and Payload

In addition to control-flow information, events carry the data required for their interpretation. The model distinguishes between *event parameter*s, which provide a lightweight, semantically defined representation of process-relevant data, and an optional *payload*, which may contain richer business documents conforming to existing standards. Event parameters define the minimal data required to interpret an event under the governing service definition and may be constrained by policy. The payload, in contrast, is optional and may follow domain-specific formats such as EDI or JSON. To bridge both representations, the model introduces a *semantic profile* that defines the mapping between payload structures and event parameters. This mapping may be expressed through a semantic overlay, enabling connectors to extract relevant parameters and perform transformations between formats. Content negotiation mechanisms further allow participants to exchange events in different representations while preserving a shared semantic interpretation, enabling interoperability across systems with varying levels of digital maturity.

5 Policy-Governed Process Execution

While Section 4 defines the operational event model, this section specifies the policy framework governing admissible interactions and data visibility. The proposed approach relies on ODRL-based policies associated with the service definition to control which actions may be

performed, under which conditions, and which information may be disclosed to participating actors.

5.1 ODRL Policy Model

Our approach extends ODRL beyond traditional data access control to govern admissible process interactions and state transitions in distributed service choreography. Under ODRL's closed-world assumption, only explicitly permitted actions are allowed. In the proposed model, actions correspond to event types, while constraints capture process-relevant conditions such as actor roles, current state, admissible transitions, and required parameters. Domain-specific ODRL extensions further support structural validation of events, enabling both behavioral and data-level requirements to be enforced through a unified policy framework.

```
{
  "@type": "odrl:Permission",
  "odrl:action": { "@id": "pilots:completeWeighing" },
  "odrl:constraint": [ {
      "odrl:leftOperand": "pilots:actorRole",
      "odrl:operator": "odrl:eq",
      "odrl:rightOperand": "pilots:serviceProvider"
    },{
      "odrl:leftOperand": "pilots:currentState",
      "odrl:operator": "odrl:eq",
      "odrl:rightOperand": "pilots:containerReceived"
    },{
      "odrl:leftOperand": "pilots:nextState",
      "odrl:operator": "odrl:eq",
      "odrl:rightOperand": "pilots:weighingCompleted"
    },{
      "odrl:leftOperand": "pilots:parameters",
      "odrl:operator": "pilots:mustContain",
      "odrl:rightOperand": [
        "pilots:containerReference",
        "pilots:weighingTimestamp",
        "pilots:measuredGrossMass",
        "pilots:massUnit"
    } ]
}
```

***Listing 2**: Example Service Definition Rule.*

Listing 2 illustrates a permission rule defined at the service-definition level. The rule permits action **`completeWeighing`** for the **`serviceProvider`** role, constrained by the current and next process state. It further requires the presence of specific event parameters through a containment constraint. This shows how ODRL is extended to encode both process semantics and structural validation requirements. It may be interesting to further align with other ODRL based research (Manab et al., 2026).

5.2 Policy Engine and Process Constraints

The enforcement of ODRL policies is delegated to a policy engine that evaluates each interaction against the applicable policy set. Policy enforcement is treated as a deterministic evaluation process that produces a structured outcome indicating whether a given action complies with the governing policies (Slabbinck et al., 2025). This outcome can be understood as a compliance decision, accepting, rejecting, or conditionally allowing the event based on the satisfaction of constraints and obligations. The receiving participant derives an evaluation context by combining the event with its local state of the world and the applicable service definition. Upon receipt, the sender identity and its binding to the claimed actor role within the service instance are ascertained. The policy engine checks if the requested action is permitted under the current process state and satisfies all associated constraints, including admissible state transitions and required parameters.

5.3 Selective Disclosure

Beyond controlling actions, policies may also govern the visibility of event data across participants. In distributed logistics processes, different actors may require access to different subsets of information, motivating mechanisms for selective disclosure. In this model, such disclosure can be expressed as policy-driven transformations on event content. However, the enforcement and implementation of these transformations are not addressed in this work and are left for future research.

6 Architecture for Federated Process Interoperability

6.1 Overview

The proposed architecture (Michiels & Schrevens, 2025) extends data space principles with process-aware interoperability, enabling actors to discover, instantiate, and execute logistics services through event-driven interaction. It combines federated trust and data exchange with a process coordination layer, allowing heterogeneous systems to participate in shared processes without tight coupling, following a choreography model rather than centralized orchestration. Participants independently validate and interpret events against shared service definitions and local process state, while global coordination emerges from policy-governed event exchange.

6.2 From Data Space to Process Space

Data spaces enable sovereign data exchange through a lifecycle of offering publication, policy-based negotiation, and controlled data transfer. While effective for data sharing, this model remains data-centric and does not directly support coordinated process execution.

We introduce logistics service offerings alongside traditional data offerings. These expose machine-readable *service definitions*, roles, interaction points, and standard terms, expressed using ODRL-based policies. In contrast to data offerings, services typically require context-specific pricing and constraints, resolved through a quotation exchange prior to execution. This results in a layered interaction model consisting of service offering, quotation, service instance, and event exchange, where the latter governs process execution through state transitions. In this model, the event API serves as the operational interface, extending the data space into a process space in which actors coordinate through policy-governed event interactions rather than isolated data transfers.

6.3 Architectural Components

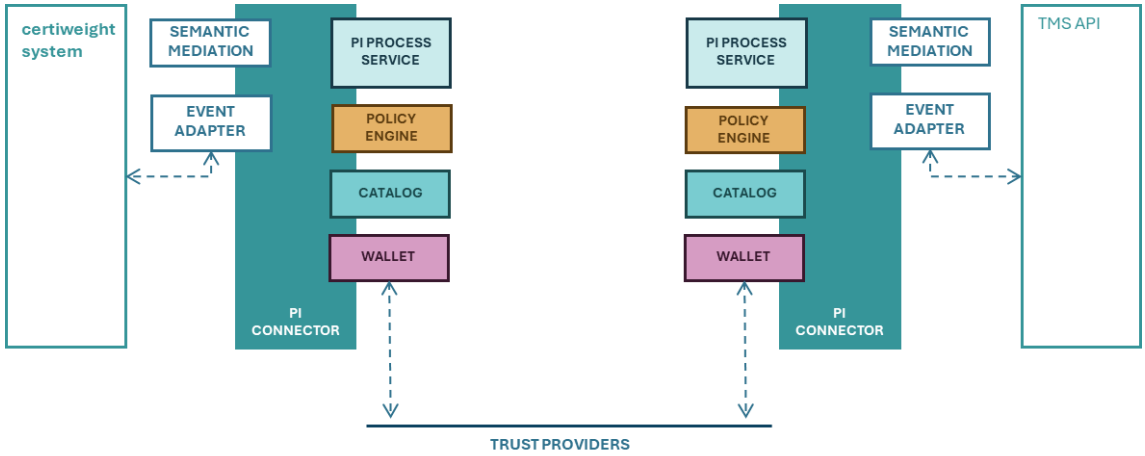


Figure 2: *A high-level overview of the process sharing architecture.*

The architecture in **Figure 2** comprises several interacting components. The *data space connector* acts as the trusted deployment unit, enabling secure, policy-governed interaction and exposing service offerings through inbound and outbound APIs. Identity and trust are handled by *wallets* and *trust providers*, which manage verifiable credentials, authentication, authorization, and revocation mechanisms. A *service catalog* supports federated discovery of service offerings, including process definitions (*service definitions*) with their policies. Process execution is managed by the *process sharing service*, which handles quotations, maintains service instances, and tracks process state through event exchange. A *policy engine* enforces ODRL-based rules locally, governing permissions, event admissibility, visibility, and state transitions. The policy engine enables distributed workflow choreography through local evaluation of admissible process interactions. *Semantic mediation* enables interoperability across heterogeneous systems through vocabulary alignment and content negotiation, while *event adapter middleware* translates between local systems and the common event model. Finally, *backend systems* (e.g. TMS, ERP) execute business logic and interact with the process through adapter-mediated events.

6.4 Interaction Flow

The interaction between components follows a structured lifecycle:

1. A service provider publishes a logistics *service offering* in the service catalog, including the *service definition* that defines the policies.
2. A consumer discovers the offering and submits a *quotation request* with context-specific parameters.
3. The provider returns a *quotation*, refining the offering into a concrete commercial and operational agreement.
4. Upon acceptance, a *service instance* is created and bound to the relevant actors, roles, and policies.
5. The process is executed through event exchange via the event-driven API. Upon receipt, each participant validates incoming events by verifying their provenance and evaluating them against the governing service definition and its local state of the world using a policy engine.
6. Backend systems interact through adapters that translate local actions into events and propagate received events into local workflows, while policy evaluation remains local to each participant. Accepted events are incorporated into the local process state and may be propagated to other authorized participants.
7. The model assumes eventual consistency of process state across authorized participants through propagation of validated events.

7 Proof-of-Concept: Certified Container Weighing Process

The certified weighing process introduced in Section 3 is expressed here using the proposed event-oriented model. Rather than restating the process flow, **Table 3** summarizes its execution as a sequence of events aligned with actor roles, state transitions, and policy constraints. This representation illustrates how process coordination, access control, and data requirements are jointly enforced through the service definition.

Table 3: *The weighing process captured as a series of events, actors, data and constraints.*

Event	Role	From → To State	Required Parameters	Policy Constraints
`requestWeighing`	serviceUser	`init → weighingRequested`	containerReference	role=user
`registerIntake`	serviceProvider	`weighingRequested → containerReceived`	containerReference, intakeTimestamp	role=provider

Event	Role	From → To State	Required Parameters	Policy Constraints
`completeWeighing`	serviceProvider	`containerReceived → weighingCompleted`	measuredGrossMass, massUnit, weighingTimestamp	role=provider mustContain
`createCertificate`	serviceProvider	`weighingCompleted → certificateCreated`	certificateId	role=provider
`purchaseCertificate`	serviceUser	`certificateCreated → certificatePurchased`	paymentReference	role=user
`issueCertificate`	serviceProvider	`certificatePurchased → certificateIssued`	certificateId	role=provider notify(carrier)

Unlike traditional API integrations, participants are not required to implement pairwise workflow logic or maintain centralized process synchronization. Coordination instead emerges from the exchange of semantically defined events evaluated locally against shared policies and process constraints. This allows heterogeneous participants to independently evolve internal systems while maintaining interoperable process execution.

While the example demonstrates how process behavior can be expressed through ODRL-based policies, requiring process designers to define such policies directly is not practical. Instead, we envision a model-driven approach in which annotated BPMN process definitions are compiled into corresponding service definitions and policy rules. Given BPMN's expressive power, future work will focus on defining a constrained subset that can be consistently mapped to event types, state transitions, and policy constraints.

The use of RDF as an underlying representation further supports extensibility of the model. By relying on shared vocabularies and ontologies, process definitions can be incrementally extended with additional activities, parameters, or constraints while preserving semantic interoperability. This is particularly relevant in logistics environments, where canonical processes often admit variations across organizations or operational contexts. RDF-based modelling enables such variations to be expressed without breaking compatibility with existing implementations, supporting gradual evolution of shared process definitions.

8 Conclusion and Future Work

This paper demonstrates how process-aware, event-driven interoperability can enable coordination across independent actors in Physical Internet environments characterized by decentralization, heterogeneous systems, dynamic collaboration, and limited cross-organizational trust. By representing process execution as a sequence of policy-governed event exchanges, the proposed approach allows heterogeneous systems to interoperate without tight coupling or centralized control.

Compared to traditional EDI and API-based integrations, the connector-based approach reduces integration effort through reuse of a common event model and policy framework, shifting from pairwise integrations to a federated model. Initial results indicate a significant reduction in integration effort from weeks to days, and a structural decrease in complexity. Further implementations beyond the proof-of-concept will allow us to quantify these advancements.

The approach extends federated data-sharing principles toward process-aware interoperability, enabling distributed process choreography through policy-governed event exchange rather than centralized workflow control. Validation in larger-scale deployments is required to assess scalability and generalizability. Future work will focus on automated derivation of service definitions from process models, event analytics, and coordinating processes across multiple services, as well as validation in larger-scale Physical Internet pilots.

9 References

Andres, B., Sanchis, R., & Poler, R. (2016). A Cloud Platform to support Collaboration in Supply Networks. *International Journal of Production Management and Engineering*, *4*(1), 5. https://doi.org/10.4995/ijpme.2016.4418

Bader, S., Pullmann, J., Mader, C., Tramp, S., Quix, C., Müller, A. W., Akyürek, H., Böckmann, M., Imbusch, B. T., Lipp, J., Geisler, S., & Lange, C. (2020). *The International Data Spaces Information Model – An Ontology for Sovereign Exchange of Digital Content* (pp. 176–192). https://doi.org/10.1007/978-3-030-62466-8_12

Bonatti, P. A., Fornara, N., & Harth, A. (2025). *Towards a Formal Semantics of the Open Digital Rights Language (ODRL 2.2)*. Https://Www.W3.Org/1999/02/22-Rdf-Syntax-Ns.

Cassan, C., Michiels, P., Sun, S., Lemos, V., Van Bever, D., Cant, A., Fink, J., & Macharis, C. (2023). Data Sharing in the Physical Internet: A Capability-Based Approach for Trustless Logistic Networks. *Proceedings of the 9th International Physical Internet Conference*, 154–165. https://doi.org/10.35090/gatech/7986

Charalabidis, Y., Gionis, G., Hermann, K. M., & Martinez, C. (2008). *Enterprise Interoperability Research Roadmap Update (Version 5.0)*.

Etzion, O. (2010). Event Processing-past, present and future. *Proceedings 36th International Conference on Very Large Data Bases, September 13-17, 2010, Singapore*. http://www.dagstuhl.de/de/programm/kalender/semhp/?semnr=10

Garbers, A., Muthukrishnan, P., & Montreuil, B. (2025). *Pragmatic Hyperconnected Freight Systems Design: A Data-Driven Case Study in the US Southeast Megaregion 2. Coca-Cola Material Handling & Distribution Chair*. https://doi.org/10.35090/gatech/13240

Harth, A., Käfer, T., Rula, A., Calbimonte, J.-P., Kamburjan, E., & Giese, M. (2024). Towards Representing Processes and Reasoning with Process Descriptions on the Web. *Transactions on Graph Data and Knowledge*, *2*(1). https://doi.org/10.4230/TGDK.2.1.1

Kovach, A., Montalvillo, L., Lanza, J., Sotres, P., & Urbieta, A. (2026). Understanding data spaces: A Systematic Mapping Study of foundations, technical building blocks, and sectoral adoption. *Computer Science Review*, *59*, 100819. https://doi.org/10.1016/j.cosrev.2025.100819

Manab, M. A., Quaranta, M., Amantea, I. A., Leyva-Sánchez, S., & Rodríguez-Doncel, V. (2026). *Extracting ODRL Policies from Business Process Models: A Graph Traversal Approach to Compliance-by-Extraction*. https://github.com/manabcodes/bpmn2odrlandhttps://bpmn.linkeddata.es/.

Michiels, P., & Schrevens, B. (2025). A Technical Blueprint for the Physical Internet. *Proceedings of the 11th International Physical Internet Conference*, 246–255. https://doi.org/10.35090/gatech/13258

Niemann, G., Franklin, R., & Spanuth, M. (2025). *Addressing the Reachability Problem in the Physical Internet: A PI-Link Protocol Inspired by the Internet*. https://doi.org/10.35090/gatech/13227

Slabbinck, W., Rojas Meléndez, J., Esteves, B., Colpaert, P., & Verborgh, R. (2025). *Interoperable Interpretation and Evaluation of ODRL Policies* (pp. 192–209). https://doi.org/10.1007/978-3-031-94578-6_11

Steinbuss, S., Koen, P., Kollenstart, M., Marino, J., Pampus, J., Turkmayali, A., & Weiß, A. (2025, November 12). *Dataspace Protocol 2025-1*. Https://Eclipse-Dataspace-Protocol-Base.Github.Io/DataspaceProtocol/2025-1-Err1/.

Sun, S., Michiels, P., Macharis, C., Cant, A., & Mommens, K. (2024). Unlocking the Potential of the Physical Internet: a Trust-enabling Decentralized Process Sharing Connector. *Proceedings of the 10th International Physical Internet Conference*, 6–15. https://doi.org/10.35090/gatech/11062

Van Gessel, T., & Hofman, W. (2023). An Interaction Pattern Ontology for Data Sharing about Logistics Activities. *Joint Workshop Proceedings of the 5th International Workshop on Semantic Data Space For Transport (Sem4Tra)*. https://ceur-ws.org/Vol-3510/paper\sem4tra\2.pdf

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE
UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

Roadmap of Logistics Data Platforms for the Physical Internet: A DTLF-Aligned Assessment of Open Standard Data Exchange, Federated Data Spaces, and Governance for Adoption

Paulo Cantillano-Lizana[1,2], Gloria Giammei[1,2,*] and Matteo Magnani[1,2]
University of Modena and Reggio Emilia, Reggio Emilia, Italy
Interuniversity Consortium for Optimization and Operation Research (ICOOR), Reggio Emilia, Italy
Corresponding author: gloria.giammei@unimore.it

Abstract: *This paper proposes an evaluation of the fitness of the Physical Internet (PI) roadmap in facilitating the development of a Logistics/Commercial Data Platform. The evaluation is based on two main innovations from Hologistics Eu-project: (i) DTLF-aligned open standard data exchange for digital freight information, and (ii) data space governance approaches exemplified by EMDS/LaaS. A roadmap suitability framework is proposed and based on four qualitative dimensions: capability coverage, interoperability readiness, governance adequacy, and adoption viability. The results suggest that there is considerable alignment in terms of functions, particularly for interoperability and data exchange, focusing on compliance. However, in terms of ecosystem, the results are less promising, given the absence of universal APIs, APIs fragmentation, and cybersecurity risks. In terms of governance, it has been found that governance-by-design, as based on IDS/IDSA, is consistent with data sovereignty and trust, although it is constrained by factors such as those of socio-technical change, including organizational change, liability, and data sensitivity. Overall, it has been found that adoption is the main constraint on the development of a Logistics/Commercial Data Platform. This paper concludes by suggesting a checklist for deployment based on PI, as well as a potential research agenda focusing on interoperability, adaptive governance, and security-by-design, as a means of moving from standards to deployment.*

Keywords: *Physical Internet; logistics data platform; DTLF; eFTI; federated data spaces; data sovereignty; interoperability; governance; API fragmentation; cybersecurity.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☐ *PI Networks,* ☐ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☒ *Logistics/Commercial Data Platform,* ☒ *Access and Adoption,* ☒ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☐ *In-Person presentation*

1 Introduction

The vision of the Physical Internet (PI) is to create an open, modular, and interoperable network of logistics services using standardised physical and digital interfaces and the Logistics/Commercial Data Platform is a key enabler of this vision. The European policy and standardisation landscape have recently come together on digital freight information exchange via the Digital Transport and Logistics Forum (DTLF) and the (electronic Freight Transport Information) eFTI Regulation. In addition, there are complementary horizontal measures such

as the Data Governance Act and the Data Act that define the framework for secure data sharing, access, and usage across various sectors (Ruohonen & Mickelsson, 2023).

The adoption of digital freight information exchange is still delayed by issues such as workforce digital skill gaps, system interoperability problems, and varying company capacities that complicate uniform implementation of electronic freight transport information regulations (Dasaklis, et al., 2024) but also encounter resistance due to fears of increased market transparency and interface standardization, which affect adoption differently among service providers and users (Bossong, Reinhardt, & Elbert, 2024). Additionally, unwillingness to data sharing and integration challenges limits the full potential of digital freight forwarders despite their advanced functionalities like AI analytics and blockchain connectivity (Zeng, Chen, Xu, Chan, & Xu, 2025). The issues facing digital freight information exchange are more of a governance and adoption nature.

This paper has three major contributions, first, by providing a formalised framework for the assessment of the fitness of the PI roadmap, second by synthesising evidence on DTLF-aligned data exchange using open standards and federated governance models such as European Mobility Data Space (EMDS)/ Logistics as a Service (LaaS), and third, by derivation targeted research and deployment agenda to close remaining gaps. We specifically refer to the European PI implementation context (DTLF/eFTI), while maintaining compatibility with the broader Physical Internet roadmap families developed in Europe and Japan, which structure PI development around roadmap areas such as Logistics/Commercial Data Platforms, Governance, and Access & Adoption. Our operationalisation and evidence base are intentionally anchored in the European implementation context, DTLF and eFTI, together with EU data governance instruments, while the proposed fitness framework is designed to be transferable to other PI roadmap variants by re-mapping the roadmap areas and evidence sources

2 Background and related work

The PI roadmap areas of interest can be defined as: (i) Logistics/Commercial Data Platform (shared semantic foundations, interoperable interfaces, discoverability, compliance), (ii) Governance (trust frameworks, sovereignty, accountability, access control, liability), and (iii) Access & Adoption (incentives, onboarding, organisational readiness, standard uptake and lifecycle management). These are interdependent, that is, platform interoperability without governance can stall due to lack of trust or legal clarity, and governance without interfaces can impede adoption. DTLF-Compliant Open Standard Data Exchange is based upon semantic models, APIs, federated models, and EU standardisation, and key enable technologies include UN/CEFACT and GS1 standardised data modelling, open protocols, and interoperability tooling, with examples such as the ISO standard Core Components Technical Specification (CCTS) meta-model for data exchange standard lifecycle management and tooling, Score for life-cycle management and integration of enterprise systems (Ivezic, et al., 2022). For integration with cyber-physical systems and digital twin systems, Open Platform Communications Unified Architecture (OPC UA) can be used for the transformation of digital twin definitions into communication protocols (Cavalieri & Gambadoro, 2023).

DTLF-Compliant Open Standard Data Exchange involves real-time sharing of ETAs, Intermodal Transport Unit (ITU) status updates, and single window concepts for B2B and B2A data exchange. Examples include the federated architecture and port single window implementations of the FENIX project, which have demonstrated improved electronic data exchange and lower administrative burdens (Apruzzese & Bresseleers, 2023). Key elements include data sovereignty, trust, and multiple-layered data concepts such as data stewardship,

shared ownership, and self-sovereign identity, with the International Data Spaces approach highlighted as governance-by-design concepts, as opposed to economic coordination concepts (Von Scherenberg, Hellmeier, & Otto, 2024); (Otto & Jarke, 2019); (Janssen, Brous, Estevez, Barbosa, & Janowski, 2020). The key adoption barriers are identified as socio-technical issues such as trust, company resistance, technical limitations, and the sharing of sensitive data, especially where usage control and access governance are difficult to apply across companies (Hellmeier, Pampus, Qarawlus, & Howar, 2023). The key identified standards and frameworks are eFTI (2020/1056), IDSA/IDS governance guidelines, and the European Interoperability Framework (EIF), with the need for evaluation of the impact of such frameworks on trust, compliance, and adoption (Shabani, 2022); (Korca & Costa, 2021); (Prifti & Fosch-Villaronga, 2024); (Gyard, et al., 2025). For DTLF-aligned exchange, the key limitations are identified as differential adoption, the absence of universal APIs, security risks, and the cost of API compatibility.

The key elements point to the existence of systemic evidence of the impact of API evolution and fragmentation on integration costs (Milčiuvienė, Amilevičius, & Raskinis, 2024). It also points to the significant cybersecurity risk posed using APIs, which has driven the need for security and data governance standards such as ISO/IEC 25012:2008, ISO 20000, and ISO 27001, as identified in the standards mapping.

3 Method

This paper applies a qualitative roadmap-fitness assessment anchored in the European PI implementation context (DTLF/eFTI and EU data governance instruments) and mapped to PI roadmap areas that are also present in the Europe/Japan PI roadmap families: Logistics/Commercial Data Platform, Governance, and Access and Adoption

The unit of analysis is the “approach” (DTLF-aligned exchange; EMDS/LaaS governance) mapped against the PI roadmap areas and evaluated using a common rubric.

Roadmap fitness is defined as the degree to which an approach provides the capabilities and conditions required to implement a PI roadmap area in real-world, multi-actor logistics settings. Fitness is assessed across four dimensions:

1. Capability coverage: extent to which required functions exist (e.g., event exchange, semantic modelling, access control, auditability).
2. Interoperability readiness: maturity of semantic, syntactic, and process interoperability across heterogeneous actors and systems.
3. Governance adequacy: presence and maturity of mechanisms for data sovereignty, trust, accountability, usage control, and liability handling.
4. Adoption viability: plausibility of scaled uptake considering integration burden, incentives, organisational readiness, regulatory fit, and security risks.

For each PI roadmap area, each of the two approaches was assessed (DTLF-aligned exchange and EMDS/LaaS governance) against the four fitness dimensions using a three-level scale: Low (L), Medium (M), High (H). Ratings are assigned through indicator-based judgement grounded in extracted evidence (features, benefits, barriers, and implementation examples). To reduce subjectivity, the rubric follows threshold logic:

- High: multiple indicators are satisfied and at least one implementation or standardised mechanism demonstrates feasibility beyond isolated bilateral integration; barriers are present but not dominant.

- Medium: partial indicator satisfaction; feasibility demonstrated in limited contexts or pilots; key gaps remain (e.g., fragmentation, unclear liability, missing conformance artefacts) that materially constrain scaling.
- Low: weak indicator satisfaction; evidence is mainly conceptual or isolated; barriers dominate and scaling appears infeasible without major additional mechanisms.

Indicators used per dimension are:

1. Capability coverage:
    - High: supports end-to-end exchange primitives (events/status/ETA), standardised data models, and operational workflows (e.g., single-window patterns).
    - Medium: supports core exchange functions but lacks consistent workflows, lifecycle tooling, or critical platform functions (e.g., conformance testing, monitoring).
    - Low: lacks core functional primitives or relies on bespoke/closed mechanisms.
2. Interoperability readiness:
    - High: explicit semantic alignment (shared standards/ontologies), interoperable APIs, and evidence of cross-system interoperability with conformance guidance
    - Medium: semantic standards exist but implementations vary; API fragmentation/versioning issues are significant; interoperability depends on bilateral integration effort
    - Low: no stable semantic baseline; interoperability is ad hoc.
3. Governance adequacy:
    - High: enforceable sovereignty/trust mechanisms (roles, access control, audit, usage control) plus clear accountability and dispute/liability handling logic.
    - Medium: governance concepts are present but enforcement/cross-organisation operability is incomplete; liability/accountability remains ambiguous.
    - Low: governance relies mainly on informal trust or lacks enforceable controls.
4. Adoption viability:
    - High: manageable integration and maintenance burden, clear incentives/value proposition, feasible compliance, and security-by-design practices; onboarding is plausible for SMEs and large actors.
    - Medium: value exists but integration/maintenance burdens (API evolution, semantic complexity), organisational resistance, or security concerns limit scaling; incentives/onboarding require strengthening.
    - Low: adoption requires high bespoke effort, uncertain benefits, major trust/security blockers, or unclear regulatory feasibility.

The assessment proceeds in three steps.

- Step 1: Evidence extraction. For each approach, statements on enabling technologies, key features, benefits, barriers/challenges, and standards/regulatory anchors are extracted. Implementation examples (e.g., pilots, federated architectures, single-window deployments) are captured as “demonstration evidence.”
- Step 2: Mapping. Each extracted item is mapped to one or more PI roadmap areas (Logistics/Commercial Data Platform; Governance; Access and Adoption) and one or more fitness dimensions (A–D).
- Step 3: Scoring and justification. For each roadmap area × approach, dimension scores (L/M/H) are assigned using the rubric previously. Mixed evidence defaults to the lower score unless demonstration evidence supports the higher level.

Table 1 summarises the resulting qualitative ratings for each roadmap area and fitness dimension, reported separately for the two approaches (DTLF-aligned exchange vs EMDS/LaaS governance models). The table is intended as a compact artefact linking the evidence base to the paper's comparative claims; detailed narrative justifications are provided in the Results section.

Table 1: PI roadmap fitness matrix (qualitative ratings)

PI roadmap area	Capability coverage	Interoperability readiness	Governance adequacy	Adoption viability
Logistics/Commercial Data Platform	H / M	M / M	M / H	M / L–M
Governance	M / H	M / M	M / H	L–M / L–M
Access and Adoption	M / M	M / L–M	L–M / M	L–M / L

This framework increases transparency versus purely narrative assessments by making evaluation dimensions explicit and applying a consistent rubric. Nonetheless, scores remain qualitative and reflect the strength and coverage of available evidence rather than measured performance. Future work can extend the approach with quantitative indicators (e.g., integration cost proxies, conformance test pass rates, onboarding time, adoption metrics, security incident rates) and validate ratings through expert elicitation or comparative case studies across corridors and stakeholder types.

4 Results

This section reports the roadmap-fitness results for two building blocks, DTLF-aligned open standard data exchange ("DTLF") and federated data space governance models (EMDS/LaaS, "EMDS"), across the three PI roadmap areas (Logistics/Commercial Data Platform; Governance; Access and Adoption) and the four dimensions defined in Chapter 3. Scores follow the logic of Table 1 and are justified through the underlying evidence on features, standards anchor, barriers, and implementation examples.

DTLF shows high capability coverage, supported by standardised semantic models, open interfaces, event/status exchange (e.g., ETA/ITU), and single-window patterns. EMDS is medium: it enables sharing under sovereignty and trust constraints but does not provide domain-specific logistics semantics or exchange primitives by itself. Interoperability readiness is medium for both: DTLF is constrained by heterogeneous implementations and API evolution, while EMDS depends on conformance artefacts and shared domain models to reduce semantic and process variability. Governance adequacy is stronger for EMDS, which directly targets sovereignty, access control, auditing, and usage control, whereas DTLF contributes more indirectly via compliance-aligned exchange. Adoption viability is medium for DTLF (regulatory momentum but fragmentation/security burdens) and low–medium for EMDS (governance overhead, onboarding costs, and socio-technical barriers). A representative implementation example is the FENIX project, which operationalized DTLF-aligned federated architectures and port single-window concepts for freight information exchange. The project demonstrated the feasibility of interoperable shipment-event, while also highlighting persistent challenges related to API heterogeneity and onboarding complexity.

EMDS has high capability and governance adequacy because it is explicitly a governance model; DTLF is medium as an exchange approach that requires additional layers for enforceable sovereignty, liability allocation, and dispute handling. Interoperability remains medium for both, since policy/process interoperability and consistent enforcement are still

challenging. Adoption viability is low–medium for both due to implementation heterogeneity (DTLF) and trust/sensitivity/overhead concerns (EMDS).

Both approaches have medium capability coverage, but neither is adoption-complete: DTLF reduces administrative friction, while EMDS enables controlled access. Interoperability is medium for DTLF and low–medium for EMDS because practical, low-friction onboarding depends on mature registries, connectors, and conformance tooling. Governance adequacy is low–medium for DTLF and medium for EMDS. Adoption viability is the limiting factor overall (DTLF low–medium; EMDS low), reflecting the need for a "translation layer" that turns standards and governance into scalable practice through interoperability assurance, usable/adaptive governance, incentives and onboarding support, and security-by-design for API ecosystems.

5 Discussion

The fitness matrix (Table 1) suggests a consistent and practically important pattern. DTLF-aligned open standard data exchange scores highest where the PI Logistics/Commercial Data Platform requires concrete operational primitives: shared semantics, event exchange, and compliance-oriented workflows that can be instantiated through APIs and "single-window" interaction patterns. In contrast, EMDS/LaaS governance models score highest for the Governance roadmap area because they focus explicitly on data sovereignty, trust, role-based access, and auditable sharing. This division of strengths supports the core interpretation that the two approaches are complementary layers in a PI-aligned digital stack rather than competing alternatives. However, this complementarity also introduces trade-offs. Federated governance models (EMDS) strengthen data sovereignty, trust, and usage control, but may increase onboarding effort and operational complexity across heterogeneous actors. Conversely, more centralized approaches can reduce integration friction and simplify coordination, albeit at the cost of reduced data control, lower sovereignty, and potential trust concerns. Hybrid models, combining centralized interoperability services with federated governance enforcement, may represent a pragmatic pathway for balancing scalability, interoperability, and trust within PI-aligned ecosystems.

At the same time, Table 1 shows that adoption viability remains the limiting dimension across roadmap areas. This finding indicates that PI implementation is currently constrained less by the absence of standards or governance concepts and more by the difficulty of operationalizing them at scale across heterogeneous actors with uneven capabilities, incentives, and risk tolerance.

The combined DTLF and EMDS approach is structurally compatible with PI roadmap expectations: DTLF-aligned semantics and event exchange provide the basis for interoperability and compliance-by-design information sharing, while federated governance models provide the institutional and technical scaffolding for sovereignty and trust. However, structural compatibility does not automatically deliver ecosystem-scale uptake. The primary gap is the "middle layer" that translates standards and governance principles into adoptable products and routines.

Crucially, this middle layer is socio-technical, not purely technical. It includes not only API tooling, registries, and conformance, but also onboarding processes, liability allocation, contractual templates, and operational accountability mechanisms that make participation safe and worthwhile for each actor. Without this translation layer, ecosystems remain stuck in a "pilot plateau", where interoperability is demonstrated in controlled settings but struggles to diffuse broadly.

Based on the results and the barriers reported in the evidence base (fragmentation, uneven adoption, security exposure, trust and reluctance to share sensitive data), the middle layer can be operationalized as four requirements:

1. Interoperability assurance, beyond "having a standard" Semantic alignment (UN/CEFACT, GS1) is necessary, but scaling requires conformance artefacts that reduce bilateral integration work: test suites, validated implementation profiles, reference APIs, and versioning policies. Without these, API fragmentation and evolution costs remain a structural brake on adoption.
2. Discoverability and verification of interfaces and claims Actors need a way to discover what APIs exist, what versions they implement, and whether they are conformant and secure. This implies registries and verification services (or equivalent mechanisms) that turn "interoperable by design" into "interoperable in practice."
3. Governance that is enforceable, lightweight, and usable EMDS governance is strongest when sovereignty and trust controls are enforceable in real workflows (roles, access control, auditing, usage control), but it must also be usable. If governance introduces excessive overhead or unclear accountability, participation declines, especially for SMEs. Liability and responsibilities must be explicit, because liability uncertainty is an adoption barrier even when technical controls exist.
4. Security-by-design as an adoption prerequisite Expanded API-based exchange increases the attack surface and risk perception. A credible cybersecurity posture therefore becomes a precondition for adoption, not a downstream improvement. This includes secure API engineering, continuous monitoring, and alignment with recognised information security and data quality standards where relevant.

To make the discussion actionable, Table 2 (deployment checklist) can be positioned as an implementation translation of Table 1's findings. It should be presented as the minimum set of actions that directly target the lowest-scoring dimension (adoption viability) while preserving the strengths of DTLF (platform primitives) and EMDS (trust/sovereignty).

Table 2: PI-aligned deployment checklist (from standards and governance to adoption)

Platform layer	Governance layer	Adoption layer
Implement semantic models aligned with UN/CEFACT/GS1 and maintain model lifecycle governance (CCTS-based tooling where relevant).	Adopt data sovereignty and usage-control mechanisms (role-based access, auditing, identity/credentials consistent with IDS/IDSA patterns.	Pair governance with incentives and onboarding support (address organisational and emotional reluctance)
Provide real-time event exchange primitives (ETA, ITU status) and single-window submission patterns for B2A/B2B.	Clarify liability and governance responsibilities across actors (explicitly flagged as adoption constraints).	Treat API security and evolution as first-class adoption risks; invest in secure, scalable API infrastructure and maintenance processes (Lamothe et al., 2021; ISO/IEC 25012; ISO 27001).
Establish interoperability testing and conformance artefacts to mitigate fragmented implementations (gap identified via "lack of universal APIs").		

The research questions should be explicitly framed as closing the "middle layer" gap identified above:

- R1. Regulation–governance interaction: How do binding regulations (eFTI, Data Governance Act, Data Act) interact with non-binding guidance (EIF, IDSA) to shape compliance outcomes and willingness to participate in logistics data sharing ecosystems?
- R2. Trust and participation: Which governance patterns measurably increase trust and participation under heterogeneous stakeholder incentives and commercially sensitive data?
- R3. Interoperability assurance: Which mechanisms (conformance suites, registries, API discovery/verification, implementation profiles) most reduce integration costs and mitigate fragmentation in open-standard ecosystems?
- R4. Security operationalization: How can cybersecurity be operationalized as continuous assurance (certification, monitoring, incident response) for PI-aligned data platforms without raising participation barriers?

In summary, the results support a layered interpretation of PI roadmap implementation: DTLF provides the best basis for platform functionality, EMDS provides the best basis for governance, and the decisive bottleneck is the socio-technical middle layer that makes both adoptable at scale. Positioning future work and implementation guidance around interoperability assurance, usable governance, and security-by-design is therefore the most direct route to improving adoption viability, the weakest dimension identified in Table 1.

5.1 Business and academic implications

Business implications differ according to the role of actors within the logistics chain. For logistics operators, shippers, and platform providers, the results suggest that value creation depends less on the mere adoption of standards and more on reducing participation frictions within multi-actor ecosystems. This implies prioritizing interoperable APIs, conformance tooling, onboarding support, and governance mechanisms that simplify integration across heterogeneous actors.

For SMEs, the findings highlight the importance of lightweight onboarding mechanisms and shared interoperability services capable of reducing technical and organizational barriers to participation. Adoption support mechanisms, such as implementation templates, certified connectors, and managed compliance services, may be necessary to enable participation even in the absence of advanced internal IT capabilities.

For public authorities and regulators, the results indicate that policy objectives such as eFTI adoption require not only regulatory alignment but also operational assurance mechanisms capable of reducing fragmentation and increasing trust. Without interoperability assurance, governance clarity, and cybersecurity support, the risk remains that deployments will evolve into fragmented and non-interoperable ecosystems despite the existence of common standards.

From an academic perspective, this paper contributes to the debate by conceptualizing PI digitalization as a coupled problem involving interoperability, governance, and adoption simultaneously, rather than treating them as isolated dimensions. Methodologically, the proposed fitness framework provides a structured approach for comparing PI-enabled

initiatives across different implementation contexts and roadmap variants using common evaluation dimensions and evidence-based scoring logic.

The results also suggest that adoption viability should become a primary research priority. Future research should therefore focus on measuring adoption frictions, evaluating the effectiveness of governance patterns under commercially sensitive data-sharing conditions, and quantifying the impact of interoperability assurance mechanisms on participation and ecosystem scalability.

6 Conclusion and next steps

This research assessed the fitness of the PI roadmap to support the Logistics/Commercial Data Platform through an analysis of two foundational components: DTLF-aligned open standard data exchange and federated data space governance models (EMDS/LaaS). Using a qualitative model of fitness consisting of capability coverage, interoperability readiness, governance adequacy, and adoption viability, and integrating the results from Table 1, this analysis has demonstrated that both approaches are structurally consistent with the PI roadmap. DTLF offers the strongest basis for platform-level interoperability through common semantics and operational exchange patterns, while EMDS offers the strongest basis for sovereignty and trust through governance-by-design.

Adoption viability, however, remains the key limiting factor in all domains of the roadmap. Persistent barriers to adoption, including API fragmentation and cost of evolution, cybersecurity risk in API-based systems, semantic complexity, and socio-technical factors like trust, sharing of sensitive information, and accountability, continue to impede the transition from standards and governance concepts to practical and accessible implementation. Therefore, to support the successful implementation of the PI, there is a need to develop a "translation layer" to operationalize interoperability and governance into accessible and practical products and processes. This includes conformance by design, discoverability and verification, lifecycle governance to support evolving standards, and security by design with liability allocation and onboarding.

The future direction for work will be on transitioning from qualitative fitness measures toward measurable validation and deployment-oriented experimentation. The priority measures include: (i) the development and testing of interoperability assurance mechanisms (implementation profiles, conformance test suites, registries, verified APIs); (ii) exploring lightweight, enforceable governance patterns that balance sovereignty with usability, such as explicit liability and accountability models; (iii) operationalizing cybersecurity as continuous assurance (secure API engineering, monitoring, incident response); and (iv) quantifying adoption economics and organizational constraints via corridor-level case studies (e.g., onboarding time, integration cost proxies, conformance pass rates, participation metrics). The above measures will enhance the evidentiary base required for transitioning from pilot initiatives toward ecosystem-scale PI adoption.

AI and ML integration, represents a natural extension of PI data platforms, enabling collaborative analytics without compromising data sovereignty. However, challenges such as bias propagation, data quality heterogeneity, and governance of model outputs remain open research questions.

7 Acknowledgment

The development of the proposed methodology is part of the EU-funded project, HOLOGISTICS, which received funding from the Horizon Europe Program under the Grant Agreement number 101202556.

References

- Apruzzese, M., & Bresseleers, P. (2023). Business requirements for logistics data sharing: recommendations from the FENIX project. *Transportation Research Procedia*. doi:10.1016/j.trpro.2023.11.488
- Bossong, P., Reinhardt, A., & Elbert, R. (2024). Adoption drivers and barriers of digital freight transport platforms—An intermodal case study. *Electronic Markets, 35*. doi:10.1007/s12525-025-00780-0
- Cavalieri, S., & Gambadoro, S. (2023). Proposal of Mapping Digital Twins Definition Language to Open Platform Communications Unified Architecture. *Sensors (Basel, Switzerland), 23*. doi:10.3390/s23042349
- Dasaklis, T., Kopanaki, E., Chountalas, P., Rachaniotis, N., Voutsinas, T., Giannakis, K., & Chondrokoukis, G. (2024). Exploring the Implementation Challenges of the Electronic Freight Transport Information (eFTI) Regulation: An Empirical Perspective from Greece. *Logistics*. doi:10.3390/logistics8010030
- Gyard, A., Abedian, S., Gribbon, P., Manias, G., Van Nuland, R., Zatloukal, K., . . . Martí-Bonmatí, L. (2025). Lessons Learned from European Health Data Projects with Cancer Use Cases: Implementation of Health Standards and Internet of Things Semantic Interoperability. *Journal of Medical Internet Research*. doi:10.2196/66273
- Hellmeier, M., Pampus, J., Qarawlus, H., & Howar, F. (2023). Implementing Data Sovereignty: Requirements & Challenges from Practice. *ARES '23: Proceedings of the 18th International Conference on Availability, Reliability and Security.* doi:10.1145/3600160.3604995
- Ivezic, N., Jelisic, E., Jankovic, M., Kulvatunyou, B., Kehagias, D., & Marjanovic, Z. (2022). Advancing Data Exchange Standards for Interoperable Enterprise Networks. *Workshop of I-ESA'22.* Valencia, Spain.
- Janssen, M., Brous, P., Estevez, E., Barbosa, L., & Janowski, T. (2020). Data governance: Organizing data for trustworthy Artificial Intelligence. *Government Information Quarterly, 37*. doi:10.1016/j.giq.2020.101493
- Korca, B., & Costa, E. (2021). Directive 2014/95/EU: building a research agenda. *Journal of Applied Accounting Research*. doi:10.1108/JAAR-05-2020-0085
- Milčiuvienė, S., Amilevičius, D., & Raskinis, G. (2024). Research Data Management Within the EU Data Policy Framework. *Baltic Journal of Law & Politics*, 168-187. doi:10.2478/bjlp-2024-00020
- Otto, B., & Jarke, M. (2019). Designing a multi-sided data platform: findings from the International Data Spaces case. *Electronic Markets, 29*, 561-580. doi:10.1007/s12525-019-00362-x
- Prifti, K., & Fosch-Villaronga, E. (2024). (2024). Towards experimental standardization for AI governance in the EU. *Computer Law Security Review*. doi:10.1016/j.clsr.2024.105959
- Ruohonen, J., & Mickelsson, S. (2023). Reflections on the Data Governance Act. *Digital Society, 2*. doi:10.1007/s44206-023-00041-7
- Shabani, M. (2022). Will the European Health Data Space change data sharing rules? *Science, 375*, 1357-1359. doi:10.1126/science.abn4874
- Tijan, E., Jović, M., Jardas, M., & Gulić, M. (2019). The Single Window Concept in International Trade, Transport and Seaports. *Pomorstvo, 33*, 130-139. doi:10.31217/P.33.2.2
- Von Scherenberg, F., Hellmeier, M., & Otto, B. (2024). Data Sovereignty in Information Systems. *Electronic Markets*. doi:10.1007/s12525-024-00693-4

- Zeng, F., Chen, A., Xu, S., Chan, H., & Xu, L. (2025). Digitalization of the maritime shipping service: defining the digital freight forwarder. *International Journal of Physical Distribution & Logistics Management*. doi:10.1108/ijpdlm-04-2024-0165

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Place your logo here

H-CPI-ARP: A Hierarchical Proxy Mechanism Bridging High-Tech Hubs and Low-Tech Logistics Units

Wenqing Lei 1[1], Yong-Hong Kuo 2[1], and George Q. Huang 3[2]
1. Department of Data and Systems Engineering, The University of Hong Kong, Hong Kong SAR
2. Department of Industrial and Systems Engineering, The Hong Kong Polytechnic University, Hong Kong SAR
Corresponding author: gq.huang@polyu.edu.hk

Abstract: *As the Cyber-Physical Internet (CPI) emerges, logistics networks are rapidly evolving toward hyperconnectivity. However, integrating massive volumes of low-value, passive logistics units expose traditional flat address resolution protocols to severe scalability challenges, often triggering catastrophic network broadcast storms. Building upon the Converged Address Resolution Protocol (CPI-ARP), this paper proposes a Hierarchical Address Resolution and Proxy Mechanism (H-CPI-ARP). This mechanism leverages high-tech mobile hubs equipped with active network capabilities (e.g., smart CPI-Trucks) to act as edge computing proxies. Crucially, it provides aggregated identity mapping and direct location resolution for both their nested Physical Shipment Units (PSUs) and the massive number of individual low-tech parcels contained within them. High-fidelity discrete-event simulations demonstrate that this 1:1: N hierarchical approach effectively eliminates MAC-layer broadcast overhead and resolves the resolution latency bottleneck in extreme high-concurrency scenarios. Ultimately, it provides a highly robust and cost-effective technological pathway for achieving granular, global-scale logistics asset visibility.*

Keywords: *Physical Internet, Cyber-Physical Internet, Address Resolution, Proxy Mechanism*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan: ☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories), ☐ Transportation Equipment, ☒ PI Networks, ☒ System of Logistics Networks, ☐ Vertical Supply Consolidation, ☐ Horizontal Supply Chain Alignment, ☐ Logistics/Commercial Data Platform, ☐ Access and Adoption, ☐ Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☒ Poster, ☐ Flash Video, ☒ In-Person presentation*

1 Introduction

The Physical Internet (PI) has emerged as a disruptive vision for global logistics transformation, aiming to replace traditionally fragmented supply chains with an open, hyperconnected network system based on standardized modular containers and unified protocols (Ballot, 2019; Montreuil et al., 2013; Treiblmaier et al., 2020). To operationalize this interconnectivity goal, the Cyber-Physical Internet (CPI) has been introduced. By constructing an advanced cyber layer, the CPI tightly integrates physical logistics operations with digital twin environments, thereby providing full-lifecycle, end-to-end visibility for massive logistics assets (Tao et al., 2019; Wu et al., 2024, 2025)

However, achieving such granular, real-time tracking at the global scale envisioned by the PI imposes severe addressing and scalability challenges on the underlying network. Integrating physical entities into the cyber layer relies on robust address resolution mechanisms, such as the Converged Address Resolution Protocol (CPI-ARP), which maps logical Physical Internet

Protocol (PIP) addresses to Physical Media Access Control (PMAC) identifiers (Lei et al., 2025). Although this protocol theoretically synchronizes cyber-physical identities, traditional flat addressing architectures exhibit critical vulnerabilities in hyperconnected, high-density Internet of Things (IoT) environments. Address resolution heavily depends on network broadcasting; when massive nodes concurrently initiate routing queries, they rapidly exhaust channel capacity, triggering catastrophic "broadcast storms" (Al-Fuqaha et al., 2015; Glaroudis et al., 2020). Such massive redundant traffic not only leads to severe packet collisions and loss but also exacerbates end-to-end tracking latency to unacceptable levels.

This network communication bottleneck is further amplified within real-world logistics ecosystems (Ouyang et al., 2025). On one hand, high-tech "Asset Hosts" (e.g., smart CPI-Trucks) function as mobile hubs equipped with active network protocol stacks and edge computing capabilities. On the other hand, low-tech "Shipment Units" (e.g., physically nested parcels) typically carry only passive identifiers and lack the capacity for independent communication. As the density of passive tracking entities grows exponentially within the PI network—such as in large-scale automated warehousing or mega-hubs (Beniwal & Singhrova, 2022; Leng et al., 2021)—relying on a flattened architecture where every micro-level parcel must independently trigger network handshakes and resolution responses becomes mathematically and operationally unsustainable.

This dilemma highlights a critical research gap: how to manage the network visibility of massive, nested, low-tech logistics entities through edge computing proxy mechanisms, without triggering broadcast storms or compromising the stability of the CPI infrastructure. Specifically, this paper investigates the following research questions:

- What are the fundamental requirements for a multi-tier identity hierarchy that effectively reflects physical aggregation while maintaining individual traceability for each parcel?
- How can proxy-based intercepted queries and aggregated replies mitigate the resolution overhead in high-density logistics networks?
- How does a dynamic lease-based mechanism ensure real-time synchronization between physical loading/unloading events and the digital address state?

To address these challenges, this study proposes a Hierarchical Address Resolution and Proxy mechanism (H-CPI-ARP). The core logic of this mechanism is grounded in the "loading-as-binding" principle: during the physical loading phase, parcels are logically nested within a Physical Shipment Unit (PSU), which is subsequently nested under the transport host's PIP. By empowering the high-tech asset host to act as an edge computing surrogate, the system can accomplish the location resolution of thousands of nested entities through a single aggregated proxy response, fundamentally eliminating the network conditions that breed broadcast storms (Qiu et al., 2020; Shi et al., 2016). Based on the investigation, the main contributions of this work are as follows:

- We establish a three-tier binding model that bridges high-tech hubs and low-tech units, ensuring "network-level visibility" for passive assets.
- Introduced an extended H-CPI-ARP protocol featuring specialized mode flags and dynamic lease fields, enabling mobile hubs to perform autonomous proxy resolution and proactive address recycling.
- Developed a recursive binding logic that synchronizes the physical loading process with digital identity mapping, ensuring "plug-and-play" visibility for nested logistics assets.

2 Hierarchical Address Resolution and Proxy Mechanism

The core of the proposed **H-CPI-ARP** mechanism is a transition from a flattened address space to a structured, three-tier hierarchy designed to synchronize physical logistics aggregation with network-layer resolution. To address the scalability limits of baseline protocols, this section defines the hierarchical mapping model, the extended packet structure, and the proxy-based resolution workflow.

2.1 Multi-tier Hierarchical Identity Model

The operational foundation of the mechanism is the Hierarchical Surrogate Model, as illustrated in Figure 1. This model categorizes logistics entities into three functional tiers based on their communication capabilities and physical roles.

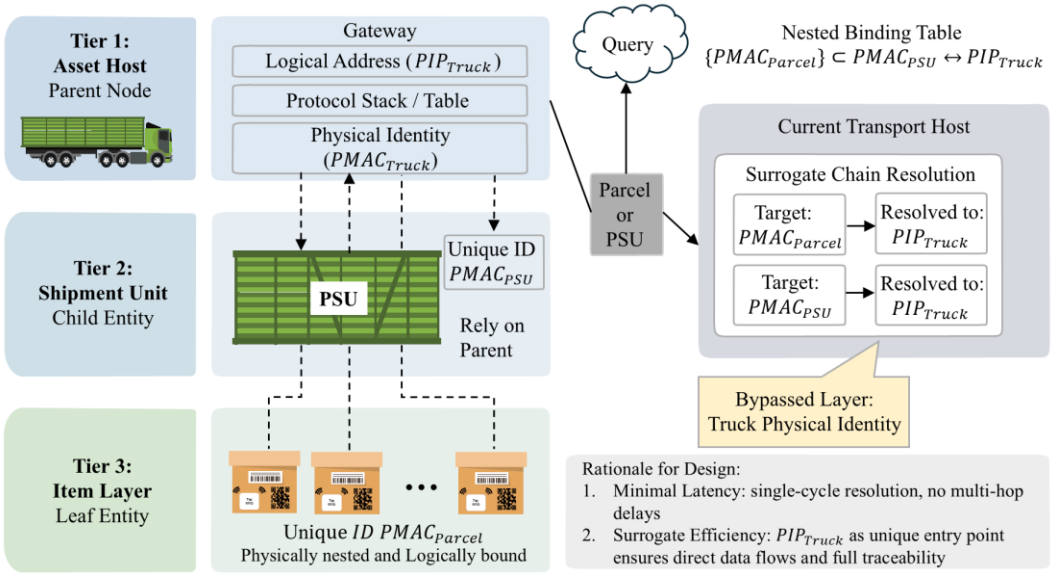


Figure 1. Muti-tier Hierarchical Surrogate Model

1. Tier 1: Asset Host (Parent Node): Represented by the CPI-Truck, this layer serves as the communication gateway. It possesses a logical address (PIP_{Truck}) and the necessary protocol stack to intercept and respond to network queries.
2. Tier 2: Shipment Unit (Child Entity): Represented by the PSU, this layer acts as the primary addressable container. It carries a unique $PMAC_{PSU}$ but relies on the Parent Node for network access.
3. Tier 3: Item Layer (Leaf Entities): Represented by individual Parcels, each maintaining a unique $PMAC_{Parcel}$. These units are physically nested within the PSU and logically bound to the Truck's PIP through the surrogate chain.

As depicted in Figure 1, the mechanism deliberately bypasses the intermediate physical identity of the truck (PIP_{Truck}) during active resolution. Instead, it establishes a direct functional mapping: $\{PMAC_{Parcel}\} \subset PMAC_{PSU} \leftrightarrow PIP_{Truck}$. This design ensures that any external resolution request targeting a parcel or a PSU is directly routed to the PIP of the current transport host, minimizing lookup latency and preventing the multi-hop resolution delays characteristic of traditional hierarchical networks.

The rationale for bypassing the $PMAC_{Truck}$ layer in the resolution chain is twofold:

1. Minimal Latency: By directly binding the shipment units to the communication-capable PIP of the asset host, the system avoids redundant multi-hop lookups, ensuring a single-cycle resolution for nested items.

2. Surrogate Efficiency: In the CPI environment, the PIP functions as the unique entry point to the logistics subnet. Since the Truck host's physical identity and logical address are inherently synchronized, linking the PSU directly to the PIP provides the most direct path for data flows while maintaining full traceability of the underlying PMAC identifiers.

2.2 H-CPI-ARP Protocol Design and Field Extensions

To support the hierarchical proxy logic, the standard CPI-ARP header is extended with specialized control fields that enable mobile hubs to orchestrate resolution for nested entities. The bit-level structure of the H-CPI-ARP packet is illustrated in Figure 2.

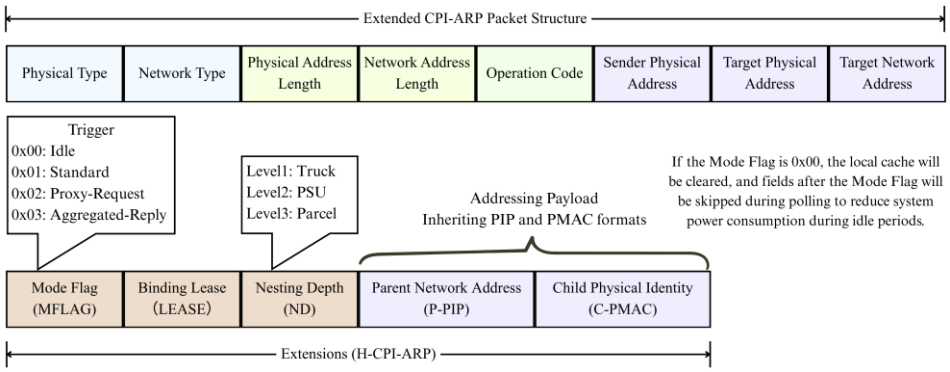


Figure 2. Structure Extension of CPI-ARP

The protocol introduces four critical extensions to the baseline frame:

- Mode Flag (MFLAG): This is a 1-byte logical switch for the resolution engine. A value of 0x00 represents initial or idle mode. This indicates that the host is unbound/expired; operates as a standard node without interception logic. 0x01 represents a standard peer-to-peer query. 0x02 (Proxy-Request) signals a query targeting a nested entity, prompting host interception. 0x03 (Aggregated-Reply) indicates a proxy-generated response containing host connectivity data.
- Binding Lease (LEASE): A 2-byte field defining the temporal validity of the Parent-Child mapping. This ensures that the global resolution cache is automatically updated as the physical topology changes, such as when a PSU is unloaded from a truck.
- Nesting Depth: A 4-bit indicator used to specify the granularity of the target entity within the logistics hierarchy (e.g., Level 1 for the Truck host, Level 2 for the PSU, and Level 3 for the individual Parcel).
- Addressing Payload: The packet carries the Parent Logical Address (P-PIP) and the Child Physical Identity (C-PMAC). This dual-addressing scheme allows the resolver to associate the identity of a low-tech unit directly with the communication endpoint of its high-tech host.

When the MFLAG transitions to 0x00 (Idle Mode) upon lease expiration or a physical unloading event, the corresponding entries in the Nested Binding Table are purged to prevent stale address resolution. In this state, the extended protocol fields (LEASE, Nesting Depth, and Addressing Payload) are retained in the header structure to maintain bit-alignment but are logically nullified (zero-padded) to ensure system-wide consistency.

2.3 Proxy-Based Resolution and Surrogate Logic

The operational execution of H-CPI-ARP is characterized by an edge-interception mechanism that prevents low-tech units from participating in complex network handshakes. As depicted

in the sequence diagram in Figure 3, the surrogate workflow is strictly governed by the state of the mobile hub's Nested Binding Table, divided into three distinct phases:

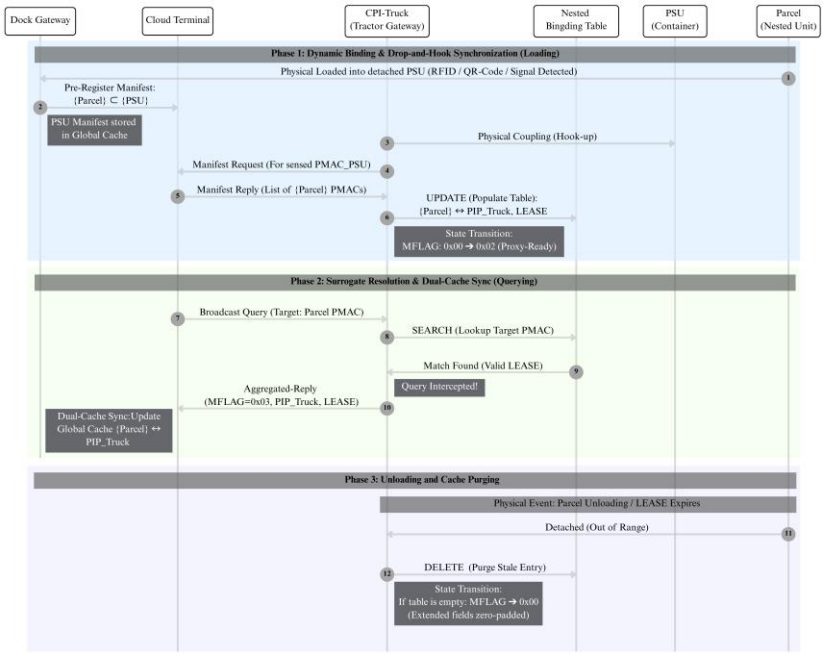


Figure 3. H-CPI-ARP Proxy Addressing and State Switching Sequence Diagram

Phase 1: This phase aligns with realistic "drop-and-hook" logistics operations, accommodating the physical separation of the high-tech tractor (Host) and the low-tech container (PSU) housing passive nested items. During dock operations, nested items are loaded into a detached PSU, and the dock gateway pre-registers this $\{PMAC_{Parcel}\} \subset PMAC_{PSU}$ mapping into the Global Resolution Cache. Subsequently, when the CPI-Truck connects with the loaded PSU, its edge gateway senses the macroscopic $PMAC_{PSU}$. The truck then synchronizes with the global cache to download the nested manifest, automatically populating its Local Nested Binding Table to complete the $\{PMAC_{Parcel}\} \leftrightarrow PIP_{Truck}$ mapping. Consequently, the gateway transitions from the Idle state (MFLAG = 0x00) to the Proxy-Ready state (MFLAG = 0x02), signifying its readiness to orchestrate resolution.

Phase 2: Surrogate Resolution Flow (Querying & Dual-Cache Sync). When an external resolver (e.g., a Cloud Terminal) broadcasts a query for a specific $PMAC_{Parcel}$, The resolution is handled through local surrogate logic:

1. Intercept: The CPI-Truck acts as a mobile proxy gateway. Instead of flooding the internal subnet, it executes a SEARCH operation within its local Binding Table.
2. Aggregated Proxy Reply & Global Caching: Upon a successful match, the query is intentionally intercepted. The truck generates a single aggregated response containing its own logical address PIP_{Truck} , $PMAC_{Truck}$ and the LEASE duration, setting the MFLAG to 0x03. Crucially, upon receiving this reply, the external resolver updates its Global Resolution Cache with the $\{PMAC_{Parcel}\} \subset PMAC_{PSU} \leftrightarrow PIP_{Truck}$ mapping. This dual-cache synchronization ensures that all subsequent queries within the LEASE window are routed via direct unicast, effectively eliminating redundant broadcast storms across the hub.

Phase 3: Unloading and State Purging (Unloading). To maintain absolute synchronization between the physical topology and the digital twin, the protocol strictly enforces a cache

purging mechanism. When a physical unloading event is detected, or if the LEASE timer expires without subsequent physical sensing, the gateway triggers an Unbind Event. The corresponding PMAC entries are actively DELETED from the Local Nested Binding Table, preventing the generation of stale addresses. If the binding table becomes completely empty, the CPI-Truck automatically reverts to the standard Idle Mode (MFLAG = 0x00), and the extended protocol fields are zero-padded until a new loading event occurs.

3 Experimental Evaluation

To rigorously validate the proposed H-CPI-ARP mechanism, a high-fidelity discrete-event simulation was implemented. Going beyond idealized mathematical models, the simulation incorporates realistic physical network constraints typical in dense PI-hubs, including Carrier-Sense Multiple Access with Collision Avoidance (CSMA/CA) backoff penalties, edge gateway serialization bottlenecks, and background channel noise.

3.1 Scalability and Protocol Overhead

The spatial complexity of the resolution mechanisms was first evaluated by measuring the total control of traffic under varying cargo densities (number of nested parcels, N), as illustrated in Figure 4.

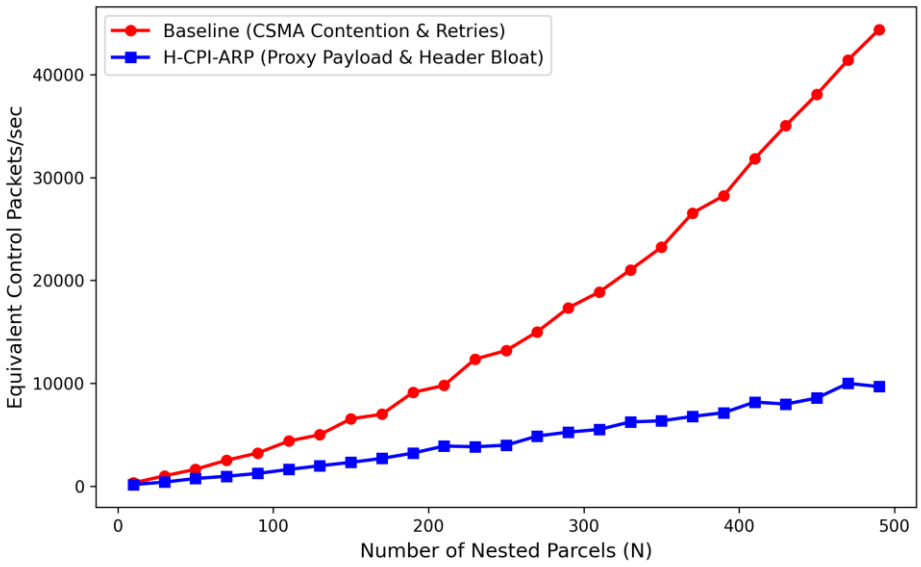


Figure 4. Network Overhead Accounting for Protocol Overhead

As N increases, the simultaneous responses trigger severe MAC-layer contention and subsequent avalanche retransmissions, pushing the equivalent control packets to unsustainable levels (exceeding 40,000 packets/sec).

In contrast, the proposed H-CPI-ARP effectively decouples network overhead from cargo density. By utilizing the proxy interception mechanism, the control packet volume is drastically reduced. H-CPI-ARP overhead exhibits a non-linear upward curve (Figure 4). This precision modeling reflects the Protocol Header Bloat during MTU (Maximum Transmission Unit) fragmentation: as hundreds of PMACs are aggregated, the resulting fragments require duplicated IP/UDP headers. However, even with this explicit fragmentation penalty pushing the overhead to approximately 10,000 packets/sec at N=500, the hierarchical proxy still achieves a 75% reduction in total network traffic compared to the baseline, ensuring a highly robust scalability margin.

3.2 End-to-End Latency and Edge Compute Trade-offs

To assess temporal robustness, Figure 5 demonstrates the average resolution latency under an increasing query rate (Queries Per Second, QPS) for a high-density host (N=400). The results reveal a profound system-level trade-off. As QPS surpasses 50, the Baseline mechanism suffers from catastrophic channel contention, causing latency to spike exponentially and hit the 200 ms timeout threshold, rendering the network completely unresponsive. Conversely, H-CPI-ARP absorbs the query stress by leveraging edge computing at the mobile gateway (Truck). While this introduces a baseline computational delay for gateway lookups, it successfully shields the core network from broadcast storms. As QPS approaches 200, the edge CPU queueing delay gradually pushes the latency to roughly 60 ms. This explicitly demonstrates the design philosophy of trading edge compute power for network channel availability, successfully preventing network-wide congestion collapse and ensuring system availability under extreme loads.

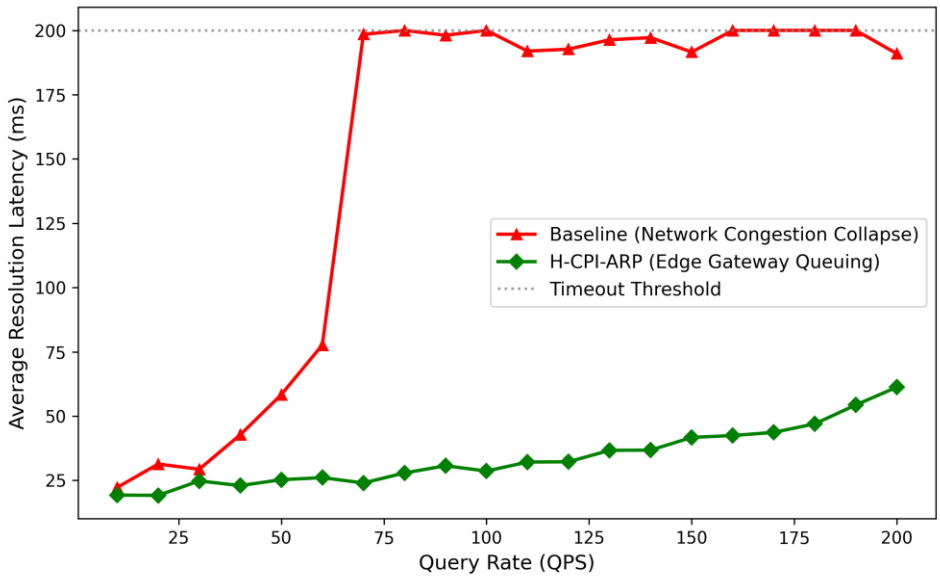


Figure 5. End-to-End Latency and Edge Compute Trade-offs

3.3 Impact of Binding Lease on Digital Twin Synchronization

A critical parameter in H-CPI-ARP is the LEASE duration, which dictates the state synchronization between the physical parcel and its digital twin. To rigorously determine the optimal binding lease (L_{opt}), we conducted a Monte Carlo simulation of over 1,000 iterations to account for stochastic network jitters. Dwell times during unloading events were modeled using a Weibull distribution ($\lambda = 300, k = 1.5$) to reflect realistic logistics operations. The resulting 95% confidence intervals are depicted as shaded regions in Figure 6. To balance the Renewal Traffic Overhead (cost of maintaining the binding) against the Stale Address Rate (probability of resolving an outdated mapping after unloading), we formulated a Normalized Joint Cost Function $J(L)$:

$$J(L) = \alpha \cdot \hat{T}_{renew}(L) + \beta \cdot \hat{P}_{Stale}(L)$$

Where $\hat{T}_{renew}(L)$ and $\hat{P}_{Stale}(L)$ represent the min-max normalized traffic overhead and stale address probability, respectively. Given the constrained channel capacity in realistic wireless logistics hubs, we assigned a higher penalty weight to network congestion ($\alpha = 0.6$) than to routing fidelity($\beta = 0.4$).

As depicted by the purple dashed curve in Figure 6, the joint cost $J(L)$ exhibits a strict global minimum at $L_{opt} \approx 70$ seconds. This result is highly significant: $a \approx$ 1-minute binding lease allows the protocol to bypass the exponential region of renewal traffic while keeping the expected stale address rate strictly under 15%. Remarkably, this mathematically derived 70-second optimum aligns perfectly with standard IoT keep-alive timer configurations (e.g., MQTT), ensuring seamless compatibility with existing industrial wireless infrastructures.

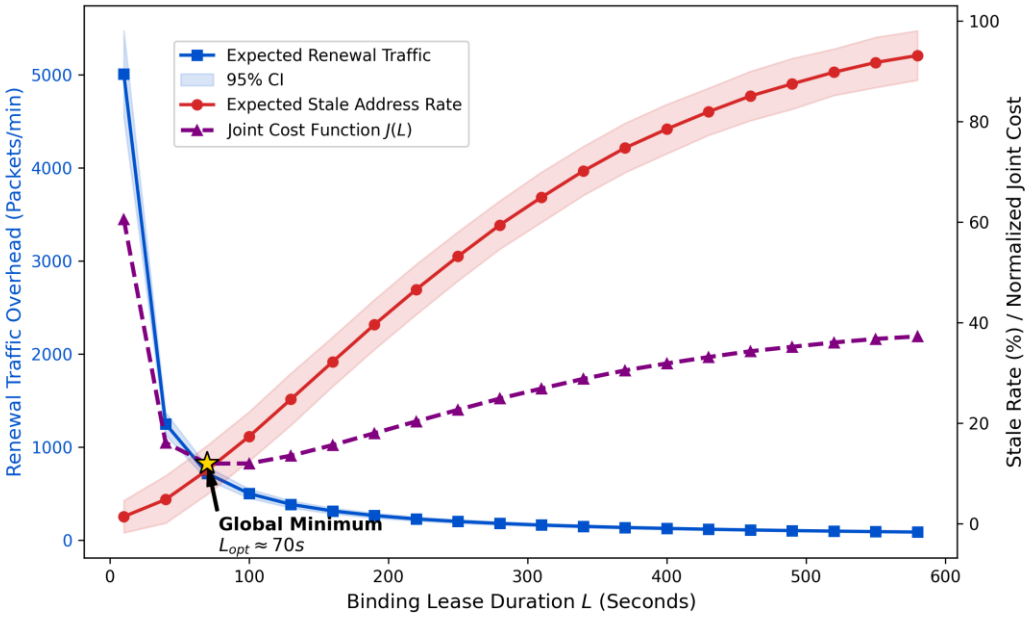


Figure 6. Monte Carlo Simulation and Joint Cost Optimization of Binding Lease

4 Discussion

The experimental results confirm H-CPI-ARP's robustness under physical constraints. Beyond raw metrics, this approach introduces a paradigm shift in CPI addressing and visibility.

4.1 Scalability via Edge Intelligence: Trading Compute for Bandwidth

Scaling visibility to billions of logistics units is a foundational PI challenge. Baseline flat architecture inevitably failed in high-density hubs due to broadcast storms. H-CPI-ARP solves this via an Edge Intelligence paradigm. By delegating resolution authority to mobile asset hosts (e.g., PI-Trucks), the system translates $O(N)$ network-wide communication overhead into localized, $O(1)$ edge-compute queuing tasks. While introducing marginal gateway processing delays, it acts as a critical shock absorber. Sacrificing negligible edge compute power to preserve core channel availability is the prerequisite for infinite horizontal scalability.

4.2 Context-Aware Cyber-Physical Synchronization

The $L_{opt} \approx 70$ s global minimum balances digital routing fidelity and network overhead. While mathematically optimal for rapid hub unloading, a static timer wastes signaling during long-haul transit. Thus, H-CPI-ARP advocates for Context-Aware Dynamic Leasing. By integrating GPS/geofencing, surrogate hosts dynamically adjust synchronization based on operational context. During highway transit, gateways aggressively extend the LEASE to hours to preserve bandwidth. Upon breaching a PI-hub geofence, the lease dynamically

shrinks back to the optimized 70s window, allocating high-resolution fidelity exactly where physical state transitions occur.

4.3 Bridging the Technological Divide in Global Logistics

A prevailing misconception is that granular PI visibility requires equipping every parcel with high-energy transceivers (e.g., 5G). 1:1: N hierarchical proxy democratizes visibility. By relying on high-tech surrogate hosts to intercept tracking queries, passive parcels (equipped with merely basic RFID/QR codes) seamlessly "piggyback" on the host's advanced network capabilities. This inclusive architecture guarantees that low-tech assets achieve real-time cyber-physical visibility, providing a cost-effective, incremental path for industry adoption without unsustainable hardware overhauls.

5 Conclusion

The realization of the global CPI hinges on the ability to maintain granular, real-time visibility over billions of ubiquitous logistics assets. However, traditional flat addressing mechanisms expose the core network to unsustainable MAC-layer broadcast storms and congestion collapse.

To overcome this scalability bottleneck, this paper proposed the H-CPI-ARP. By establishing a dynamic 1:1: N surrogate binding relationship, the mechanism leverages the edge computing capabilities of mobile PI-Hubs (e.g., smart trucks) to intercept and aggregate tracking queries. Our high-fidelity simulations, which rigorously accounted for MTU fragmentation and edge gateway serialization constraints, demonstrated that H-CPI-ARP effectively mitigates $O(N)$ network broadcast storms, yielding a 75% reduction in control traffic while maintaining bounded resolution latency (< 80 ms) under extreme concurrency.

Furthermore, through a Monte Carlo simulation and a Normalized Joint Cost Function, we mathematically identified an optimal binding lease $L_{opt} \approx 70\ s$. This metric inspired the proposal of a Context-Aware Dynamic Leasing strategy, ensuring high-fidelity state synchronization in docking zones while preserving bandwidth during long-haul transit. Ultimately, H-CPI-ARP bridges the technological divide, allowing low-tech parcels to seamlessly piggyback on the advanced network infrastructure of their physical hosts, paving a cost-effective path for inclusive PI deployment.

While H-CPI-ARP presents a robust foundation for PI visibility, several avenues remain for future exploration. First, the current model assumes a two-tier hierarchy (Parcel to Truck). Future research will extend this protocol to support Multi-hop Nested Hierarchies (e.g., $Parcel \in Tote \in Pallet \in Container \in Truck$), addressing the complex recursive proxy logic required for deeply nested physical encapsulations.

Secondly, we aim to integrate Machine Learning algorithms into the mobile gateway to predict physical unloading events. A predictive, AI-driven contextual engine could autonomously adjust the LEASE timer based on historical transit data and real-time traffic conditions, further optimizing the trade-off between network efficiency and digital twin fidelity. Finally, transitioning from simulated environments to a small-scale, real-world hardware pilot in an industrial warehouse will be critical to validating the physical MAC-layer interactions of the protocol.

References

- Al-Fuqaha, A., Guizani, M., Mohammadi, M., Aledhari, M., & Ayyash, M. (2015). Internet of Things: A Survey on Enabling Technologies, Protocols, and Applications. *IEEE Communications Surveys & Tutorials*, *17*(4), 2347–2376. https://doi.org/10.1109/COMST.2015.2444095
- Ballot, E. (2019). The Physical Internet. In H. Zijm, M. Klumpp, A. Regattieri, & S. Heragu (Eds.), *Operations, Logistics and Supply Chain Management* (pp. 719–734). Springer International Publishing. https://doi.org/10.1007/978-3-319-92447-2_31
- Beniwal, G., & Singhrova, A. (2022). "A systematic literature review on IoT gateways." *Journal of King Saud University - Computer and Information Sciences*, *34*(10), 9541–9563. https://doi.org/10.1016/j.jksuci.2021.11.007
- Glaroudis, D., Iossifides, A., & Chatzimisios, P. (2020). Survey, comparison and research challenges of IoT application protocols for smart farming. *Computer Networks*, *168*, 107037. https://doi.org/10.1016/j.comnet.2019.107037
- Lei, W., Li, M., Kuo, Y.-H., & Huang, G. Q. (2025). Converged Address Resolution Protocol for Traceability and Visibility in Cyber–Physical Internet. *IEEE Internet of Things Journal*, *12*(14), 27322–27338. https://doi.org/10.1109/JIOT.2025.3562803
- Leng, J., Yan, D., Liu, Q., Zhang, H., Zhao, G., Wei, L., Zhang, D., Yu, A., & Chen, X. (2021). Digital twin-driven joint optimisation of packing and storage assignment in large-scale automated high-rise warehouse product-service system. *International Journal of Computer Integrated Manufacturing*, *34*(7–8), 783–800. https://doi.org/10.1080/0951192X.2019.1667032
- Montreuil, B. (2011). Toward a Physical Internet: Meeting the global logistics sustainability grand challenge. *Logistics Research*, *3*(2–3), 71–87. https://doi.org/10.1007/s12159-011-0045-x
- Montreuil, B., Meller, R. D., & Ballot, E. (2013). Physical Internet Foundations. In T. Borangiu, A. Thomas, & D. Trentesaux (Eds.), *Service Orientation in Holonic and Multi Agent Manufacturing and Robotics* (Vol. 472, pp. 151–166). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-35852-4_10
- Ouyang, Z., Yuan, Z., Chen, Q., Li, M., Zhao, Z., & Huang, G. Q. (2025). Shipment Scheduling and Routing Protocols in Cyber-Physical Internet for Prefabricated Construction Modules Logistics. *IEEE Transactions on Intelligent Transportation Systems*, 1–15. https://doi.org/10.1109/TITS.2025.3581569
- Qiu, T., Chi, J., Zhou, X., Ning, Z., Atiquzzaman, M., & Wu, D. O. (2020). Edge Computing in Industrial Internet of Things: Architecture, Advances and Challenges. *IEEE Communications Surveys & Tutorials*, *22*(4), 2462–2488. https://doi.org/10.1109/COMST.2020.3009103
- Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge Computing: Vision and Challenges. *IEEE Internet of Things Journal*, *3*(5), 637–646. https://doi.org/10.1109/JIOT.2016.2579198
- Tao, F., Qi, Q., Wang, L., & Nee, A. Y. C. (2019). Digital Twins and Cyber–Physical Systems toward Smart Manufacturing and Industry 4.0: Correlation and Comparison. *Engineering*, *5*(4), 653–661. https://doi.org/10.1016/j.eng.2019.01.014
- Treiblmaier, H., Mirkovski, K., Lowry, P. B., & Zacharia, Z. G. (2020). The physical internet as a new supply chain paradigm: A systematic literature review and a comprehensive framework. *The International Journal of Logistics Management*, *31*(2), 239–287. https://doi.org/10.1108/IJLM-11-2018-0284
- Wu, H., Huang, L., Li, M., & Huang, G. Q. (2024). Cyber-Physical Internet (CPI)-enabled logistics infrastructure integration framework in the greater bay area. *Advanced Engineering Informatics*, *60*, 102551. https://doi.org/10.1016/j.aei.2024.102551
- Wu, H., Li, M., Yu, C., Ouyang, Z., Lai, K., Zhao, Z., Pan, S., Wang, S., Zhong, R. Y., Kuo, Y.-H., Zhang, F., Huang, W., Shen, Z.-J. M., Ballot, E., & Huang, G. Q. (2025). Towards cyber-physical internet: A systematic review, fundamental model and future perspectives. *Transportation Research Part E: Logistics and Transportation Review*, *197*, 104051. https://doi.org/10.1016/j.tre.2025.104051

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE

Operationalising Physical Internet Principles in Neighbourhood Urban Logistics: The IKIGAI Hybrid Fixed–Mobile Micro-Hub Living Lab

Stefania Pesavento[1]; and Dimos Toulomidis[2]; Sofoklis Dais[2]

1. PNO Innovation, Valencia, Spain
2. CERTH/HIT, Thessaloniki, Greece
Corresponding author: stefania.pesavento@pnoinnovation.com

Abstract: *The operationalisation of Physical Internet (PI) principles in dense urban environments remains limited, particularly at neighbourhood scale where fragmented retail systems and constrained infrastructure hinder efficient logistics consolidation. This paper presents the design and evaluation of a hybrid fixed–mobile micro-hub system implemented in the Can Clota neighbourhood of Esplugues de Llobregat (Barcelona Metropolitan Area) within the IKIGAI project. The proposed system integrates local markets, a temporary consolidation hub, and an electric van (e-van) acting as a mobile PI node to support three delivery modes: smart lockers, mobile hub collection, and door-to-door delivery. A scenario-based methodology is applied to compare the hybrid configuration with fragmented and fixed-only logistics systems under consistent demand assumptions. Key performance indicators include vehicle kilometres travelled, energy consumption, CO_2 emissions, and service accessibility. The results demonstrate the potential of mobile micro-hubs to enhance consolidation efficiency, reduce environmental impact, and improve service flexibility. The study contributes a replicable framework for implementing interoperable, low-carbon urban logistics systems aligned with Physical Internet principles.*

Keywords: *Physical Internet, Urban Logistics, Micro-Hubs, Mobile Logistics Hubs, Smart Contracts, Synchromodal Transport, Last-Mile Delivery*

Physical Internet (PI) Roadmap Fitness: *Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),☐ Transportation Equipment, ☒ PI Networks, ☒ System of Logistics Networks, ☐ Vertical Supply Consolidation, ☐ Horizontal Supply Chain Alignment, ☒ Logistics/Commercial Data Platform, ☒ Access and Adoption, ☒ Governance.*

Targeted Delivery Mode-s: *☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

1 Introduction

Urban logistics systems are undergoing increasing pressure to improve efficiency, reduce environmental impact, and support the growing demand for local and e-commerce-based distribution. In dense metropolitan environments, last-mile delivery operations are often characterised by fragmented flows, low vehicle utilisation, and a high number of delivery trips, resulting in increased congestion and emissions. These inefficiencies are particularly evident in neighbourhood-scale logistics systems, where dispersed demand, limited infrastructure, and heterogeneous stakeholders hinder effective consolidation and coordination.

The Physical Internet (PI) paradigm has emerged as a transformative approach to addressing these challenges by promoting the interconnection and interoperability of logistics networks

through standardised, modular, and open systems (Montreuil, 2011; Montreuil, 2013). By enabling shared use of logistics assets, dynamic routing, and flow consolidation across networks, PI aims to significantly improve the economic, environmental, and operational performance of supply chains. However, despite its strong conceptual foundation, the operationalisation of PI principles in real-world urban contexts remains limited, particularly at the neighbourhood scale where infrastructure constraints and governance complexity are most pronounced. One promising approach to implementing PI principles in urban environments is the use of micro-hubs and decentralised logistics nodes. Fixed micro-hubs and locker systems have been widely studied as mechanisms for consolidating deliveries and reducing last-mile inefficiencies. However, such fixed infrastructure may lack flexibility in responding to spatially and temporally dynamic demand patterns, especially in areas where demand density varies or where permanent infrastructure deployment is constrained. In this context, mobile logistics assets, such as electric vans acting as temporary or circulating hubs, offer a potential extension of the PI paradigm by enabling dynamic, adaptive consolidation and distribution.

This paper investigates the integration of fixed and mobile logistics nodes within a hybrid system designed to operationalise PI principles on the neighbourhood scale. The study is developed in the context of the IKIGAI project, which aims to deploy smart, decentralised orchestration mechanisms for sustainable urban ecosystems. The empirical setting is the Can Clota neighbourhood in Esplugues de Llobregat (Barcelona Metropolitan Area), where a pilot is being designed to transform the existing local commerce platform into an interoperable, data-driven logistics network. The proposed system integrates local markets as supply nodes, a temporary consolidation hub, and an electric van (e-van) operating as a mobile Physical Internet node. The e-van performs multiple daily routes, collecting goods from local markets and distributing them within the neighbourhood through a combination of delivery modes: (i) fixed smart lockers, (ii) mobile hub operations where the vehicle acts as a temporary collection point, and (iii) optional door-to-door delivery. This hybrid configuration enables the combination of infrastructure-based and service-based logistics solutions, supporting both consolidation efficiency and service accessibility. In doing so, it extends traditional micro-hub concepts by introducing mobility and flexibility into the network structure.

The research adopts a scenario-based methodological approach to evaluate the performance of the proposed system. Three configurations are considered: (i) a fragmented baseline scenario characterised by independent delivery flows, (ii) a fixed infrastructure scenario relying on hubs and lockers, and (iii) the proposed hybrid system integrating a mobile micro-hub. These scenarios are analysed under consistent demand assumptions to isolate the impact of consolidation strategies and network design. Performance is assessed using key indicators including vehicle kilometres travelled, energy consumption, CO_2 emissions, and service accessibility.

The contribution of this paper is threefold. Firstly, it proposes a hybrid fixed–mobile micro-hub architecture that extends Physical Internet concepts to neighbourhood-scale urban logistics. Secondly, it provides a structured evaluation framework for comparing static and dynamic consolidation strategies in urban contexts. Thirdly, it offers insights into the role of mobile logistics nodes and multimodal delivery configurations in improving both environmental and operational performance. The remainder of the paper is structured as follows. Section 2 presents the system architecture and operational design of the proposed logistics network. Section 3 describes the scenario definition and modelling assumptions. Section 4 outlines the methodological approach and performance indicators. Section 5 presents and discusses the results. Finally, Section 6 concludes the paper and highlights implications for future research and urban logistics deployment.

2 System Architecture and Operational Design

2.1 System Overview

The proposed system is designed as a hybrid Physical Internet (PI)-enabled urban logistics network integrating multiple supply sources, a flexible consolidation process, and a mobile distribution layer. The system serves the Can Clota neighbourhood in Esplugues de Llobregat (Barcelona Metropolitan Area) and is developed within the IKIGAI project.

Goods originate from local markets located outside the Can Clota area, including Can Vidalet Market and La Plana Market, which act as primary supply nodes. In addition, the system is designed to integrate additional retail and service sectors, such as pharmacies and selected non-food retailers, as well as fresh food products including perishable items (e.g. fish). This multi-sector integration enables the aggregation of heterogeneous demand streams within a unified logistics network, while introducing differentiated handling requirements for time-sensitive goods. Customers can place orders through multiple channels, including the JoComproEsplugues platform, WhatsApp-based communication with retailers, or directly at the market. All orders are registered and integrated into the system to ensure coordination and traceability across actors. Order consolidation is performed by CityLogin at a designated collection point located outside Can Clota, potentially in proximity to the markets or municipal facilities. This consolidation point is not fixed and may vary depending on operational and spatial constraints. Operationally, orders are prepared as small, individually traceable load units, each with a unique ID and handling-status updates, enabling transfer between the consolidation point, the mobile node and the final delivery interface

Last-mile distribution is carried out using an electric van (e-van), which operates as a mobile logistics node transporting consolidated flows from external supply points into the Can Clota neighbourhood. Through this configuration, the system connects upstream supply networks with neighbourhood-level distribution, supporting flow consolidation, interoperability, and flexible service provision in line with Physical Internet principles.

To enable horizontal cooperation without exposing commercially sensitive information, data exchange is managed through the FACTUAL-based local logistics data space, which shares only the operational data needed for orchestration and traceability, while customer, pricing and full order-history data remain under each actor's control and municipal actors access only aggregated KPIs..

2.2 Hybrid Node Configuration

In this study, PI operationalisation is understood in a partial but practical sense. The proposed system does not yet implement the full Physical Internet vision of universal modular containerisation and open inter-network routing. Instead, it operationalises four core PI-compatible features at neighbourhood scale: (i) shared logistics nodes, (ii) coordinated consolidation across multiple supply actors, (iii) interoperable information exchange, and (iv) adaptive allocation of delivery modes across a common service network. As the pilot matures, these features could be extended through standardised parcel units, handling protocols and wider horizontal network integration.

The system architecture is based on a hybrid configuration of Physical Internet nodes, combining multiple types of logistics assets across different spatial layers. The configuration includes:

1. Supply nodes: local markets and additional retail and service providers (e.g. pharmacies and selected non-food retailers) located outside the Can Clota service area
2. Flexible consolidation node: a temporary or variable collection point outside Can Clota where orders from multiple sources are aggregated
3. Mobile node: an electric van (e-van) acting as a circulating micro-hub
4. Delivery nodes: smart lockers and temporary mobile hubs located within the Can Clota neighbourhood

This configuration establishes a clear separation between upstream supply and downstream distribution, while ensuring connectivity between them through the mobile node. The e-van is conceptualised as a mobile Physical Internet node, enabling the transfer of consolidated flows from external supply networks into the neighbourhood. Unlike fixed micro-hubs, the mobile node introduces spatial and temporal flexibility, allowing the system to adapt to demand variability and operational constraints.

Within the service area, smart lockers function as fixed delivery nodes, enabling asynchronous parcel collection and reducing delivery failures. In parallel, the system incorporates mobile hub operations, whereby the e-van temporarily stops at predefined locations to serve as a dynamic pickup point for users.

The hybrid node configuration also supports the integration of heterogeneous demand streams with different service requirements. In particular, perishable goods such as fresh food require shorter delivery lead times and are therefore prioritised for direct delivery or mobile hub collection, while non-perishable items can be efficiently allocated to locker-based delivery.

This differentiation enables the system to balance consolidation efficiency with service quality. The combination of lockers and mobile hub operations allows deliveries to be adapted to different user needs and product types within a single system. This reduces the need for multiple delivery trips while maintaining service flexibility, particularly in a neighbourhood context where fixed infrastructure is limited.

3 Scenario Definition

A scenario-based approach is adopted to evaluate the performance of alternative logistics system configurations under consistent demand conditions. This approach enables the isolation of the impact of consolidation strategies, node configurations, and delivery modalities on operational and environmental performance.

The analysis focuses on the Can Clota neighbourhood as the service area, with supply nodes located outside the area. Demand is assumed to be identical across all scenarios in terms of volume, spatial distribution, and composition. It includes heterogeneous product categories such as food, pharmaceutical goods, and retail items, with differentiated service requirements for perishable products. Three scenarios are defined, representing increasing levels of coordination and system integration.

3.1 Scenario 0: Fragmented Delivery System

The baseline scenario represents a fragmented logistics system in which deliveries are performed independently by multiple supply actors, with limited coordination and low levels of consolidation. Each retailer or supplier organises its own delivery or customer collection process without integration into a shared logistics network. In this configuration, goods are transported directly from supply nodes to end users through individual delivery flows or made available for collection at existing pickup points. These include fixed locker or collection facilities located at the library and at local markets. However, the use of these facilities is not coordinated across actors and does not follow an optimised consolidation strategy. As a result, the system is characterised by:

- multiple uncoordinated delivery trips entering the service area
- low vehicle utilisation and limited load consolidation
- dispersed collection patterns at existing pickup points
- lack of integration between delivery and collection operations

Delivery and collection activities occur in parallel, with no centralised routing, shared infrastructure management, or coordinated allocation of delivery modes. This results in

increased vehicle kilometres travelled and reduced operational efficiency compared to more integrated system configurations.

3.2 Scenario 1: Fixed Infrastructure System

The second scenario introduces a coordinated logistics system based on fixed infrastructure. In this configuration, goods from multiple supply nodes are aggregated prior to entering the service area and distributed through a structured network of fixed delivery points, such as smart lockers.

Unlike the baseline scenario, where pickup points are used in a fragmented and uncoordinated manner, this configuration assumes centralised planning and consolidation of flows. Deliveries are organised to optimise routing and reduce the number of vehicle trips required to serve the area. In this scenario, goods are transported from external supply nodes to designated locker locations, including both existing facilities (e.g. library and market lockers) and additional lockers within the Can Clota neighbourhood. End users collect their orders from these fixed points, eliminating the need for direct delivery in most cases.

Key characteristics of this configuration include: (1) coordinated consolidation of flows prior to last-mile distribution; (2) optimised routing of delivery trips to fixed locker locations; (3) reduced number of vehicles entering the service area compared to Scenario 0; (4) reliance on static infrastructure for final delivery.

While this approach improves operational efficiency compared to the fragmented baseline, it remains limited in its ability to adapt to demand variability, user preferences, and product-specific requirements. In particular, the exclusive reliance on fixed delivery points reduces flexibility for time-sensitive or perishable goods and may constrain service accessibility for certain user groups.

3.3 Scenario 2: Hybrid Fixed–Mobile System

The third scenario represents the proposed hybrid logistics system integrating both fixed and mobile nodes. In this configuration, goods from multiple supply sources are consolidated outside the Can Clota service area and transported into the neighbourhood using an electric van (e-van) acting as a mobile Physical Internet node.

The system combines centralised consolidation with dynamic last-mile distribution. The e-van performs multiple daily routes, collecting aggregated orders from the consolidation point and distributing them within the service area through a multi-stop routing structure.

Unlike Scenario 1, which relies exclusively on fixed delivery infrastructure, this configuration supports multiple delivery modalities within a unified operational framework:

1. Locker delivery, enabling asynchronous collection at fixed points
2. Mobile hub collection, where the e-van temporarily stops at predefined locations, allowing users to collect orders within a defined time window
3. Door-to-door delivery, provided for selected use cases

This hybrid approach enables dynamic allocation of delivery modes based on demand characteristics, user preferences, and product requirements. In particular, time-sensitive and perishable goods, such as fresh food products, are prioritised for rapid delivery through mobile hub or door-to-door services, reducing delivery lead times and avoiding prolonged storage.

The integration of a mobile node introduces additional flexibility compared to fixed-only systems. The e-van enables the redistribution of consolidated flows within the neighbourhood without requiring additional permanent infrastructure, while also adapting to spatial and temporal variations in demand.

Key characteristics of this scenario include: (1) consolidation of flows outside the service area; (2) use of a mobile node to distribute goods within the neighbourhood; (3) integration of multiple delivery modalities within a single system; (4) adaptive allocation of delivery modes based on product and user requirements; (5) improved balance between operational efficiency and service accessibility. Compared to the previous scenarios, the hybrid system is expected to achieve higher consolidation efficiency, reduce vehicle kilometres travelled, and provide greater flexibility in meeting diverse user and product needs.

4 Methodology

4.1 Modelling approach

A comparative scenario-based modelling approach was adopted to assess the operational and environmental performance of alternative neighbourhood logistics configurations. The purpose of the analysis was to isolate the effect of consolidation logic, node configuration and delivery mode allocation under a common demand structure. Accordingly, the same order volume, spatial distribution and product mix were maintained across all scenarios, whereas only the organisation of the logistics system was allowed to vary.

The analysis represents the Can Clota case as a local distribution network linking external supply points, an intermediate consolidation stage and final delivery or collection within the neighbourhood. The modelling logic follows the operational sequence of the proposed service, from order aggregation and consolidation to neighbourhood distribution and final receipt by the user. Three alternative system configurations were examined: a fragmented delivery structure, a coordinated fixed-infrastructure structure and a hybrid fixed–mobile structure centred on an electric van acting as a mobile micro-hub.

The methodological focus was placed on the logistics performance of these alternative configurations rather than on detailed behavioural forecasting. For this reason, the model was designed as an ex-ante comparative framework that estimates relative differences in transport effort, energy use, emissions and accessibility across scenarios. This approach is appropriate for the present study because the objective is to evaluate the contribution of the hybrid architecture to the operationalisation of PI principles at neighbourhood scale.

4.2 System Boundaries and Assumptions

The system boundary was defined around the last-mile and neighbourhood distribution segment. Upstream procurement and long-distance freight transport were excluded, since the study aims to compare alternative urban distribution structures rather than the full supply chain. The spatial boundary includes the Can Clota neighbourhood as the service area and the external interface through which consolidated orders enter the area from participating markets and retailers.

A representative operating period was assumed for all scenarios. Demand was held constant in terms of the number of orders, their spatial allocation and their composition. The demand mix includes food products, pharmaceutical goods and selected retail items. Perishable products were assumed to require shorter lead times and reduced storage duration within the system; this affects the allocation of orders to delivery modes in the hybrid configuration.

Several simplifying assumptions were adopted to ensure comparability. All delivery nodes were assumed to be operational and accessible during the analysis period. Failed deliveries, missed pickups, returns and reverse logistics were not quantified in the ex-ante comparison; operationally, parcels not collected within the defined time window are flagged in the coordination platform and can be reassigned to a later mobile pickup slot, redirected to a locker or other collection point, or returned to the consolidation point, while failed home deliveries generate a proof-of-attempt and a follow-up delivery action. Locker deliveries were assumed

to result in successful user collection. Travel distances were measured along the road network for vehicle movements, whereas service accessibility was assessed through the distance or availability of collection and delivery options from the user perspective. In the hybrid scenario, one electric van was assumed to perform the circulating distribution function within the neighbourhood.

As the pilot is in a pre-operational phase, the model relies on scenario-based assumptions derived from the planned system configuration and urban context. These assumptions are used consistently across scenarios to ensure comparability rather than to predict exact operational outcome.

4.3 Scenario Parameterisation

The scenarios were parameterised using the same demand base and the same service area, while varying the degree of coordination, the type of delivery nodes used and the routing structure.
In Scenario 0, the fragmented baseline, deliveries are carried out independently by supply actors with limited or no coordination. Vehicle trips are generated separately and goods either reach end users directly or are made available through existing pickup points without integrated planning. This scenario represents low consolidation, duplicated route structures and limited interoperability between actors.

In Scenario 1, goods from multiple supply sources are first aggregated and then distributed through fixed delivery infrastructure, primarily smart lockers. Routing is therefore organised around consolidated trips to designated locker locations. Compared with the baseline, this configuration reduces duplicated entry into the service area and improves vehicle utilisation; however, it remains dependent on static infrastructure and provides limited flexibility for time-sensitive goods and differentiated service needs.

In Scenario 2, the proposed hybrid fixed–mobile system, consolidated orders, are transported into the neighbourhood by an electric van operating as a mobile micro-hub. Orders are then allocated to one of three delivery modes: locker delivery, mobile hub collection and door-to-door delivery. The allocation is rule-based and reflects product and service characteristics. Non-perishable goods are considered suitable for locker delivery, whereas perishable and time-sensitive goods are prioritised for mobile hub collection or direct delivery. The key variable in this scenario is therefore not only the routing pattern but also the dynamic distribution of orders across delivery modes.

This parameterisation makes it possible to evaluate how progressively higher levels of coordination and flexibility affect system performance. The comparison is therefore not limited to vehicle movement alone, but also captures the interaction between node type, service logic and accessibility.

4.4 Performance Indicators (KPIs)

The comparative assessment was conducted using four core key performance indicators: vehicle kilometres travelled, energy consumption, CO_2 emissions and service accessibility. These indicators were selected because they capture the principal dimensions through which the proposed system is expected to generate value, namely transport efficiency, environmental performance and user-oriented service provision.

Vehicle kilometres travelled represents the total transport effort required to serve the demand under each scenario and provide the main indicator of consolidation efficiency. Energy consumption captures the operational energy required by the delivery system and enables comparison between fragmented and coordinated configurations. CO_2 emissions translate transport activity into environmental impact and allow the benefit of consolidation and electrified neighbourhood distribution to be assessed. Service accessibility reflects the extent

to which the system remains convenient for users while seeking efficiency gains through consolidation.

In the fixed-infrastructure scenario, accessibility depends primarily on the proximity of users to locker locations. In the hybrid scenario, accessibility also reflects the availability of temporary mobile collection points and the option of direct delivery for selected orders. The accessibility indicator therefore captures not only physical distance, but also the capacity of the system to accommodate heterogeneous delivery requirements within a unified network.

Service accessibility is quantified as the percentage of orders served either through door-to-door delivery or through a collection option meeting the predefined access criterion.

4.5 Calculation methods

For each scenario s, vehicle kilometres travelled were calculated as the sum of the distances of all delivery tours required to serve the complete order set:

$$VKT_s = \sum_{r=1}^{R_s} d_{r,s} \quad (1)$$

where $d_{r,s}$ denotes the network distance of route r in scenario s, and R_s is the total number of routes performed in that scenario.

Energy consumption was estimated by applying a vehicle-specific energy factor to the total travelled distance:

$$EC_s = \sum_{v=1}^{V_s} VK\,T_{v,s} \cdot \alpha_v \quad (2)$$

where $VKT_{v,s}$ is the distance travelled by vehicle type v in scenario s, and α_v is the corresponding energy consumption coefficient per kilometre. In the hybrid configuration, the principal vehicle is an electric van; in the fragmented and fixed scenarios, the same formulation allows the use of one or more conventional vehicle types where appropriate.

CO_2 emissions were then derived from operational energy use or distance-based emission factors:

$$CO2_s = \sum_{v=1}^{V_s} E\,C_{v,s} \cdot \gamma_v \quad (3)$$

where γ_v denotes the emission factor associated with vehicle type v or its energy carrier. This formulation allows a consistent comparison across scenarios even when the vehicle mix differs.

Service accessibility was calculated as the share of orders served under an acceptable access condition:

$$SA_s = \frac{\sum_{i=1}^{N_s} a_{i,s}}{N_s} \quad (4)$$

where N_s is the total number of orders in scenario s, and $a_{i,s}$ takes the value 1 when order $i$is served either through door-to-door delivery or through a collection point that satisfies the predefined accessibility criterion, and 0 otherwise. This indicator enables the analysis to capture the balance between consolidation and user convenience.

4.6 Comparative Evaluation Framework

The evaluation framework compares the three scenarios under identical demand conditions and interprets performance differences as the result of alternative logistics architectures. Scenario 0 serves as the reference case, because it represents the fragmented structure from which coordination gains can be measured. Scenario 1 captures the effect of fixed-point consolidation,

whereas Scenario 2 captures the additional effect of combining fixed nodes with a mobile micro-hub and differentiated delivery modes.

For each indicator, the scenarios were compared in terms of their expected relative performance against the fragmented baseline. At this stage, absolute values and percentage differences are not reported, since the pilot is pre-operational and the required observed data on route execution, demand allocation, vehicle utilisation and user collection behaviour are not yet available. This allows the results to be interpreted not only in terms of performance improvement, but also in terms of trade-offs between efficiency, emissions and service provision. In particular, the framework makes it possible to assess whether the hybrid system achieves simultaneous gains across operational, environmental and accessibility dimensions, or whether improvements in one dimension are associated with losses in another.

From a Physical Internet perspective, the comparison also reflects increasing levels of operational integration. The transition from Scenario 0 to Scenario 1 represents the effect of coordinated flow consolidation and shared fixed delivery infrastructure, whereas the transition from Scenario 1 to Scenario 2 represents the introduction of dynamic node behaviour, flexible service allocation and mobile network functionality. The framework therefore evaluates not only alternative delivery systems, but also the extent to which Physical Internet principles become operational at neighbourhood scale through the combination of shared assets, interoperable nodes and adaptive distribution logic.

5 Results and Discussion

The comparative analysis highlights clear differences in operational performance across the three scenarios, reflecting increasing levels of coordination and system integration. Table 1 summarises the relative performance of each configuration with respect to the selected key performance indicators.

Table 1: Relative performance across scenarios (S0 = 100; lower is better except accessibility)

Scenario	VKT Index	Energy Index	CO_2 Index	Accessibility (% o orders meeting access criterion)
S0: Fragmented	High	High	High	Medium
S1: Fixed Infrastructure	Medium	Medium	Medium	Low-Medium
S2: Hybrid Fixed–Mobile	Low	Low	Low	High

The fragmented baseline (Scenario 0) shows the lowest efficiency, with uncoordinated delivery flows resulting in high vehicle kilometres travelled, energy consumption, and emissions. Although existing pickup points are available, their uncoordinated use does not enable effective consolidation. The fixed infrastructure scenario (Scenario 1) improves efficiency through coordinated consolidation and routing to locker locations, reducing the number of delivery trips. However, its reliance on static infrastructure limits flexibility, particularly for time-sensitive deliveries and users located further from fixed collection points. The hybrid fixed–mobile system (Scenario 2) achieves the best overall performance by combining consolidation with flexible distribution. The use of the e-van as a mobile node enables dynamic allocation of delivery modes, allowing perishable goods to be prioritised for rapid delivery while maintaining efficient consolidation for other flows. Across scenarios, the hybrid configuration reduces

transport effort while improving service accessibility. From a Physical Internet perspective, it represents the transition from fragmented logistics to an adaptive, interoperable network combining fixed and mobile nodes.

6 Conclusion

This paper evaluated three alternative logistics configurations for neighbourhood-scale urban distribution, focusing on the role of hybrid fixed–mobile systems in operationalising selected Physical Internet principles in a neighbourhood-scale urban logistics context. The findings indicate that combining external consolidation with mobile distribution enables a more efficient and flexible logistics system compared to both fragmented and fixed-only configurations.

In particular, the integration of a mobile node allows delivery modes to be dynamically adapted to demand characteristics, improving consolidation efficiency while maintaining service accessibility, including for time-sensitive and perishable goods. These results highlight the potential of hybrid node configurations to bridge external supply networks and local distribution systems within dense urban environments. The study contributes a structured framework for analysing the impact of node configuration and delivery mode integration on urban logistics performance. From a Physical Internet perspective, it demonstrates how modular and interoperable logistics systems can be implemented at neighbourhood scale.

The analysis is based on an ex-ante modelling approach using assumed parameters, as the pilot is in a pre-operational phase. Future work will focus on validating the results through real pilot data and assessing scalability to other urban contexts within the Barcelona Metropolitan Area. In this respect, the results should be interpreted as a qualitative ex-ante comparison rather than as a quantified impact assessment. Since the pilot is still in a pre-operational phase, validated operational data on route distances, order volumes, vehicle utilisation, delivery-mode allocation and user collection behaviour are not yet available. Numerical estimates of VKT reduction, energy savings, CO_2 reduction and accessibility improvement will therefore be produced once living lab data become available.

References

- Montreuil, B. (2011). Towards a Physical Internet: Meeting the Global Logistics Sustainability Grand Challenge. Logistics Research, 3(2–3), 71–87. https://doi.org/10.1007/s12159-011-0045-x
- Montreuil, B. (2013). Physical Internet Manifesto (v1.10). http://www.physicalinternetinitiative.org
- Landschützer, C., Ehrentraut, F., Jodin, D. (2015). Containers for the Physical Internet: requirements and engineering design related to FMCG logistics. Logistics Research, 8, article 8. https://doi.org/10.1007/s12159-015-0126-3
- Buldeo Rai, H., Kang, S., Sakai, T., Tejada, C., Yuan, Q., Conway, A., Dablanc, L. (2022). 'Proximity logistics': Characterizing the development of logistics facilities in dense, mixed-use urban areas around the world. Transportation Research Part A: Policy and Practice, 166, 41–61. https://doi.org/10.1016/j.tra.2022.10.007

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Rethinking Blockchain Consensus Mechanisms: A Supply Chain Perspective through the Physical Internet

Imene Yasmine Benyelles[1], Sonia Mahjoub[2] and Hicham Abbad [1]
1. Nantes University, LEMNA, Nantes, France
2. ONIRIS, LEMNA, Nantes, France
Corresponding author: imeneyasmine.benyelles@univ-nantes.fr

Abstract: *This research paper proposes a method of achieving consensus for blockchain-enabled Supply Chain transactions that does not rely on computationally intensive work, stake-weighted participation, or pre-designated authority. Instead, this method views the process of establishing a consensus mechanism using the Physical Internet (PI) as an operational coordination issue within a network of decentralized and connected PI-Hubs. Each validator corresponds to a logistics hub controlled by a supply chain actor, and each blockchain transaction represents a standardized logistics event embedded in an event chain, including events related to π-container movement across stakeholder hubs. A bi-objective optimization model is formulated to minimize confirmation latency and validation overhead while enforcing validator eligibility, π-interface protocol compatibility, hub capacities, and strict precedence. The problem is solved using NSGA-II and evaluated through Pareto-front indicators, including hypervolume and spacing. The work delivers three contributions: a mathematical model, a Physical Internet-grounded optimization approach, and a consensus design framework that connects blockchain validation, hyperconnected logistics, and multi-objective decision-making.*

Keywords: *Physical Internet, Blockchain, Consensus Mechanism, Multi-Objective Optimization, NSGA-II, Supply Chain.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☒ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☒ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☒ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☒ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

1 Introduction

Supply chains are undergoing a major transformation as the volume of available information, the level of inter-organizational coordination, and the need for trusted traceability increase. In this context, blockchain has emerged as a promising infrastructure for creating transparency, authenticity, decentralization, and immutability between supply chain partners (Azzi et al., 2019; Joo & Han, 2021; Sarfaraz et al., 2023b). Meanwhile, the Physical Internet proposes a complementary logistics vision based on open, standardized, and hyperconnected networks, rather than fragmented transportation and warehousing operations (Montreuil, 2011; Montreuil et al., 2012).

The intersection of these two ideas is particularly relevant for supply chain applications, since logistics events are both physical and informational. They correspond to physical operations, such as the movement of goods through warehouses, vehicles, and handling systems, while also generating information about product custody, status, and quality that must be verified and shared across multiple actors. In this environment, blockchain supports the creation of a trusted and immutable network. However, blockchain performance depends strongly on the design of

the consensus algorithm, particularly in terms of latency, throughput, and coordination costs (Mahjoub & Abbad, 2025; Zhang et al., 2021).

The recent literature on supply chain consensus has started to move from generic mechanisms toward approaches involving reputation-based consensus (de Oliveira et al., 2020; Rahman et al., 2025; Sarfaraz et al., 2023a) or fairness-based consensus (Mahjoub & Abbad, 2025; Zhu et al., 2024). This shift shows the need for validation rules that are specific to supply chain operations and adapted to their operational contexts, rather than mechanically copied from cryptocurrency-oriented consensus mechanisms (Zheng et al., 2017). However, most existing studies still represent validators as abstract blockchain nodes, with limited emphasis on physical logistics structures, logistics hub interoperability, and event-chain feasibility (Rahman et al., 2025; Sarfaraz et al., 2023b).

To address this gap, this research develops a Physical Internet-based consensus mechanism framework called PI-Consensus. Within the PI-Consensus framework, blockchain validators are represented as PI-hubs controlled by supply chain actors, while transactions are represented as logistics events embedded in a multi-stage event chain. Unlike conventional consensus rules based on computational effort, stake, or predefined authority, PI-Consensus allocates validation responsibility according to operational suitability within the PI network, namely π-interface compatibility, hub capacity, and distributed participation requirements.

This paper makes three main contributions. First, it reframes blockchain consensus as a PI-native logistics coordination problem, where validators are PI-hubs and transactions are ordered logistics events linked to π-container flows. Second, it formulates validator selection as a bi-objective optimization model integrating latency, validation overhead, eligibility, capacity, precedence, π-interface compatibility, and distributed participation. Third, it evaluates PI-Consensus using NSGA-II, Pareto-front indicators, sensitivity analysis, and hub-event mapping, showing that decentralized validator participation is preserved as event-chain workloads increase.

The remainder of this paper is organized as follows. Section 2 reviews the literature on blockchain for supply chain coordination and on the Physical Internet as a logistics paradigm for interoperable and hyperconnected networks. Section 3 presents the problem description, modeling assumptions, and the bi-objective mathematical formulation of PI-Consensus. Section 4 describes the NSGA-II solution approach. Section 5 presents the experimental setting and performance indicators. Section 6 reports and discusses the numerical results. Finally, Section 7 concludes the paper and outlines future research directions.

2 Literature Review

2.1 Blockchain for Supply Chain coordination

A large number of research papers have investigated Blockchain as a solution for coordination and tracking of goods through Supply Chain systems (Saberi et al., 2019). Blockchain can be described as a decentralized and distributed ledger maintained by a peer-to-peer network of nodes that collectively share, verify and update transactions stored in cryptographically linked blocks. This structure makes the ledger transparent, tamper-resistant and independent of central authority (El Amine & Lamia, 2023; Riouali & Rabhi, 2026). Within such a system, transactions represent verifiable records of exchanges or operational events, while validators are the nodes responsible for confirming them (Wüst & Gervais, 2018). Consensus, in turn, is the method or the process by which participants in the network reach an agreement on what is currently considered to be the valid state of the network's ledger, what order they will be added in, and which validating authority or authorities will be allowed to append a new block (Jain et al., 2025; Zheng et al., 2017).

Recent work on Supply Chain Blockchain has therefore started to question the suitability of conventional consensus and to propose lighter or more specialized mechanisms. Rahman et al. (2025) developed a lightweight practical consensus mechanism for Blockchain, while Mahjoub and Abbad (2025) formulated a fair-oriented consensus framework based on reinforcement learning. Yet even in these promising studies, the validator is still not modeled as a PI-grounded logistics resource whose role depends on interoperability and networked operations.

2.2 Physical Internet

According to the Physical Internet, logistics systems should evolve from being closed and fragmented to become an open and hyper-connected network (Louzim et al., 2024). PI achieves this goal through standardized encapsulation, interoperability protocols, and interconnected logistics nodes able to collaborate across organizational boundaries (Montreuil, 2011; Montreuil et al., 2012). This work has been expanded since then to include logistics collaboration, coordination platforms and the need for digital enablers that provide trust and visibility in open networks (Bekrar et al., 2021; Meyer et al., 2019).

From this perspective, a digital trust layer is essential for open logistics networks since there needs to be reliable validation of events across many different organizations that do not have a single central authority. Many studies have therefore suggested that Blockchain can be used as an enabling layer for PI collaboration by supporting secure exchanges of information as well as smart contract coordination (Betti et al., 2019; Chargui et al., 2024; Hasan et al., 2021; Meyer et al., 2019; Pan, 2024). Yet, the consensus layer remains underdeveloped from the perspective of PI. The literature provides evidence that blockchain can support logistics collaboration and event traceability. However, little attention has been paid to the design of consensus mechanisms in which validator selection is derived from the operational structure of PI networks. In particular, existing studies do not explicitly model validators as PI-hubs constrained by π-interface compatibility, hub capacity, network reachability, and ordered π-container event chains.

3 Problem description and mathematical model

3.1 General description

This work proposes a consensus design model for a blockchain-enabled Physical Internet supply chain, in which PI-hubs act as distributed validators and blockchain transactions are represented as logistics events embedded in dynamically routed multi-stage supply chain sequences (Figure 1). The problem consists of assigning each event to an eligible hub and a validation period over a fixed planning horizon in order to minimize confirmation latency and validation overhead. This assignment is carried out under operational assumptions related to validator eligibility, hub capacity, event timing, and precedence between logistics events. The objective is therefore to define a consensus mechanism that remains consistent with both supply chain execution and Physical Internet principles.

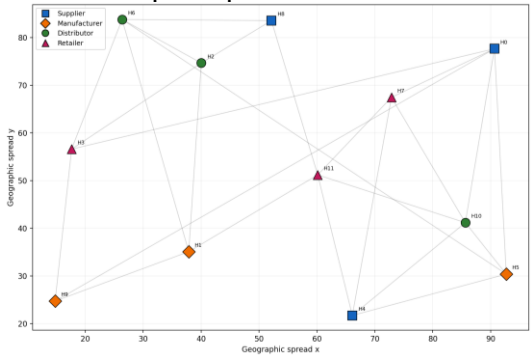


Figure 1 PI-consensus network topology

In this setting, validator selection is not treated as a purely digital mining process. A PI-hub can validate a logistics event only when it is operationally suitable for that event. Suitability depends on the hub's role in the supply chain, its available validation capacity during the considered period, its compatibility with the π-interface protocol required by the event, and its reachability within the PI network. Therefore, the selected validator set reflects both blockchain governance requirements and the physical feasibility of π-container movement across stakeholder hubs.

3.2 Assumptions

The model is developed under the following assumptions:

- Transactions are modeled as discrete logistics events.
- Events are grouped into ordered Supply Chain sequences.
- Each event belongs to one chain and may be subject to precedence constraints.
- Each event is validated exactly once by one eligible PI-hub in one time period.
- PI-hubs act as blockchain validators and correspond to supply chain stakeholder roles.
- A hub may validate an event only if role, protocol, and network compatibility are met.
- The planning horizon is discretized into multiple periods.
- Hub validation capacity is limited in each period.
- The validation of an event consumes hub capacity and generates latency and overhead.
- Precedence relations between events in the same chain must be respected.
- The PI network topology is known and remains unchanged during the horizon.

3.3 Problem Formulation

Based on the above assumptions, the consensus design is formulated as a bi-objective mixed integer optimization model.

- Notations:

Table 1: Notations

Sets/parameters	Description
I	Set of logistics events, indexed by *i*
J	Set of PI-hubs acting as validator nodes, indexed by *j*
T	Set of time periods, indexed by *t*
L	Set of event sequences, indexed by *l*
P	Set of precedence pairs *(i,i')*, where event i must be validated before event *i'*
u_i	Priority weight associated with event *i*
a_{ij}	Effective latency incurred when event *i* is validated by hub *j*
b_{ij}	Validation overhead incurred when event *i* is assigned to hub *j*
C_{jt}	Validation capacity available at hub *j* during period *t*
e_{ij}	Binary eligibility parameter equal to 1 if hub *j* can validate event *i*, and 0 otherwise
g_j	Fixed activation cost of hub *j*
m	Minimum number of hubs that must participate in consensus over the planning horizon

- Decision variable

$$x_{ijt} \begin{cases} 1, if\ event\ i\ is\ validated\ \ by\ hub\ j\ in\ period\ t \\ 0, otherwise \end{cases}$$

$$y_j \begin{cases} 1, if\ hub\ j\ is\ active\ during\ the\ planning\ horizon \\ 0, otherwise \end{cases}$$

Objective function

The problem is formulated as a bi-objective optimization model.

- **Objective 1:** Minimization of weighted validation latency

$$\min F_1 = \sum_{i \in I} \sum_{j \in J} \sum_{t \in T} u_i \, a_{ij} x_{ijt} \quad (1)$$

- **Objective 2:** Minimization of validation overhead

$$\min F_2 = \sum_{i \in I} \sum_{j \in J} \sum_{t \in T} b_{ij} \, x_{ijt} + \sum_{j \in J} g_j \; y_j \quad (2)$$

Constraints

Each event must be validated once:

$$\sum_{j \in J} \sum_{t \in T} x_{ijt} = 1 \qquad \forall i \in I \quad (3)$$

An event can only be assigned to an eligible hub:

$$x_{ijt} \leq e_{ij} \qquad \forall i \in I, \forall j \in J, \forall t \in T \quad (4)$$

The validation workload assigned to a hub in a given period must not exceed its capacity

$$\sum_{i \in I} a_{ij} x_{ijt} \leq C_{jt} \qquad \forall j \in J, \forall t \in T \quad (5)$$

For every precedence pair $(i,i') \in P$, event i' can only be validated after event i:

$$\sum_{j \in J} \sum_{t \in T} t \, x_{i'jt} \geq \sum_{j \in J} \sum_{t \in T} t \, x_{ijt} + 1 \quad \forall (i, i') \in P \quad (6)$$

Where *t* represents the period's index, and (6) ensures the correct temporal order within each sequence.

A hub is considered active if at least one event is assigned to it:

$$x_{ijt} \leq y_j \forall i \in I, \forall j \in J, \forall t \in T \quad (7)$$

To preserve distributed consensus, a minimum number of hubs must participate over the planning horizon:

$$\sum_{j \in J} y_j \geq m \quad (8)$$

Variable domains:

$$x_{ijt} \in \{0,1\} \, \forall i \in I, \forall j \in J, \forall t \in T \quad (9)$$

$$y_j \in \{0,1\} \forall j \in J \quad (10)$$

4 Resolution methodology: NSGA II solution approach

The proposed multi-objective optimization problem is combinatorial, and exact optimization becomes increasingly challenging as the number of hubs, event chains, and time periods increases. To address this issue, NSGA-II (Deb et al., 2002) is used to approximate the Pareto front while preserving solution diversity. Each individual encodes an assignment of logistics events to validator hubs and validation periods. Heuristic and random initialization, followed by selection, crossover, mutation, and repair, are used to generate feasible solutions subject to eligibility, precedence, capacity, and minimum hub participation constraints. The population is updated using non-dominated sorting and crowding distance until a set of Pareto-efficient PI-Consensus configurations is obtained.

4.1 NSGA II solution approach

Algorithm 1: NSGA II for PI-Consensus

Input: I, J, T, P, a_{ij}, b_{ij}, C_{jt}, e_{ij}, g_j, m, NSGA-II parameters

Output: Pareto set of non-dominated PI-Consensus solutions
1: Generate the PI network and logistics event data
2: Compute eligibility, latency, overhead, workload, and capacity parameters
3: Encode each individual as event-to-hub-period assignments (i,j,t)
4: Initialize population using heuristic and random feasible assignments
5: Repair infeasible individuals for eligibility, capacity, precedence, and minimum hub participation
6: Evaluate F1 and F2 for each individual
7: for generation =1 to G do
8: Rank solutions using fast non-dominated sorting
9: Compute crowding distance
10: Select parents using binary tournament selection
11: Apply assignment-based crossover with probability p_c
12: Apply hub-period mutation with probability p_m
13: Repair offspring for eligibility, capacity, precedence, and minimum hub participation
14: Evaluate F1 and F2 for each offspring
15: Merge parent and offspring populations
16: Rank merged population using non-dominated sorting and crowding distance
17: Keep the best N individuals for the next generation
18: end for
19: Return the Pareto-efficient set of feasible PI-Consensus solutions

5 Numerical simulations and performance evaluation

The proposed model was implemented in Python 3.13.5 using the Anaconda distribution and executed on a system equipped with an 11th-generation Intel Core i7-1165G7 processor running at 2.80 GHz with 32 GB of RAM. In the present study, a π-container denotes a standardized modular logistics unit in the Physical Internet, conceptually analogous to a data packet in the digital Internet. Goods are encapsulated into π-boxes, π-bins, or π-pallets and transferred through interoperable PI-hubs using standardized interfaces and protocols. An event sequence represents the complete journey of one π-container through the supply chain, where each event corresponds to a physical logistics operation such as shipment, reception, quality check, consolidation, or deconsolidation. Precedence constraints ensure that blockchain validation follows the same chronological order as the underlying material flow.

5.1 Parameters

The PI network comprises 12 hubs distributed across four supply chain roles, namely Supplier, Manufacturer, Distributor, and Retailer, over 10 time periods. Three instances with 10, 20, and 30 event sequences generate 31, 63, and 94 logistics events, respectively. Eligibility is determined by protocol compatibility, container class, and network reachability, while hub capacity is calibrated for each period with a tight buffer to promote balanced load distribution. NSGA-II is executed with a population size of 140 and 200 generations. Binary tournament selection is used for parent selection, fast non-dominated sorting is used for ranking, and crowding distance is used to preserve diversity. The crossover and mutation probabilities are set to 0.90 and 0.28, respectively. The stopping criterion is the maximum number of generations, and feasibility is handled through a repair procedure. These parameter values were selected based on the sensitivity analysis.

5.2 Performance indicators

Two standard metrics are used to evaluate the Pareto fronts (Collette & Siarry, 2005; Deb et al., 2002). The quality of the Pareto front is evaluated with the hypervolume indicator. In the

bi-objective minimization space (F_1, F_2), Hyper Volume (HV) measures the total area dominated by the non-dominated solutions and bounded by a fixed reference point r=(r_1, r_2). It is defined as

$$HV(P;r) = \lambda\left(\bigcup_{x\in P}[F_1(x), r_1] \times [F_2(x), r_2]\right) \quad (11)$$

where P is the Pareto front and λ(·) denotes the area in the two-dimensional objective space. A larger hypervolume indicates a better Pareto front, as it reflects both better convergence toward low objective values and better spread of trade-off solutions.

The Spacing (SP) measures the standard deviation of nearest-neighbour distances along the front; lower values indicate more uniform distribution. The average number of active hubs is also reported to verify the participation of distributed validators.

$$S = \left(\frac{1}{N-1}\sum_{i=1}^{N}\left(d_i - \bar{d}\right)^2\right)^{\frac{1}{2}} \quad (12)$$

Where N is the number of solutions in the Pareto front, d_i is the distance from solution i to its nearest neighbour in objective space, and $\bar{d}$ is the average of all d_i.

6 Results and discussion

The Pareto fronts (Figure 2) show that, as the number of π-container flows increases, the frontier moves upward and to the right: both latency F1 and consensus overhead F2 worsen because more events compete for limited hub-period capacity.

The small instance yields the richest front because spare capacity gives NSGA-II more freedom to trade speed against resource use. The medium instance captures realistic daily operations and the large instance reflects peak-load conditions where capacity becomes binding and the room for trade-off narrows.

A key result is that the average number of active validator hubs remains above 10 in all cases, showing that the model preserves distributed participation instead of collapsing into a small centralized validator set.

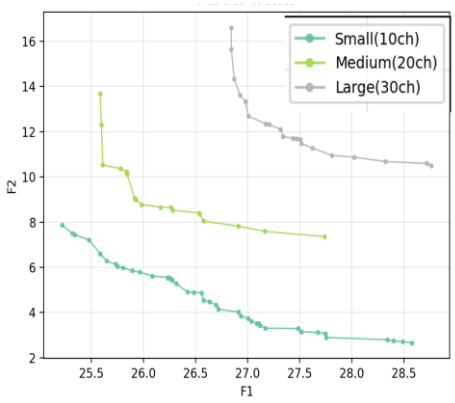


Figure 2 Pareto Front with 3 instances

Operationally, this means PI-Consensus scales coherently: adding more π-container journeys increases cost and coordination effort, but the mechanism still returns structured non-dominated solutions and keeps validator selection aligned with real logistics roles and PI-native constraints.

Table 2: NSGA-II performance across instances

Instance	Events	PF size	Best F1	Best F2	F1 range	F2 range	HV	Spacing	Avg. active hubs	CPU time (s)
Small (10 sequences)	31	42	25.2	2.7	[25.2, 28.6]	[2.7, 7.9]	147.7	0.075	10.6	62.3
Medium (20 sequences)	63	18	25.6	7.4	[25.6, 27.7]	[7.4, 13.7]	88.9	0.398	11.3	186.8
Large (30 sequences)	94	20	26.8	10.5	[26.8, 28.8]	[10.5, 16.6]	45.2	0.271	11.1	247.4

Table 2 shows that PI-Consensus generates feasible and structured Pareto fronts for all instance sizes. As the number of event sequences increases from 10 to 30, the ranges of both F1 and F2 shift toward higher values, which indicates that larger validation workloads make it more difficult to achieve low latency and low overhead simultaneously. The hypervolume decreases from 147.7 in the small instance to 45.2 in the large instance, confirming that overall Pareto quality deteriorates as the problem becomes more constrained. At the same time, the average number of active hubs remains between 10.6 and 11.3, showing that the proposed mechanism continues to preserve distributed validator participation instead of collapsing into a small centralized set of hubs. These results indicate that PI-Consensus scales coherently under increasing logistics complexity, while maintaining the decentralized character required for blockchain-enabled Physical Internet coordination.

Table 3: Comparison with PoW- and PoS-inspired proxy mechanisms on the medium PI workload

Mechanism	Validator-selection logic	F1 latency	F2 overhead	Avg. active hubs
PoW-proxy	Prioritizes the most powerful eligible hubs	44	42	1.4
PoS-proxy	Capacity-weighted selection among eligible hubs	35	25	3.6
PI-Consensus	NSGA-II selection under PI-native constraints	27	10.5	11.3

Table 3 compares PI-Consensus with PoW- and PoS-inspired proxy mechanisms on the medium PI workload. These proxies do not implement the full blockchain protocols; they reproduce their dominant validator-selection logic within the same PI assignment setting. The results show that PI-Consensus achieves about 23% lower latency and 58% lower overhead than the PoS-proxy, while preserving distributed validator participation above the required threshold.

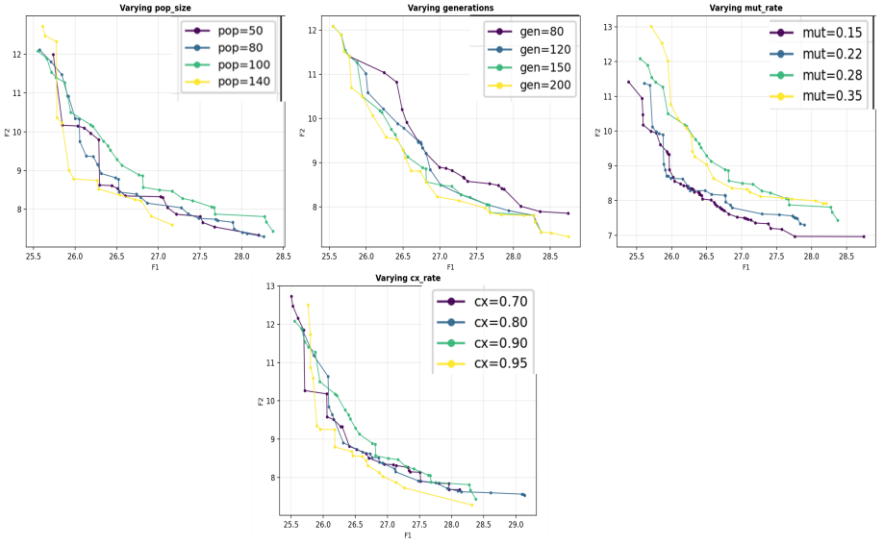


Figure 3: Pareto front sensitivity under different NSGA-II parameter settings

Figure 3 reports the one-at-a-time sensitivity of NSGA-II to population size, number of generations, mutation rate, and crossover rate. The results show that the Pareto front is more sensitive to the number of generations and the mutation rate, as these parameters visibly affect both convergence and spread. By contrast, population size and crossover rate have a more moderate influence within the tested ranges.

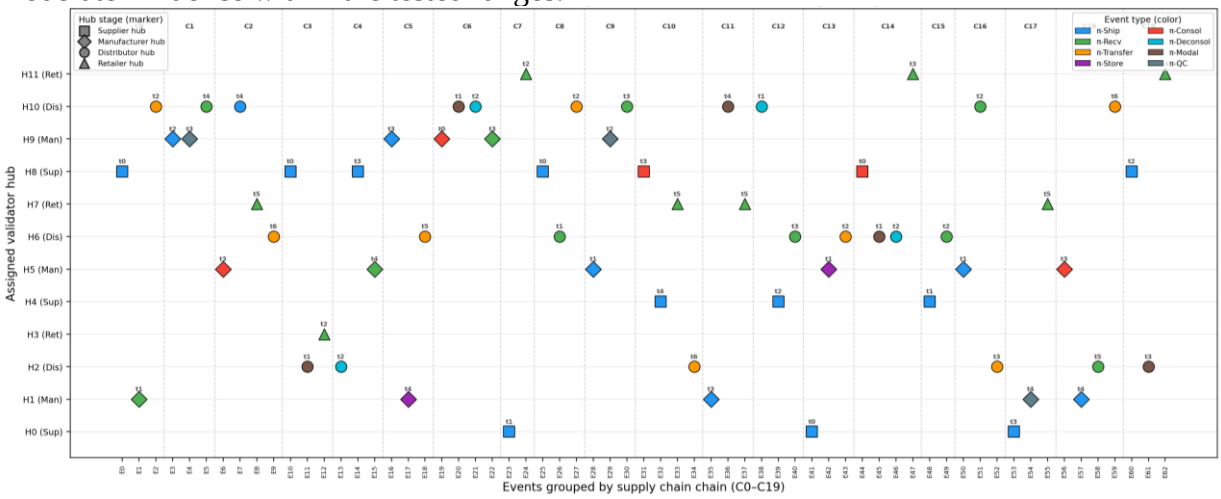


Figure 4 Hub event mapping

Figure 4 shows the event-to-validator assignment in the balanced solution. The 63 π-events are grouped by event sequence, placed on their assigned hub, colored by event type, and labeled by validation period. The strictly increasing period labels within each sequence confirm full precedence compliance, meaning that blockchain validation follows the physical order of supply chain operations. The figure also reveals a clear operational specialization: shipment events are mainly assigned to Supplier hubs, reception events to downstream hubs, and quality checks to Manufacturer hubs. This pattern emerges from PI-specific compatibility, protocol support, network structure, and optimization, rather than from generic blockchain rules.

7 Conclusion

This paper introduced PI-Consensus, a consensus framework grounded in the Physical Internet for blockchain-enabled supply chains. Validators are modeled as PI-hubs and transactions as ordered logistics events embedded in multi-stage supply chain sequences. By formulating validator assignment as a bi-objective optimization problem, the study demonstrates that consensus can be treated as an operational coordination problem driven by logistics feasibility, interoperability, capacity, and distributed participation, rather than by generic blockchain rules alone. This directly addresses the gap identified in the literature, where existing supply chain consensus approaches still represent validators as abstract nodes with limited consideration for hub interoperability and event-chain feasibility. Future work will focus on testing PI-Consensus on larger dynamic instances and prototyping it on a permissioned blockchain connected to realistic logistics event streams. Practical deployment will require progress on π-container standardization, π-interface interoperability, cross-platform data exchange, stakeholder adoption, and data-governance rules. A possible implementation path is to deploy PI-Consensus as a validator-selection layer for permissioned blockchain platforms connected to PI-hub event data.

References

Azzi, R., Chamoun, R. K., & Sokhn, M. (2019). The power of a blockchain-based supply chain. *Computers and Industrial Engineering*, *135*, 582–592. https://doi.org/10.1016/j.cie.2019.06.042

Bekrar, A., Ait El Cadi, A., Todosijevic, R., & Sarkis, J. (2021). Digitalizing the closing-of-the-loop for supply chains: A transportation and blockchain perspective. *Sustainability*, *13*(5), 2895.

Betti, Q., Khoury, R., Hallé, S., & Montreuil, B. (2019). *Improving Hyperconnected Logistics with Blockchains and Smart Contracts*. http://arxiv.org/abs/1904.03633

Chargui, T., Mahjoub, Y. I., Bekrar, A., & Trentesaux, D. (2024). Integration of Blockchain and the Physical Internet: A Review and Future Prospects. *Studies in Computational Intelligence*, *1136 SCI*, 583–592. https://doi.org/10.1007/978-3-031-53445-4_48

Collette, Y., & Siarry, P. (2005). Three new metrics to measure the convergence of metaheuristics towards the Pareto frontier and the aesthetic of a set of solutions in biobjective optimization. *Computers and Operations Research*, *32*(4), 773–792. https://doi.org/10.1016/j.cor.2003.08.017

de Oliveira, M. T., Reis, L. H. A., Medeiros, D. S. V., Carrano, R. C., Olabarriaga, S. D., & Mattos, D. M. F. (2020). Blockchain reputation-based consensus: A scalable and resilient mechanism for distributed mistrusting applications. *Computer Networks*, *179*. https://doi.org/10.1016/j.comnet.2020.107367

Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2002). A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II. In *IEEE TRANSACTIONS ON EVOLUTIONARY COMPUTATION* (Vol. 6, Number 2).

El Amine, B. M., & Lamia, S. T. (2023). Blockchain Adoption for a Sustainable Agricultural Supply Chain: Opportunities and Challenges for the Dairy Industry. *2023 International Conference on Decision Aid Sciences and Applications, DASA 2023*, 331–336. https://doi.org/10.1109/DASA59624.2023.10286692

Hasan, H. R., Salah, K., Jayaraman, R., Yaqoob, I., & Omar, M. (2021). Blockchain Architectures for Physical Internet: A Vision, Features, Requirements, and Applications. *IEEE Network*, *35*(2), 174–181. https://doi.org/10.1109/MNET.021.2000442

Jain, A. K., Gupta, N., & Gupta, B. B. (2025). A survey on scalable consensus algorithms for blockchain technology. *Cyber Security and Applications*, *3*. https://doi.org/10.1016/j.csa.2024.100065

Joo, J., & Han, Y. (2021). An evidence of distributed trust in blockchain-based sustainable food supply chain. *Sustainability (Switzerland)*, *13*(19). https://doi.org/10.3390/su131910980

Louzim, H., Maliki, F., Chargui, T., & Bekrar, A. (2024). PI-Container Flow Optimization in a Physical Internet Network. *2024 IEEE International Conference of the African Federation of Operational Research Societies, AFROS 2024 - Proceedings*. https://doi.org/10.1109/AFROS62115.2024.11036899

Mahjoub, S., & Abbad, H. (2025). An Intelligent Fair and Decentralized Consensus Mechanism for Blockchain-Based Supply Chain. *IFAC-PapersOnLine*, *59*(10), 244–249. https://doi.org/10.1016/j.ifacol.2025.09.043

Meyer, T., Kuhn, M., & Hartmann, E. (2019). Blockchain technology enabling the Physical Internet: A synergetic application framework. *Computers and Industrial Engineering*, *136*, 5–17. https://doi.org/10.1016/j.cie.2019.07.006

Montreuil, B. (2011). *Towards a Physical Internet: Meeting the Global Logistics Sustainability Grand Challenge*. www.cirrelt.ca

Montreuil, B., Meller, R. D., & Ballot, E. (2012). Physical internet foundations. *IFAC Proceedings Volumes (IFAC-PapersOnline)*, *14*(PART 1), 26–30. https://doi.org/10.3182/20120523-3-RO-2023.00444

Pan, S. (2024). Blockchain Technology for Logistics Collaboration in Physical Internet. *2024 International Conference on Automation in Manufacturing, Transportation and Logistics, ICaMaL 2024*. https://doi.org/10.1109/ICaMaL62577.2024.10919774

Rahman, M. S., Khalil, I., Atiquzzaman, M., & Bouras, A. (2025). A lightweight practical consensus mechanism for supply chain blockchain. *High-Confidence Computing*, *5*(1). https://doi.org/10.1016/j.hcc.2024.100253

Riouali, Y., & Rabhi, A. (2026). A Blockchain-Driven Approach to Enhancing Transparency and Traceability in Organic Food Supply Chain Management. *Applied Computational Intelligence and Soft Computing*, *2026*(1). https://doi.org/10.1155/acis/3114851

Saberi, S., Kouhizadeh, M., Sarkis, J., & Shen, L. (2019). Blockchain technology and its relationships to sustainable supply chain management. *International Journal of Production Research*, *57*(7), 2117–2135.

Sarfaraz, A., Chakrabortty, R. K., & Essam, D. L. (2023a). Reputation based proof of cooperation: an efficient and scalable consensus algorithm for supply chain applications. *Journal of Ambient Intelligence and Humanized Computing*, *14*(6), 7795–7811. https://doi.org/10.1007/s12652-023-04592-y

Sarfaraz, A., Chakrabortty, R. K., & Essam, D. L. (2023b). The implications of blockchain-coordinated information sharing within a supply chain: A simulation study. *Blockchain: Research and Applications*, *4*(1). https://doi.org/10.1016/j.bcra.2022.100110

Wüst, K., & Gervais, A. (2018). Do You Need a Blockchain? *Proceedings of IEEE European Symposium on Security and Privacy (EuroS&P)*, 133–157. https://doi.org/10.1109/EuroSP.2018.00014

Zhang, P., Zhou, M., Zhao, Q., Abusorrah, A., & Bamasag, O. O. (2021). A Performance-Optimized Consensus Mechanism for Consortium Blockchains Consisting of Trust-Varying Nodes. *IEEE Transactions on Network Science and Engineering*, *8*(3), 2147–2159. https://doi.org/10.1109/TNSE.2021.3079415

Zheng, Z., Xie, S., Dai, H., Chen, X., & Wang, H. (2017). *An overview of blockchain technology: Architecture, consensus, and future trends*. 557–564.

Zhu, F., Hu, F., Zhao, Y., Chen, B., & Tan, X. (2024). A Secure and Fair Federated Learning Framework Based on Consensus Incentive Mechanism. *Mathematics*, *12*(19). https://doi.org/10.3390/math12193068

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

A Physical Internet Approach to Disaster Management: Leveraging Open Distribution Networks for Humanitarian Supply Chains

M. Metin Inan [1], Matthieu Lauras[1,2], Benoit Montreuil[1,3]

1. Physical Internet Center, Supply Chain and Logistics Institute, H. Milton Stewart School of Industrial & Systems Engineering, Georgia Institute of Technology, Atlanta, United States
2. The Centre of Excellence in Supply Chain and Transportation (CESIT), Kedge Business School, France
3. Coca-Cola Material Handling & Distribution Chair

Corresponding author: minan@gatech.edu

Abstract: *Humanitarian organizations are widely criticized for fragmented logistics, duplicated readiness, and weak inter-organizational coordination, which together raise system cost and weaken surge response. This paper presents preliminary results from an exploratory study of how Physical Internet (PI) principles might address some of these concerns. We propose an Open Distribution Network (ODN), a shared three-tier depot and transport network in which organizations retain private inventory ownership through π-containers and access shared logistics capacity through explicit π-protocols, and compare it against a siloed baseline using a daily multi-agent simulator with four organizations, overlapping item portfolios, and a one-year multi-hazard timeline. The simulator uses heuristic planning, abstracted throughput, and stochastic demand realization within a deterministic disaster timeline. Initial results indicate that the ODN reduces annual operational cost (33.3% in the baseline scenario), reduces permanent shortage (4.99% to 0.75%), and reduces mean time-to-coverage (1.67 to 1.11 days), with cost advantage robust across most parameter ranges and bounded primarily by transport-cost intensity. We frame these as exploratory findings: they suggest mechanisms worth deeper study rather than universal claims and identify directions for high-fidelity validation.*

Keywords: *Physical Internet; Open Distribution Network; Humanitarian Supply Chain; Disaster Management; Shared Capacity; Inter-Organizational Collaboration; Hyperconnected Logistics; Governance Protocols*

Physical Internet (PI) Roadmap Fitness*: ☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories), ☒ PI Networks, ☒ System of Logistics Networks.*

Targeted Delivery Mode-s*: ☒ Paper, ☐ In-Person presentation*

1 Introduction

Humanitarian supply chain (HSC) performance has become increasingly consequential as natural and human-made disasters generate high-impact, multi-event operating years. Humanitarian organizations (HOs) must prepare under severe resource constraints while facing public and institutional expectations for rapid, equitable relief. Yet HSCs are persistently criticized in both academic and practitioner literature for limited inter-organizational

coordination, insufficient preparedness, weak operational flexibility, and avoidable duplication of assets and effort (Van Wassenhove, 2006; Balcik et al., 2010; Holguín-Veras et al., 2012). A central structural driver is fragmentation across actors: in many disaster-prone regions, multiple HOs serve overlapping affected populations while independently maintaining depots, transport arrangements, workforces, and pre-positioned inventories. The resulting duplication raises system-wide readiness cost without guaranteeing surge effectiveness, and the architectural choice between siloed and shared logistics infrastructure has received limited attention in the prepositioning literature.

The Physical Internet (PI) paradigm — an open, modular, protocol-governed logistics system built on shared hubs, standardized interfaces, and interoperable services (Montreuil, 2011; Montreuil et al., 2012; Pan et al., 2017) — offers a structural lens on this fragmentation. Its premise is that the relevant bottleneck in humanitarian preparedness is not simply inventory availability but the ability to stage, move, and govern scarce logistics capacity under uncertainty. PI principles are attractive because they speak directly to the operational and governance dimensions that the literature flags as obstacles to inter-organizational collaboration. However, sharing is not automatically beneficial: a shared network can introduce congestion, raise lateral transport costs, or systematically disadvantage lower-priority missions if governance rules are poorly designed. A rigorous comparison must therefore identify both the mechanisms of advantage and their limits.

This paper presents preliminary results from an exploratory study of how PI principles might mitigate some of the persistent concerns about HSC fragmentation. We propose and evaluate an Open Distribution Network (ODN): a shared three-tier network of depots and transport services in which organizations retain ownership of their inventory through π-containers, preserve mission-specific coverage commitments, and access shared logistics capacity through explicit π-protocols. We compare this architecture to a siloed baseline using a daily multi-agent simulator with four organizations, overlapping item portfolios, and a one-year multi-hazard timeline including wildfire and hurricane seasonality.

The study is positioned as an early-stage characterization rather than a definitive evaluation. The contribution is threefold. First, we instantiate PI primitives for humanitarian logistics: open π-hubs at three tiers, modular π-services, π-containers preserving private inventory ownership, shared information interfaces, and a capacity-arbitration governance π-protocol. Second, we identify operational mechanisms through which the ODN may improve cost and service performance: facility consolidation, lateral same-day flow under realized-demand variance, and capacity pooling across organizations with offset hazard profiles. Third, we provide initial paired-seed simulation results suggesting where ODN is cost- and service-superior to a siloed baseline, and where its advantage is bounded; we emphasize parameter sensitivities and limitations that motivate future high-fidelity work.

2 Literature Review and Positioning

2.1 Critique of Fragmented Humanitarian Logistics

Humanitarian logistics differs structurally from commercial logistics: demand is stochastic and lumpy, urgency is high, actors are heterogeneous, and performance is measured in timeliness and coverage as well as cost (Van Wassenhove, 2006; Holguín-Veras et al., 2012). HSCs have been systematically criticized along three recurring dimensions.

Coordination is chronically weak: Balcik et al. (2010) identify restrictive donor funding structures, differing organizational mandates, a lack of logistical standardization, and the absence of a central coordinating authority as primary barriers to relief chain integration. They argue that systemic improvement requires formalized coordination mechanisms, such as strategic partnerships and joint procurement, to manage interdependencies that goodwill alone cannot resolve. **Operational flexibility is limited**: hierarchical, organization-specific network structures reduce agility during response (Charles et al., 2010). **Avoidable duplication is widespread**: when multiple HOs serve overlapping populations, each maintains separate warehouses, fleets, and inventories with limited interoperability, raising system cost while still failing to guarantee surge effectiveness (Balcik et al., 2010; Holguín-Veras et al., 2012). The duplicated readiness model is both inefficient and operationally fragile.

2.2 Responses: Shared Infrastructure and Horizontal Cooperation

Structural responses to fragmentation cluster in two streams. The first is horizontal cooperation in inventory, most prominently the United Nations Humanitarian Response Depot (UNHRD) network, which provides shared warehousing and stock-borrowing services across more than 80 partner organizations. Toyasaki et al. (2017) study UNHRD using a non-cooperative game-theoretic newsvendor model and show that stock rationing rules affect efficiency and equity. Shared inventory infrastructure is therefore theoretically valuable and operationally established, but the governance rules under which it operates determine whether benefits are realized. The second stream, logistics service network design for shared depot infrastructure, generally finds measurable cost benefits from shared deployment relative to organization-specific solutions. A persistent insight is that heterogeneity of mission and portfolio must be respected: coordination architectures should share capacity rather than ownership.

2.3 Prepositioning and Inventory Planning

Prepositioning relief items at strategically located warehouses is the dominant preparedness lever in humanitarian operations. Balcik and Beamon (2008) provide the canonical formulation integrating facility location, inventory, and budget under stochastic disaster scenarios; subsequent work has extended this in several stochastic, location-routing, and procurement directions. However, this literature largely treats organizational architecture as fixed: the choice between siloed and shared infrastructure is rarely the unit of comparison.

2.4 The Physical Internet and Humanitarian Applications

The Physical Internet paradigm targets order-of-magnitude improvements in efficiency, sustainability, and resilience through the universal interconnection of open logistics services and networks (Montreuil, 2011; Montreuil et al., 2012). PI constituents include π-containers (modular standardized handling units), open π-hubs (shared multi-tier facilities), π-services (modular contractual interfaces), and π-protocols (governance rules for routing and allocation). Humanitarian applications remain comparatively limited but show promise: Ado et al. (2014) introduced the framing of PI-enabled interconnected humanitarian logistics; Grest et al. (2021) develop simulation-based assessments comparing baseline humanitarian practices with hyperconnected alternatives. Recent work has linked PI-enabled distribution to readiness assessment: Inan et al. (2024) evaluate hyperconnected supply chains using a quantitative readiness metric and show how open contracted storage centers can improve disruption preparedness.

2.5 Positioning of This Paper

This paper combines the four streams above. It transfers the PI architectural lens to the critique of humanitarian fragmentation, treats the architectural choice between siloed and shared infrastructure as the unit of comparison, and adds explicit governance protocols as first-class objects of analysis. The work is preliminary by design: it explores whether PI principles plausibly address the persistent concerns flagged above and characterizes the mechanisms and limits of any advantage observed.

3 Open Distribution Network Architecture

The ODN replaces duplicated private depot hierarchies with a shared three-tier distribution network. In the siloed regime, each organization operates its own main, regional, and staging depots. In the ODN regime, the physical facilities and transport services are shared, while inventory ownership remains private and auditable. Table 1 summarizes the architectural distinction.

The architecture has five PI components. Open π-hubs represent the shared depots across tiers, replacing per-organization replication. Modular π-services define standardized storage, handling, and truck-trip services contracted on a per-use basis. π-containers preserve organizational ownership and support donor and audit constraints even when inventory is co-located in shared facilities. Information π-interfaces provide a shared inventory ledger and forward-looking need projections, supporting capacity planning and anticipatory replenishment without exposing operational specifics. A governance π-protocol arbitrates capacity contention when concurrent organizational requests exceed available shared space, with a criticality-priority rule as the baseline.

The expected ODN advantage decomposes into three mechanisms. Facility consolidation reduces duplicated fixed readiness costs by replacing four parallel hierarchies with one shared hierarchy. Lateral T3 flow permits same-day movement from nearby staging depots within a cluster when realized demand exceeds projection at one zone and falls short at another, preserving same-day commitments that the canonical-chain-only siloed regime cannot. Capacity pooling applies a cross-organizational factor ρ to shared depot capacity, reflecting offset hazard profiles such as wildfire and hurricane seasons.

These mechanisms also speak directly to barriers identified in Section 2.1. Trust is preserved because π-containers maintain organizational identity and donor accountability even within shared facilities. Data sharing is bounded by the minimum-disclosure information π-interface. Contracting is decomposed into modular per-use service transactions rather than complex multi-party long-term agreements. Donor earmarking is respected because the architecture pools capacity, not ownership.

4 Simulation Model and Experimental Design

We instantiate the comparison as a daily multi-agent simulator over a one-year horizon. This simulation-based evaluation follows the methodological logic of Grest et al. (2021), who argue that the complexity of HSCs makes simulation an appropriate vehicle for assessing PI-oriented reconfiguration. Decision functions are implemented as heuristics rather than optimizations to (i) match practitioner reality and (ii) isolate architectural effect from planning sophistication.

Table 1: *Architecture comparison between siloed regime and Open Distribution Network*

Dimension	Siloed Regime	Open Distribution Network
Depot count	12 per organization × 4 organizations = 48 facilities	12 shared facilities in 3 T3 clusters
T3 sizing	Per-zone	Per-cluster, then distributed
Capacity headroom	None; projection error acts as buffer	None; projection error acts as buffer
Pooling factor on shared capacity	Not applicable	$\rho \approx 0.85$ on cross-organizational sum
Inventory ownership	Private, in private facilities	Private (π-containers), in shared facilities
Transport	Per-organization ad hoc	Shared ad hoc, with consolidation
Lateral flows	Not permitted	T3↔T3 within cluster same day; T2↔T2 in 1 day
Replenishment	Within own hierarchy	Shared hierarchy plus supplier fallback
Capacity contention	None	Criticality-priority π-protocol
Information interface	None	Shared ledger and need projections

Sets and primitives. Four organizations (A, B, C, D) serve ten demand zones using four item types. A, B, C are multi-hazard generalists; D is a wildfire-only specialist serving zones Z1–Z4 only. Item portfolios overlap: water is owned by A and C, medical by B, food by C, and wildfire kits by D. Items have criticality scores used in dispatch ordering and ODN capacity arbitration. Each regime uses 12 depot locations per hierarchy: one T1 main depot, three T2 regional depots, eight T3 staging depots grouped into three regional clusters. The siloed regime replicates this hierarchy for each organization, yielding 48 facilities; the ODN uses the same 12 physical facilities as shared infrastructure. Lead times: 2 days for T1→T2, 1 day for T2→T3, same-day for T3→Zone, same-day for intra-cluster T3↔T3 in ODN only, and 30 days for external procurement.

Service-level commitments. Each organization-item pair commits to a cumulative service-level curve over deadlines $d \in \{1, 3, 5, 7\}$ days. Let $SL_{o,i,d}$ denote the target fraction of organization o's item i demand that should be covered by deadline d.

Annual planning. Each organization, at the simulation start, independently projects total seasonal demand for each zone using zone population, climatological event rate, affected population fraction, and per-capita item requirement. The projection is perturbed by an asymmetric triangular error ε with bounds [−0.20, +0.50] and mode +0.20, reflecting the empirically observed humanitarian tendency toward conservative over-provisioning combined with irreducible projection uncertainty. The asymmetry is the central source of architectural differentiation: when an organization over-projects (the modal case), its prepositioned inventory exceeds realized need at most zones but falls short at zones where demand realization runs hot.

Inventory is allocated to tiers using a two-layer sizing logic. Per-event sizing for lower tiers (T3 and T2) allocates buffer capacity for a single expected event, sized so that the same-day delivery commitment $SL_{(o,i,1)}$ can be met from staging stock and the 1–3 day commitment increment can be met from regional stock. Seasonal-residual sizing for the main tier (T1) then allocates the remainder needed to cover the full annual commitment after T3 and T2 are sized.

The two regimes diverge sharply on T3 sizing: the architectural mechanism at the heart of the comparison. *In the siloed regime,* each zone is assigned to a specific T3 staging depot, and that T3 is sized independently for the zone's per-event demand. This means each Siloed T3 must absorb its assigned zone's full projection error: if zone z's projected event demand is over-estimated by 35%, the T3 serving z holds 35% excess inventory locally with no recourse to share that excess with other zones. *In the ODN regime,* T3 staging depots are organized into regional clusters, and the cluster's aggregate per-event demand is sized once across all zones in the cluster, then distributed across the cluster's T3 depots proportionally to each zone's projected demand share. Because zone-level projection errors partially cancel within a cluster, the cluster aggregate is smoother than the sum of independent zone projections, and ODN therefore commits less total staging-tier inventory at the same service-level commitment.

The architectural capacity reduction in ODN comes from the pooling factor ρ on shared depot capacity. Each shared depot must reserve enough storage to hold the inventory all organizations have allocated to it; in the most conservative case, the depot would reserve capacity equal to the full *sum* of organizational requirements, treating each organization's peak as if it occurred simultaneously. In reality, organizations have hazard exposures offset across the year, wildfire-driven demand peaks in late winter and spring, hurricane-driven demand peaks in late summer and fall, so the system never simultaneously requires every organization's peak capacity. The pooling factor captures the fraction of summed organizational requirements that the shared depot needs to reserve. Shared ODN capacity at depot w then aggregates organizational allocations and applies the pooling factor ρ:

$$K_w^{ODN} = \rho \sum_{o \in O} \sum_{i \in I} v_i I'_{oiw},$$

where v_i is item volume, I'_{oiw} is initial allocated inventory for organization o and item i at depot w, and $\rho \in (0, 1]$ captures the additional capacity reduction from cross-organizational temporal aggregation (organizations whose hazard exposure is offset across the year, like wildfire-driven and hurricane-driven organizations, do not require concurrent peak capacity at the shared depot). Baseline $\rho = 0.85$, sensitivity-tested in Section 5.3.

Daily simulation loop. Each day proceeds through (1) disaster realization with multiplicative N(1.0, 0.20^2) noise on projected demand and probability, 0.30 adjacent-zone partial hits, (2) regime-conditional dispatch — Siloed escalates assigned T3 → T2 → T1 with no lateral flow; ODN uses assigned T3 → other T3s in cluster (same-day lateral) → T2 → T1, (3) trip consolidation in ODN, (4) replenishment trigger for disasters ending that day, (5) in-transit advancement, and (6) demand expiration. Realization noise is paired across regimes within each replicate.

Replenishment. The day after a disaster ends, organizations re-project remaining-season need with a fresh ε draw and place orders to close the gap toward updated tier targets. When upstream stock is insufficient, the unfilled portion converts to direct external procurement at that tier (supplier fallback), incurring 30-day lead time and procurement cost. Organizations do not reshuffle pre-positioned inventory between depots between disasters; once allocated, stock remains at its assigned depot until consumed or replenished from upstream. This reflects current operational practice: inter-depot reshuffling is logistically expensive and complicated by donor-accountability constraints, but it is a meaningful limitation of the present model that warrants

investigation in future work, since a more permissive reshuffling regime would likely extend the architectural advantage we observe.

Experimental design. Paired-seed replicates within each condition. We sweep seven parameters one at a time around baseline, summarized in Table 2.

Table 2: *Experimental design*

Run	Replicates	Architecture	Sweep	Variants
0	5 each	Siloed, ODN	–	No-disaster baseline (+100% projection error)
1	30 each	Siloed, ODN	–	Main result
2	10 each	Siloed, ODN	Population scale	$\times 0.5$, $\times 2.0$, $\times 4.0$
3	10 each	Siloed, ODN	Storage cost	$\times 0.25$, $\times 4.0$
4	10 each	Siloed, ODN	Depot fixed cost	$\times 0.5$, $\times 2.0$
5	10 each	Siloed, ODN	Transport cost	$\times 0.5$, $\times 2.0$, $\times 4.0$
6	10 each	Siloed, ODN	Procurement lead time	60d, 90d, 180d
7	10 each	ODN only	Pooling factor ρ	$\{1.00, 0.85, 0.75, 0.70, 0.60\}$
8	10 each	Siloed, ODN	Demand noise σ	$\{0.10, 0.30, 0.40\}$

5 Initial Results

We report initial findings from the experimental matrix above. We deliberately frame these as preliminary evidence about plausible mechanisms rather than definitive performance claims; Section 6 discusses limitations that motivate high-fidelity validation.

5.1 Baseline Comparison

In the main 30-replicate run, ODN reduces annual operational cost from \$12.39M to \$8.26M (a 33.3% reduction). Table 3 decomposes this. The largest contributor is depot fixed cost: four private hierarchies cost \$6.96M annually, the shared hierarchy costs \$1.74M. Depot capacity cost decreases modestly through pooling and cluster sizing. Transport cost rises in ODN by \$1.27M because lateral same-day movements activate to maintain coverage. This is not a failure of consolidation, it is the operational price paid for service recovery, with the fixed-readiness savings dominating.

Table 3. *Baseline operational cost decomposition, Run 1 (mean ± 1σ across 30 paired replicates).*

Component	Siloed	ODN	Difference
Depot fixed	\$6.96M	\$1.74M	–\$5.22M
Depot capacity	\$1.12M	\$0.96M	–\$0.16M
Transport	\$0.63M	\$1.90M	+\$1.27M
Handling	\$0.14M	\$0.15M	+\$0.01M
Procurement	\$3.54M	\$3.52M	–\$0.02M
Total	**\$12.39M**	**\$8.26M**	**–\$4.13M**

Aggregate service performance also improves: permanent shortage falls from 4.99% to 0.75% and mean time-to-coverage decreases from 1.67 days to 1.11 days. The aggregate advantage is, however, distributed unevenly across (organization, item) lines. ODN substantially closes the gap on multi-organization items where the cluster-mediated lateral mechanism has the most leverage: A water final miss drops from 8.4% to 0.2%, B medical from 4.4% to 0.0%, and C water from 6.0% to 0.6%. On single-organization or low-urgency items, the architectural advantage is smaller or absent: C food and D wildfire_kit show comparable performance across the two regimes, and the aggregate ODN advantage is driven by the dominance of the water and medical gains in the demand mix rather than by uniform improvement across all lines. We discuss the structural reasons for this unevenness in Section 6.

5.2 Sensitivity Analysis

Across sweeps over population scale, depot fixed cost, transport cost, supplier lead time, and demand realization variance $\sigma \in [0.10, 0.40]$, ODN's operational cost advantage holds in all tested ranges, scaling primarily with depot fixed cost and shrinking only at high transport multipliers (break-even $\approx \times 4.24$). On permanent shortage, ODN remains below Siloed across the demand-noise range.

5.3 Pooling and Governance Tradeoff

Figure 1 reports the ODN-only pooling sweep across four panels: capacity readiness cost, total operational cost, π-protocol contention events, and per-(organization, item) service rate, as the pooling factor ρ varies from 1.00 (no pooling, the shared depot reserves capacity equal to the full sum of organizational requirements) to 0.60 (aggressive pooling, the shared depot reserves only 60% of the summed requirements, relying heavily on cross-organizational seasonal offset).

The pooling sweep yields three observations. First, facility consolidation is the dominant cost mechanism: even at $\rho = 1.0$ with no pooling reduction applied, ODN costs \$8.75M, well below the \$12.39M Siloed baseline. Most of the architectural cost advantage comes from replacing four parallel hierarchies with one, not from pooling capacity within the shared hierarchy. Second, lower ρ produces additional capacity savings but at the cost of governance load: reducing ρ from 1.00 to 0.60 lowers operational cost from \$8.75M to \$7.49M, but π-protocol contention events more than double (from 432 to 591) — the protocol must arbitrate competing organizational requests more frequently when shared capacity is tight. Third, service degradation under aggressive pooling concentrates on the lowest-criticality item: at $\rho = 0.60$, C food coverage falls from 0.996 to 0.971 while D wildfire_kit holds steady around 0.99 and high-criticality lines (A water, B medical, C water) remain at full coverage. Tighter pooling shifts service burden toward lower-priority missions through the criticality-priority protocol, a design choice with explicit equity implications. The pooling factor is therefore not a free tuning knob but a real cost-vs-equity tradeoff lever.

6 Discussion and Limitations

The preliminary results suggest several mechanisms worth investigating further, but they should be read as exploratory rather than definitive.

Cost advantage is mechanism specific. In the baseline, the cost advantage is not driven by transport consolidation: ODN transport is *more* expensive because lateral movements activate to maintain coverage. The cost advantage comes from eliminating duplicated fixed depot

readiness and from the seasonal pooling factor ρ on shared depot capacity. The architecture is service-superior across the tested range but cost-superior only when transport costs are not extreme.

Service advantage is uneven across (organization, item) lines. Where multiple organizations or multi-event-rate zones produce smoothing across the cluster, ODN's lateral T3↔T3 mechanism converts realization variance into preserved coverage: substantial gains on water (A and C) and medical (B). On items concentrated upstream by their service-level commitment (C food) or owned by a single organization in a single cluster (D wildfire_kit), the lateral mechanism has limited leverage and ODN performs comparably to Siloed. The aggregate advantage in permanent shortage reduction therefore reflects the specific portfolio mix; a different mix of items, organizations, and hazard exposures could produce different aggregate magnitudes.

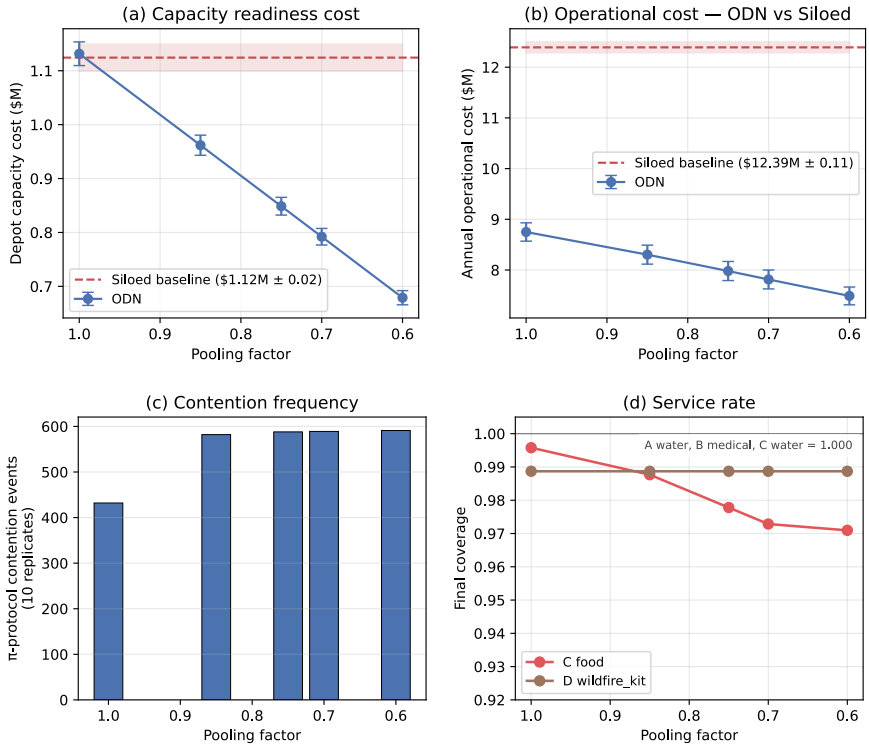


Figure 1: *ODN sensitivity to the pooling factor ρ.*

Governance is central, not auxiliary. The criticality-priority π-protocol protects high-criticality items but can disadvantage lower-criticality portfolios under tight pooling. A proportional-fairness protocol would likely produce different equity-efficiency tradeoffs. This points to protocol design as a first-class research object alongside network design.

Limitations. The simulator abstracts depot throughput, treats demand realization as instantaneous, applies heuristic planning rather than rolling-horizon optimization, and assumes unconstrained fleet availability with per-trip cost. The disaster timeline is deterministic; only demand realization within that timeline is stochastic. Cost parameters are stylized within order-of-magnitude bounds rather than sourced from a single operational document. Inventory remains strictly private and does not reshuffle between depots between disasters; cross-organizational borrowing under π-protocol terms — which would represent a stronger

expression of PI mutual aid — is not implemented and would be expected to extend the architectural advantage further. Each of these limitations represents a degree of operational realism that future work should add.

Directions for future work. Three directions follow from the present results. First, high-fidelity validation in a discrete-event simulator with rolling-horizon planning, explicit throughput constraints, and network-disrupting events. Second, exploration of cross-organizational inventory borrowing under π-protocol terms, preserving donor-restricted ownership while permitting time-bounded mutual aid; this represents a stronger expression of PI principles than the strict private-ownership architecture studied here. Third, empirical estimation of the pooling factor ρ from partner-organization data and historical disaster timelines, replacing the calibratable parameter with a measured one. Across all three directions, the empirical study of allocation protocols that sustain participation in shared humanitarian infrastructure remains an open and operationally consequential question.

References

- Ado, A., Ben Othmane, I., Matei, M., Montreuil, B. (2014): Towards Physical Internet enabled Interconnected Humanitarian Logistics. Proceedings of the 1st International Physical Internet Conference, Québec City, Canada.
- Balcik, B., Beamon, B. M. (2008): Facility location in humanitarian relief. *International Journal of Logistics Research and Applications*, 11(2), 101–121.
- Balcik, B., Beamon, B. M., Krejci, C. C., Muramatsu, K. M., Ramirez, M. (2010): Coordination in humanitarian relief chains: Practices, challenges and opportunities. *International Journal of Production Economics*, 126(1), 22–34.
- Charles, A., Lauras, M., Van Wassenhove, L. N. (2010): A model to define and assess the agility of supply chains: building on humanitarian experience. *International Journal of Physical Distribution & Logistics Management*, 40(8/9), 722–741.
- Grest, M., Inan, M. M., Cohen, Y. M., Barenji, A. V., Dahan, M., Lauras, M., Montreuil, B. (2021): Design of a Simulation-Based Experiment for Assessing the Relevance of the Physical Internet Concept for Humanitarian Supply Chains. Proceedings of IPIC 2021, Online, Greece. HAL: hal-03395146.
- Holguín-Veras, J., Jaller, M., Van Wassenhove, L. N., Pérez, N., Wachtendorf, T. (2012): On the unique features of post-disaster humanitarian logistics. *Journal of Operations Management*, 30(7–8), 494–506.
- Inan, M. M., Montreuil, B., Lauras, M., Dahan, M. (2024): Assessing and Enhancing Readiness in Hyperconnected Supply Chains. Proceedings of IPIC 2024, Savannah, GA, USA. DOI: 10.35090/gatech/11078. HAL: hal-04709114.
- Montreuil, B. (2011): Toward a Physical Internet: Meeting the Global Logistics Sustainability Grand Challenge. *Logistics Research*, 3(2–3), 71–87. https://doi.org/10.1007/s12159-011-0045-x
- Montreuil, B., Meller, R. D., Ballot, E. (2012): Physical Internet Foundations. CIRRELT-FSA-2012-56.
- Toyasaki, F., Arikan, E., Silbermayr, L., Falagara Sigala, I. (2017): Disaster relief inventory management: Horizontal cooperation between humanitarian organizations. *Production and Operations Management*, 26(6), 1221–1237.
- Van Wassenhove, L. N. (2006): Humanitarian aid logistics: Supply chain management in high gear. *Journal of the Operational Research Society*, 57(5), 475–489.

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE

A roadmap-aligned decision-support methodology to identify scaling mechanisms for Physical Internet-compliant logistics innovations: the IKIGAI approach

Filippo Mauro[1], Fabio Cartolano[1], Paola Astegiano[1], Paola Cossu[1], Dimos Touloumidis[2], Sofoklis Dais[2], Georgia Ayfantopoulou[2], Alkis Pitelis[3], Ioanna Fergadiotou[3]

1. FIT Consulting srl, Rome, Italy
2. Centre for Research and Technology Hellas (CERTH/HIT), Thessaloniki, Greece
3. Inlecom Innovation, Athens, Greece

Corresponding author: Filippo Mauro - mauro@fitconsulting.it

Keywords: *Physical Internet Roadmap, PI-Norm, Scaling Mechanisms.*

Conference Domain Fitness: *This contribution is positioned at the intersection of Hyperconnected supply chain & logistics systems, Synchromodal and multimodal transportation, Collaborative logistics and horizontal cooperation, Transition strategies and implementation roadmaps within the IPIC 2026 domains. Drawing on the IKIGAI project's Logistics Innovations experience, it shows how Physical Internet principles can be operationalised in real operational contexts through complex multicriteria analysis and decision-making supporting tools. The alignment of IKIGAI methodologies benchmarked by project's Logistics Innovations within the ALICE and PI ZEL Roadmaps provides concrete evidence of how complex supply chains can overcome barriers in systematizing operational processes and reach compliance to the PI-Norm to finally achieve sustainability and efficiency through process standardization.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),☐ Transportation Equipment, ☒ PI Networks, ☒ System of Logistics Networks, ☐ Vertical Supply Consolidation, ☒ Horizontal Supply Chain Alignment, ☐ Logistics/Commercial Data Platform, ☒ Access and Adoption, ☒ Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

1 Introduction

The decarbonisation of freight transport and logistics is now a central objective of climate and industrial transition strategies, yet the sector remains difficult to transform because emissions are tied to growing freight demand, operational complexity and the persistence of fragmented service structures (McKinnon, 2018). Logistics systems are still organised through heterogeneous flows, nodes and networks that often operate with limited coordination, low interoperability and weak process standardisation, which makes scale-up difficult even when promising innovations already exist (Hesse and Rodrigue, 2004). In this context, the Physical Internet (PI) has been proposed as a systemic response to the structural inefficiencies of conventional logistics, promoting a transition towards open, standardised, interoperable and hyperconnected logistics networks rather than isolated

improvements at the level of single actors or single services (Montreuil, 2011; Ballot et al., 2014; Crainic and Montreuil, 2016). Around this vision, a set of European roadmap initiatives has progressively emerged to translate PI principles into a more structured transition pathway, identifying strategic directions, enabling conditions and priority areas for implementation. However, a practical gap remains between this roadmap-level guidance and its operational implementation. While PI roadmaps clarify where logistics systems should move, they do not usually provide a clear pathway for guiding innovations from pilot application to scalable and PI-aligned deployment. This paper addresses that missing methodological layer within the context of the EU-funded IKIGAI project (GA 101202912), which provides a concrete setting in which PI-oriented scale-up can be analysed across different Logistics Innovations (LIs). Building on the Zero Emission Logistics (ZEL) roadmap (ETP, 2019) and PI roadmap (ALICE-ETP, 2020), IKIGAI advances five complementary LIs that address different but interrelated barriers to scale-up, including technological readiness, governance, regulatory conditions and market uptake. Within this setting, the PI Norm is conceived as the process-standardisation backbone that operationalises PI principles through three independent components:

- standardised logistics processes
- interoperable digital and documentation flows
- governance and compliance mechanisms ensuring auditability and replicability.

PI Norm compliance is therefore treated not as a formal add-on, but as a condition for making LIs transferable across cases and consistent with the broader PI transition direction. Figure 1 summarises the main elements of IKIGAI, including the five LIs and the main project catalysers, such as the PI Rulebook, PI Norm, pilots audit, PI Coalition, Twin Transition Genome, and Impact Augmenter Kit.

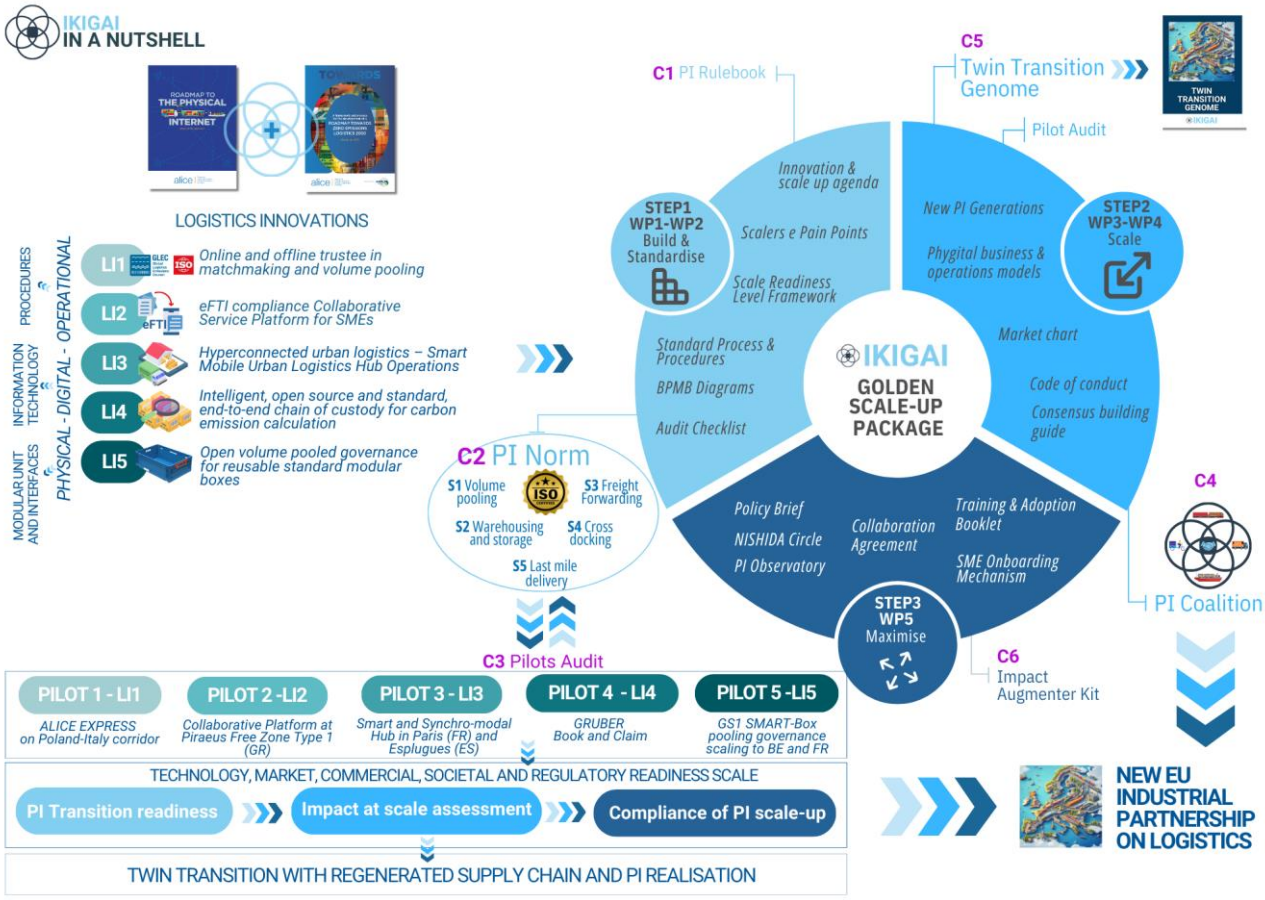


Figure 1 – IKIGAI in a Nutshell

Against this background, this paper proposes a methodology to identify scaling mechanisms that can help LIs move from pilot application towards more scalable and PI-compliant deployment. The next section presents the methodological framework used to assess the five LIs in a structured and comparable way and to support the identification of these mechanisms.

2 Methodology

The paper proposes a methodology to identify scaling mechanisms for LIs that aim to move towards Physical Internet compliance. The method is based on a simple principle: scale-up should not be understood only as growth in size or reach, but as progression towards a logistics configuration that is more standardised, interoperable and replicable. For this reason, the methodology combines target definition, current assessment, gap analysis and scaling mechanism identification within one structured framework.

The framework is organised into four steps. Step 1 defines the TO-BE condition, that is, the maturity profile that each LI aims to reach. Step 2 defines the AS-IS condition through a structured current assessment. Step 3 calculates and prioritises the gap between the current state and the target state. Step 4 translates the prioritised gaps into candidate scaling mechanisms. The same logic can be applied to different LIs because it is based on a common assessment structure rather than on one specific service model.

2.1 Step 1: TO-BE analysis through scaling dimensions and target scenarios

The first step defines the desired future condition of each LI. Its purpose is to identify the maturity profile that the innovation intends to reach through its scale-up path.

To do so, the methodology uses six scaling dimensions. Each dimension is expressed through four ordered maturity states. Each state is associated with a score from 1 to 4, where 1 represents the lowest maturity level and 4 the highest. This makes it possible to represent the target condition of each LI as a structured combination of scores.

The six dimensions are the following.

- **Operational.** This dimension captures the geographic and functional scope of the LI. It reflects how widely the solution is deployed and how far it is integrated into transport networks. The four maturity states are: *Isolated*, *Local/Regional*, *Network-integrated*, *Scaled & optimised*.
- **Technological.** This dimension assesses digital maturity and interoperability. It captures the extent to which the innovation relies on proprietary systems, basic interoperability, compliance-ready interoperability, or advanced digital capabilities. The four maturity states are: *Proprietary*, *Interoperable (basic)*, *Compliant & interoperable (advanced)*, *Intelligent & automated*.
- **Regulatory, policy and governance.** This dimension assesses the degree of alignment with regulatory conditions, policy frameworks, governance arrangements, and compliance structures. The four maturity states are: *Unclear/informal*, *Pilot-based*, *EU-aligned*, *Standardised & scalable*.
- **Market and economic.** This dimension reflects demand conditions and the level of market uptake. It captures whether the innovation still faces weak demand, is starting to emerge, is growing in a stable way, or has reached a mature and expansive position. The four maturity states are: *No/weak demand*, *Emerging*, *Stable growth*, *Mature/expansive*.
- **Societal and organisational.** This dimension captures stakeholder engagement, organisational readiness, and the degree to which the innovation is embedded in shared ways of working. The four maturity states are: *Siloed*, *Engaged*, *Coordinated*, *Institutionalised*.
- **Sustainability and circularity.** This dimension measures the extent to which environmental performance and circular economy principles are integrated into the innovation. The four maturity states are: *Initial*, *Developing*, *Advanced*, *Systemic*.

These dimensions were selected because a PI-oriented transition cannot be assessed through technology alone. A logistics service may show good technical performance and still remain difficult to scale if its governance is weak, its operating model is not repeatable or its market conditions are still fragile. For each dimension, the LI selects the maturity state it aims to reach. The selected levels should define a credible and achievable target condition, not an idealised or utopian end state.
The resulting combination of selected maturity states identifies the TO-BE profile of the innovation. This TO-BE profile is interpreted through five target scenarios:

- **Baseline.** The innovation remains confined to pilots or to a limited corridor. Integration is partial. Adoption is slow and uneven. Sustainability actions exist, but they are not yet embedded end-to-end.
- **Collaborative.** Scale-up is driven mainly by coordination among shippers, logistics service providers, SMEs, and other partners. Shared workflows and trust mechanisms support broader participation and reduce fragmentation.
- **Accelerated.** Adoption increases because interoperability and compliance readiness reduce barriers to scale. Digital processes become repeatable across contexts. Sustainability moves closer to standard practice.
- **Disruptive.** Rapid change is driven by advanced digital capabilities and strong optimisation logic. The innovation can scale quickly, but uptake may remain uneven across actors or territories.
- **Transformative.** The innovation becomes part of a fully integrated ecosystem. Operations, governance, data sharing and sustainability work together end-to-end. Replication at scale becomes feasible.

These scenarios are not forecasts. They are structured target configurations that help ensure coherence across the selected maturity levels. In this sense, the TO-BE analysis is not limited to choosing six separate scores. It defines a consistent direction of development for the LI as a whole.

2.2 Step 2: AS-IS assessment through relevance weighting and current scoring

The second step assesses the current condition of each LI. This step has two components.
The first component is relevance weighting. For each of the six dimensions, the LI assigns a relevance level: *Low*, *Medium*, or *High*. Relevance does not describe maturity. It expresses the strategic importance of that dimension for the scale-up of the specific innovation. This distinction matters because the same maturity gap may be far more critical in one case than in another.
The second component is the current scoring. This is carried out through a customised checklist. The checklist is organised by dimension, and within each dimension the questions are grouped by the same four maturity states used in Step 1. This gives the method a clear internal coherence: the categories used to define the target state are also used to assess the current state.
In practical terms, the checklist allows each LI to read the statements associated with the four maturity states and recognise which one best reflects its present condition. The selected state becomes the AS-IS score for that dimension. This makes the assessment easier to apply and reduces ambiguity in the later comparison between current state and target state.
The use of dimension-specific checklists is important because the maturity level has to emerge from a more structured reading of how the innovation currently operates. The exact wording of the checklist can vary according to the characteristics of the innovation, but the scoring logic remains stable across applications.
At the end of this step, each LI has one selected maturity state for each dimension in the AS-IS condition and one relevance level for each dimension. This provides the current status required for the gap analysis.

2.3 Step 3: Gap calculation and prioritisation

The third step measures the distance between the AS-IS profile and the TO-BE target profile. For each dimension, the gap is calculated as the difference between the target score and the current score. If the current level is already equal to or higher than the target, the gap is treated as zero.
This gap is then weighted through the relevance assigned in Step 2. In this way, the method does not only show where the innovation is distant from the target. It also shows where that distance matters most for the scale-up pathway of that specific case.
The output of this step is a prioritised gap profile. This profile identifies the main bottlenecks that constrain the transition towards the target configuration. The purpose of the step is diagnostic. It turns the maturity assessment into an ordered representation of where intervention is most needed.

2.4 Step 4: Translation of prioritised gaps into scaling mechanisms

The fourth step translates prioritised gaps into scaling mechanisms. A scaling mechanism is defined as a replicable package of levers that supports one concrete transition step towards a more scalable logistics configuration. It is not a generic label such as "better governance" or "more interoperability". It is a more structured response that may include process rules, governance arrangements, digital supports, operating routines or standard artefacts.
The method does not assign one mechanism to each single gap. Instead, the prioritised gaps are read together as a pattern. When several gaps point to the same structural problem, they are grouped and translated into one or more candidate mechanisms. For this reason, the mechanism emerges from the gap analysis, but not through a one-to-one correspondence. It emerges from the combination, weighting and interpretation of the most relevant gaps across dimensions.
The proposed mechanisms remain anchored to the PI Norm core. Most of them are expected to relate primarily to one of three areas: process standardisation, interoperable information flows, or governance and compliance. Other dimensions, such as market conditions, organisational readiness, or sustainability performance, influence the urgency or feasibility of the mechanism, but they do not necessarily generate a separate mechanism on their own.
This final step allows the method to move from diagnosis to action. At case level, it supports the identification of transition packages that address the main constraints of a specific innovation. Across cases, it creates the basis for a reusable library of mechanisms that can also be tested in logistics settings beyond the original application context.

Step	Methodological sequence	Main output
1. TO-BE analysis	Dimensions → Prospect state (maturity state) for each dimension → Target scenario	Target Scenario
2. AS-IS assessment	Dimension → Relevance → Checklist → Maturity state selection → Actual State	Actual State
3. Gap analysis	TO-BE vs AS-IS → Gap per dimension → Weighted gap → Prioritisation	Prioritised gap(s)
4. Mechanism identification	Prioritised gaps → Gap patterns → Candidate scaling mechanisms	Mechanism shortlist

Table 1 – Four steps methodology

3 Analysis

This section examines how the methodological framework was operationalised into a spreadsheet-based assessment tool designed to support the application of the method across different LIs. The tool translates the four methodological steps into a common analytical structure and makes it possible to define target conditions, assess current maturity, calculate prioritised gaps and prepare the identification of scaling mechanisms in a consistent way. In this sense, the analysis does not focus on the outcomes of individual case applications, but on the structure of the tool itself and on the way the methodological logic was embedded into it.
The chapter follows the same sequence introduced in the methodology.

3.1 Step 1: Operationalisation of the TO-BE analysis

In the tool, the TO-BE analysis is operationalised through the selection of one target maturity state for each of the six scaling dimensions. Each selected state is converted into a score from 1 to 4, and the resulting combination defines the target scenario of the LI.

The target scenario is not selected directly by the user. It is derived automatically from the TO-BE profile through a set of predefined scenario rules. These rules combine three elements: threshold values in specific dimensions, the average of the six target scores and the number of dimensions with a score of at least 3. As shown in Figure 2, this rule-based structure is used to determine whether the selected maturity profile corresponds to one of the five target scenarios: Baseline, Collaborative, Accelerated, Disruptive, or Transformative.

Scenario rules						
ScenarioID	Metric	Operator	Value1	Value2	Comment	Hint
TRANS	Sustainability and circularity	=	4		Transformative requires systemic sustainability (Systemic level)	Set Sustainability = 4
TRANS	Count(>=3)	>=	5		Transformative requires broad maturity: at least 5 dimensions scored ≥3	Reach COUNT(>=3) ≥ 5
TRANS	Average	>=	3,2		Transformative requires high overall maturity (average score threshold)	Raise Average ≥ 3.2
DISR	Technological	=	4		Disruptive requires cutting-edge technology (level 4)	Set Technological = 4
DISR	Regulatory, policy and governance	=	4		Disruptive requires standardised & scalable governance/compliance (level 4)	Set Regulatory, policy and governance = 4
DISR	Average	>=	3		Disruptive also requires overall maturity above a minimum threshold	Raise Average ≥ 3.0
ACC	Technological	>=	3		Accelerated requires advanced interoperability/compliance-ready tech (≥3)	Raise Technological ≥ 3
ACC	Sustainability and circularity	>=	3		Accelerated requires advanced sustainability readiness (≥3)	Raise Sustainability ≥ 3
ACC	Operational	>=	3		Accelerated implies network-level operations (not local pilots)	Raise Operational ≥ 3
ACC	Average	BETWEEN	2,7	3,4	Accelerated fits mid-high maturity (keeps it below fully transformative levels)	Adjust Average to 2.7–3.4
COLL	Societal and organisational	>=	3		Collaborative requires coordinated organisational adoption (≥3)	Raise Societal and organisational ≥ 3
COLL	Market and economic	>=	3		Collaborative assumes stable growth / repeatable demand (≥3)	Raise Market and economic ≥ 3
COLL	Average	BETWEEN	2,2	3	Collaborative fits low-mid maturity (above baseline, below accelerated)	Adjust Average to 2.2–3.0
BASE	Average	<=	2,2		Baseline corresponds to early-stage/limited maturity overall	Keep Average ≤ 2.2

Figure 2 – Scenario Rules

To avoid ambiguous classification, the scenarios are also ordered by priority, from Transformative to Baseline (1 to 5). If more than one scenario satisfies the rule set, the tool assigns the one with the highest priority. If no scenario is fully matched, the tool returns a near-match result and indicates the condition that is still missing (Hint). This makes the TO-BE analysis more robust, because the scenario label is always tied to the actual maturity configuration selected for the innovation.

Resulting Scenario:	Comment:
Baseline	The innovation remains confined to pilots or a limited corridor. Integration across actors and processes is still partial, and adoption is slow and uneven. Sustainability actions exist but are not yet embedded end-to-end.

Resulting Scenario:
Near match (Baseline -> Collaborative): Raise Societal and organisational ≥ 3

Figure 3 – Examples of Resulting Scenario

3.2 Step 2: Operationalisation of the AS-IS assessment

The second step operationalises the AS-IS assessment of each LI. In the tool, this is done through two linked actions for each of the six dimensions: the selection of a relevance level and the identification of the Current State.

First, the LI assigns a relevance level to each dimension. The available options are Low, Medium and High. This choice does not describe maturity. It expresses how important that dimension is for the scale-up path of the specific innovation. The result is a first weighting layer that is later used in the gap calculation.

Second, the LI identifies its Current State through a dimension-specific checklist. The checklist is not structured as a generic questionnaire. For each dimension, the questions are organised around the same four maturity states used in the TO-BE analysis. In practical terms, the LI reads the checklist statements associated with the four maturity states and selects the state that best reflects its actual situation. This is the checklist example for LI1, Operational dimension:

- Isolated
  - Is the service currently limited to a single pilot flow or a single corridor setup with minimal replication?

    - Are onboarding and booking handled mainly case-by-case (e.g., via emails/calls) rather than through a standard process?
    - Is coordination between shippers, MTOs and LSPs mostly bilateral, without a shared operational cycle?
- Local/Regional
    - Are there repeatable procedures for onboarding and booking, but only with a small set of partners or within one corridor context?
    - Is the corridor service stable for a limited number of flows/volumes, but difficult to extend to additional terminals/legs?
    - Are exceptions (delays, capacity changes) managed with partial standardisation but still heavy manual effort?
- Network-integrated
    - Are LI1 processes (onboarding, booking, data sharing, procurement/tendering) standardised and used across multiple partners in the corridor?
    - Is volume pooling coordinated through a shared operating model (trustee roles + agreed interfaces between shippers, MTOs and LSPs)?
    - Can new shippers/LSPs be onboarded using a defined procedure with predictable steps and responsibilities?
- Scaled & optimised
    - Is the operating model designed to replicate to additional corridors with limited rework (same process cycle and templates)?
    - Are performance KPIs tracked (e.g., lead time, frequency, capacity utilisation, empty trips) and used to improve operations?
    - Is re-planning (capacity changes, disruptions) handled through structured rules and routines rather than ad-hoc coordination?

The selected label is then converted into a numerical score from 1 to 4. In this way, the tool produces an AS-IS maturity profile that is directly comparable with the TO-BE profile defined in Step 1.
As a result, Step 2 produces two outputs for each dimension: a relevance weight and a Current State score. Together, these provide the basis for the weighted gap analysis developed in the next step.

3.3 Step 3: Operationalisation of the weighted gap and priority ranking

The third step translates the comparison between the TO-BE and AS-IS profiles into a prioritised diagnostic output. For each dimension, the tool first calculates the gap as the positive difference between the target score and the current score. If the current score is equal to or higher than the target score, the gap is set to zero. This ensures that only the maturity distance that still needs to be closed is carried forward in the analysis.
The gap is then multiplied by the relevance weight assigned in Step 2. The result is the weighted gap, which is the core output of this stage. Its role is to combine two pieces of information: the size of the maturity distance and the strategic importance of the dimension for the scale-up path of the specific innovation. In this way, the tool does not only show where the innovation is distant from the target, but where that distance matters most.

A first visual representation of this comparison is provided by the radar chart in Figure 4, which shows the relative position of the TO-BE and AS-IS profiles across the six dimensions.

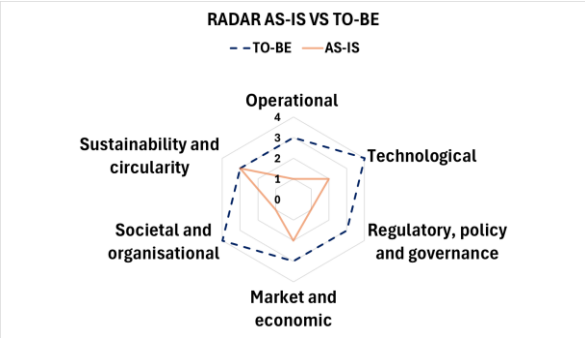


Figure 4 – Radar representation of GAP Analysis

However, this visual comparison is only the first level of interpretation. The actual prioritisation is carried out through the weighted gap logic. To make the output directly usable, the tool ranks the six dimensions by combining two elements: the size of the gap and the relevance assigned to each dimension. In this way, the ranking does not simply show where a maturity distance exists, but where that distance is most critical for the scale-up path of the specific LI.

	STEP 1 - DEFINITION OF TARGET SCENARIO (TO BE)		STEP 2 - DEFINITION OF AS-IS				STEP 3 - GAP ANALYSIS
Scaling Dimensions	Prospect State (To Be)	TO-BE	Relevance for LI	Checklist	Current State	AS-IS	Comment
Operational	Network-integrated	3	High	**Isolated** •Is the service currently limited to a single pilot flow or a single corridor setup with minimal replication? •Are onboarding and booking handled mainly case-by-case (e.g., via emails/calls) rather than through a standard process? •Is coordination between shippers, MTOs and LSPs mostly bilateral, without a shared operational cycle? **Local/Regional** •Are there repeatable procedures for onboarding and booking, but only with a small set of partners or within one corridor context? •Is the corridor service stable for a limited number of flows/volumes, but difficult to extend to additional terminals/legs? •Are exceptions (delays, capacity changes) managed with partial standardisation but still heavy manual effort? **Network-integrated** •Are LI1 processes (onboarding, booking, data sharing, procurement/tendering) standardised and used across multiple partners in the corridor?	Isolated	1	Not aligned - Gap is 2 - High priority (2)

Figure 5 – Example of gap analysis and priority ranking for the Operational dimension of LI1

3.4 Step 4: Translation of prioritised gaps into candidate scaling mechanisms

In the final step, the prioritised gaps are translated into candidate scaling mechanisms. Rather than linking one mechanism to each individual gap, the logic is to read the most relevant gaps together and interpret them as a gap pattern. This makes it possible to move from a dimension-by-dimension diagnosis to a more usable transition logic.

Table 2 illustrates this translation logic by showing how recurrent combinations of prioritised gaps can be associated with candidate scaling mechanisms. At this stage, the table should be read as a first structured draft of the mechanism library. Its purpose is to indicate the direction of the translation logic and to provide an initial basis for consistent interpretation across cases. Further refinements are expected as the assessment is progressively applied to the different LIs and discussed with the pilot actors. The table is therefore intended as an evolving analytical component that will be updated, adjusted, and expanded as the empirical understanding of recurrent gap patterns becomes more robust.

Prioritised gap pattern	Main issue identified	Candidate scaling mechanism	Primary PI Norm anchor
Operational + Regulatory, policy and governance	The service works in a pilot setting but is not yet formalised as a repeatable operating model	Standard operating model package	Process standardisation
Technological + Regulatory, policy and governance	Data exchange exists, but it is not yet robust enough for scalable and compliant deployment	Interoperable data-sharing and compliance package	Interoperable flows
Operational + Technological + Regulatory, policy and governance	The innovation lacks an integrated structure combining repeatable processes, interoperable flows and formal coordination	Integrated scaling package	Governance and compliance
Market and economic + Societal and organisational	The innovation is technically viable, but adoption remains weak and actor participation is not yet stabilised	Onboarding and adoption package	Governance and compliance
Sustainability and circularity + Operational	Sustainability goals are present, but they are not yet embedded in routine service execution	Sustainability integration package	Process standardisation

Table 2 – Example of Gap patterns and candidate scaling mechanisms

4 Results

At the submission stage of this paper the main result of this work is the operationalisation of the proposed methodology into a structured assessment tool that makes the scale-up logic directly usable. The tool translates the conceptual framework into a stepwise procedure that supports target definition, baseline assessment, gap prioritisation, and the first identification of candidate scaling mechanisms.

A first result is the construction of a multidimensional maturity structure based on six scaling dimensions, each expressed through four ordered maturity states. This makes it possible to describe each LI through a comparable TO-BE and AS-IS profile.

A second result is the introduction of a scenario-based target logic, in which the target scenario is not selected manually but derived from the maturity configuration chosen for the TO-BE analysis. This makes the target-setting phase more disciplined and internally coherent.

A third result is the implementation of a checklist-based AS-IS assessment linked to the same maturity states used in the TO-BE analysis. This creates direct comparability between current and target conditions.

A fourth result is the development of a weighted gap analysis, in which the maturity distance is combined with the relevance assigned to each dimension. This allows the tool to rank the dimensions and to highlight where intervention is most urgent for the specific LI.

A fifth result is the definition of an initial translation logic from prioritised gaps to candidate scaling mechanisms. At this stage, this translation is expressed through a first draft mechanism table that links recurrent gap patterns to possible transition packages. Taken together, these elements provide a coherent basis for a repeatable assessment process and for the progressive consolidation of a PI-oriented mechanism library.

5 Discussion and Conclusions

The main contribution of this work lies in the translation of broad roadmap principles into an operational assessment logic. Existing roadmap documents indicate the direction of transition, but they do not usually provide a structured procedure to assess where a LI currently stands, define a realistic target condition, measure the distance to that target and connect that distance to possible scaling actions. The framework proposed in this paper addresses this gap by linking PI-oriented target setting, checklist-based maturity assessment, weighted gap analysis and the identification of candidate scaling mechanisms within one coherent structure. In this sense, the contribution is not only conceptual. It is also methodological and operational.

A second relevant point concerns applicability. Although the framework is rooted in the PI Norm logic and in roadmap trajectories developed in the European context, its analytical structure is not limited to one specific pilot or one specific geography. The use of common dimensions, ordered maturity states, relevance weighting and pattern-based mechanism identification makes the approach transferable to different logistics settings. What may vary across applications is the content of the checklist or the interpretation of certain governance conditions, but the core assessment logic can remain stable. This is important because it means that the framework can support comparison across heterogeneous innovations while still preserving case sensitivity.

At the same time, the current contribution should be read as the formalisation of an assessment framework and of its first operational logic. The mechanism library introduced in this paper should therefore be understood as an initial structured basis rather than as a final and closed taxonomy. Its further refinement will depend on repeated application across different LIs and on the progressive consolidation of recurring gap patterns. This does not reduce the value of the current work. On the contrary, it shows that the framework has been designed not as a one-off exercise, but as a reusable structure that can become more robust through application.

Overall, the paper shows that scaling can be analysed in a more disciplined way. By doing so, the framework provides a practical bridge between PI-oriented transition principles and the concrete assessment of LIs. This is also the basis on which a more mature and reusable library of scaling mechanisms can be developed in future applications.

Disclaimer: "IKIGAI is funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them."

6 References

- ALICE-ETP (2020): Roadmap to the Physical Internet, Alliance for Logistics Innovation through Collaboration in Europe, Brussels.
- Ballot E., Montreuil B., Meller R. D. (2014): The Physical Internet, La Documentation Française, Paris.
- Crainic T. G., Montreuil B. (2016): Physical Internet Enabled Hyperconnected City Logistics, Transportation Research Procedia, 12, 383–398.
- Hesse M., Rodrigue J.-P. (2004): The Transport Geography of Logistics and Freight Distribution, Journal of Transport Geography, 12(3), 171–184.
- McKinnon A. (2018): Decarbonizing Logistics: Distributing Goods in a Low-Carbon World, Kogan Page, London.
- Montreuil B. (2011): Towards a Physical Internet: Meeting the Global Logistics Sustainability Grand Challenge, Logistics Research, 3(2–3), 71–87.
- ALICE-ETP (2019): Roadmap Towards Zero Emissions Logistics 2050, Alliance for Logistics Innovation through Collaboration in Europe, Brussels.

Shipment Planning Protocols in Cyber-Physical Internet for B2C Cross-Border E-commerce Logistics

Xinyue Xie, Xuegang Li, Bo Liu, Zhiheng Zhao and George Q. Huang

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026
Bordeaux, FRANCE

The Hong Kong Polytechnic University
香港理工大學

INTRODUCTION

We study the shipment planning problem in locker-based last-mile delivery for B2C cross-border e-commerce logistics: how orders are consolidated into dispatchable shipments and when they are dispatched.

This planning problem is challenged by (1) obstructed information sharing and collaboration, (2) uncertainties such as dynamic order arrivals and time-varying locker available capacity, and (3) order fragmentation and growing demand for timeliness.

PI and CPI provide a promising foundation, but CPI-enabled shipment planning protocols for locker-based B2C cross-border e-commerce logistics remain underexplored.

OBJECTIVES

This work aims to develop a CPI framework for locker-based B2C cross-border e-commerce logistics planning problem by addressing the following research questions:

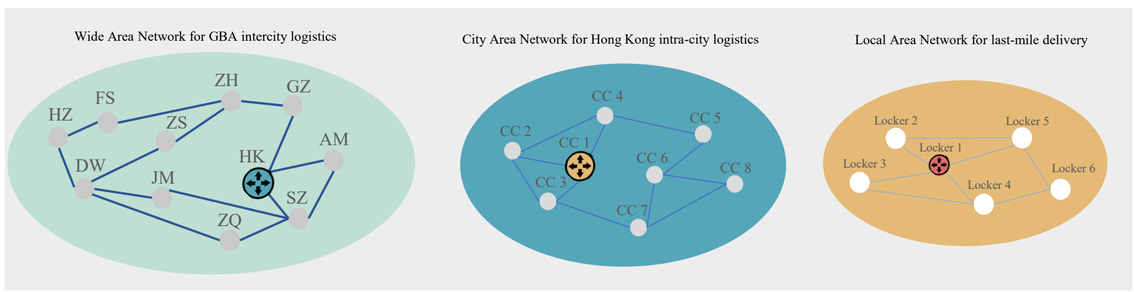
- How can the CPI framework be adapted to model B2C cross-border e-commerce logistics in a way that supports shipment planning under obstructed information sharing and collaboration?
- How can real-time planning-relevant information be inquired and represented through protocol mechanisms in CPI?
- How can CPI shipment planning protocols be designed to transform fragmented orders into feasible and efficient dispatchable shipments under dynamic demand arrivals, time-varying locker available capacity and limited fleet resources?

METHODOLOGY

- CPI Network Configuration

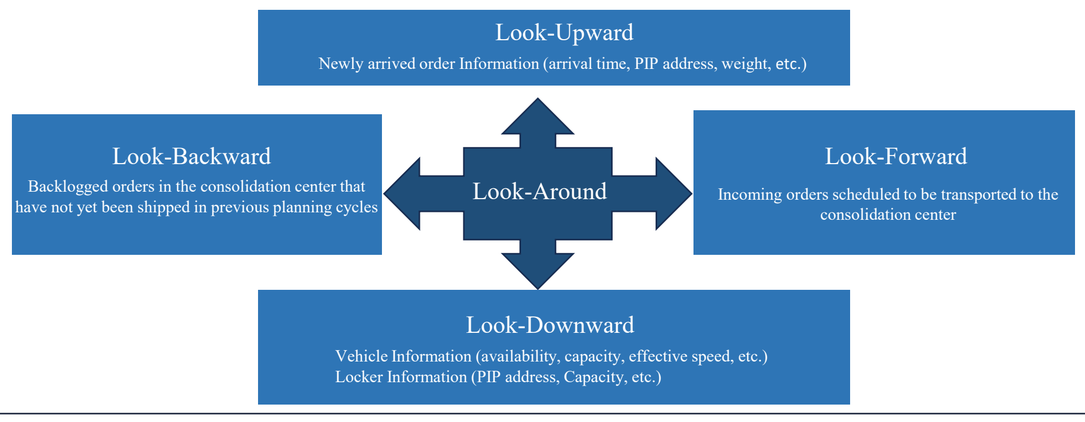
A three-layer CPI network: WAN (intercity) → CAN (intra-city consolidation centers) → LAN (last-mile to lockers).

Protocol Shipment Unit (PSU) is introduced as the standardized dispatchable shipment unit.



- Real-Time "Look-Around" Mechanism

Invoked at each planning epoch to inquire real-time planning-relevant information from CPI routers.



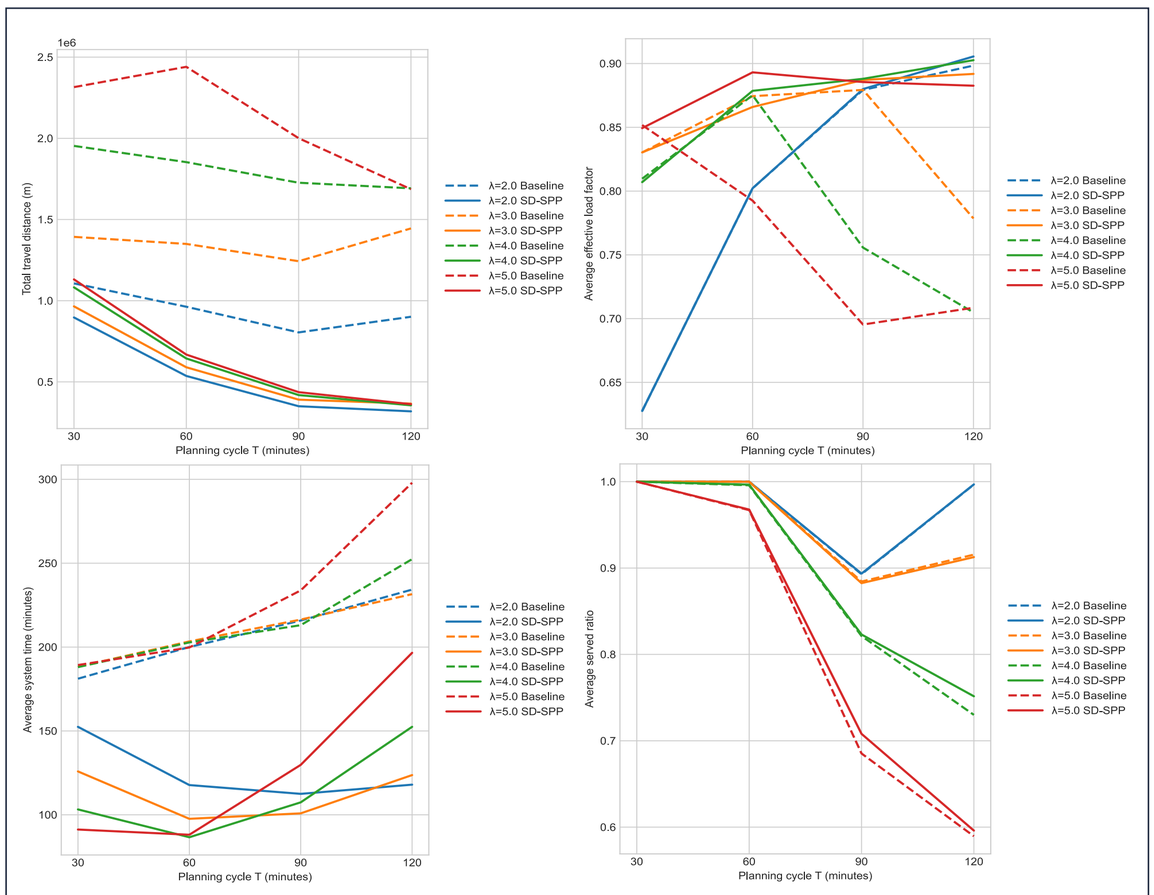
- Standardized Information Tables:

6 tables (Current Order, Locker, Vehicle, Backlog Order, Locker Aggregation, PSU) serve as standardized information representation in CPI routers.

- Spatial-Delay Shipment Planning Protocol (SD-SPP)

The protocol transforms the currently available orders into dispatchable PSUs through four main steps: locker-based order aggregation and locker-capacity-aware screening, vehicle-capacity-feasible splitting, spatial clustering, and PSU generation with vehicle assignment.

RESULTS



Baseline: Periodic FIFO dispatch without CPI-enabled information.

- Distance ↓ 19%–79%: SD-SPP consistently outperforms the baseline approach, with larger gains under longer planning cycles and higher demand.
- Effective load ↑: Capacity-aware screening avoids infeasible deliveries and improves utilization.
- System time ↓ with comparable service: Faster fulfillment without sacrificing served ratio.

CONCLUSION

- Extended CPI paradigm to locker-based B2C cross-border e-commerce logistics with multi-layer router-based architecture.
- Developed shipment planning protocols integrating Real-Time "Look-Around" Mechanism, Standardized Tables, and SD-SPP for real-time, locker-capacity-aware PSU generation.
- Demonstrated that CPI-enabled information access leads to substantial improvements in distance, load utilization, and service performance.

Future work will extend the framework to model within-cycle state dynamics, and validate the approach with richer real-world data.

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE

ISO 14083 Guidance for Comparable CO_2e Reporting in Urban Road Freight Operations

Dimos Touloumidis[1], Sofoklis Dais[1] and Georgia Ayfantopoulou[1]

1. Centre for Research and Technology Hellas (CERTH) – Hellenic Institute of Transport (HIT), Thessaloniki, Greece

Corresponding author: dtouloumidis@certh.gr

Abstract: *This study develops a framework for calculating CO_2e emissions from urban freight operations performed by light commercial vehicles and heavy goods vehicles. The framework is aligned with ISO 14083 and the GLEC Framework and is restricted to direct vehicle activity on urban roads. Two operational states are considered, namely vehicle movement between service points and stop-related stationary vehicle operation during dwell. Adjacent logistics processes, including loading/unloading, cargo-handling equipment, depot energy use, warehousing and non-propulsion auxiliary equipment, are acknowledged but excluded in order to preserve a clear and auditable system boundary. The methodology combines a three-tier data hierarchy with procedure-specific calculations that distinguish vehicle use emissions from vehicle energy provision emissions and support application under different levels of data availability. Measured fuel or electricity use is prioritised, whereas modelled parameters and default emission intensities are used when primary data are unavailable. The proposed structure supports consistent reporting at round, vehicle and fleet level and differentiates between LCV and HGV operations. It is concluded that a narrowly defined and ISO 14083-consistent framework can improve the comparability and transparency of urban freight emissions assessment, while future research should focus on empirical validation, refinement of dwell-related modelling and gradual expansion toward adjacent processes.*

Keywords: *Urban Logistics Decarbonization, ISO 14083, Physical Internet, CO2 Emissions Toolbox, Urban Logistics, Green Routing.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☒ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☒ *PI Networks,* ☐ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☒ *Logistics/Commercial Data Platform,* ☒ *Access and Adoption,* ☒ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

1 Introduction and gap

Urban freight is essential for the economic and social functioning of cities, yet it also intensifies congestion, energy use and emissions because deliveries are performed in dense street networks through repeated stop-and-go movements, short trips, curbside dwell and parking-related frictions. Recent European-oriented research shows that urban freight generates environmental pressures disproportionate to its traffic share, which is why city logistics has increasingly been treated as a core sustainability and planning issue rather than solely as an operational one (Russo & Comi, 2016; Pirra et al., 2022; Ghizzawi et al., 2024).

Over the last decade, the literature has advanced in several relevant directions. Studies of tested city-logistics measures in European cities have examined the environmental effects of infrastructure, regulation, consolidation and governance interventions, while also showing that outcomes remain highly context-dependent (Russo & Comi, 2016; Galambos et al., 2024). Operational studies based on vehicle traces and urban freight datasets have improved the monitoring of accessibility and delivery activity, but they also reveal a persistent methodological problem, namely that urban freight data are often scarce, heterogeneous, outdated and costly to obtain, which limits comparability across operators and cities (Pirra et al., 2022; Dablanc & Adoue, 2025). Broader reviews on sustainable road-freight transitions emphasise electrification, routing and digitalization, but also indicate that the field remains fragmented and methodologically uneven at the level of practical emissions accounting for routine urban operations (Alarcón et al., 2023; Ghizzawi et al., 2024).

In parallel, emissions-accounting practice has been converging toward common standards. ISO 14083 provides requirements and guidance for the quantification and reporting of greenhouse gas emissions arising from transport chain operations, whereas the Global Logistics Emissions Council (GLEC) Framework v3 translates these principles into an operational well-to-wheel approach for transport and logistics activities. Recent comparative work further suggests that the most reliable freight-emission tools are those aligned with ISO 14083 and GLEC (ISO, 2023; Smart Freight Centre, 2025; Olivari et al., 2025). Nevertheless, a gap remains in urban road logistics, as relatively limited research has translated these standards into a compact and auditable method for routine van and truck rounds within cities, especially when the analytical focus is restricted to vehicle movement and stop-related operation rather than broader logistics and facility processes (Royo et al., 2025).

Against this background, the objective of the paper is to develop an ISO 14083-aligned framework for calculating CO_2e emissions from regular urban road freight operations performed by vans and trucks within cities. The scope is intentionally narrow and covers only two operational states, namely vehicle movement on the urban network and stop-related vehicle operation during dwell. On this basis, the paper addresses the following research questions:

- RQ1. What is the appropriate system boundary for calculating emissions from routine urban delivery rounds performed by vans and trucks?
- RQ2. How can ISO 14083 and GLEC principles be translated into a practical and auditable calculation method for movement and stop-related emissions under different levels of data availability?
- RQ3. Which activity data and reporting indicators are most suitable for comparable emissions reporting at round, vehicle and fleet level?

To address these questions, the paper develops a methodological framework rather than a full case-specific emissions inventory. Its contribution lies in defining a narrow but transparent accounting structure for routine urban van and truck rounds, with explicit system boundaries, reproducible calculation procedures and a documented hierarchy of data sources. In this way, the study seeks to improve comparability and auditability in urban freight CO_2e reporting while remaining consistent with ISO 14083 and GLEC principles. The framework does not replace ISO 14083 or GLEC, but operationalises their accounting logic for a specific urban road-freight use case by defining the reporting object, the movement and dwell-related calculation procedures, and the data hierarchy required when primary operational data are incomplete.

2 Study scope and boundaries

The framework adopts an ISO-consistent boundary by defining a specific transport operation and separating included from excluded processes. The unit of analysis is the urban delivery round performed by a van or truck within the city. In the present study, the urban delivery round is the reporting object and is treated either as a documented transport operation or as a documented aggregation of transport operations performed by one vehicle within a bounded service cycle. The scope is limited to the direct road operation of the vehicle during that round and therefore covers two analytical operational states only, namely movement on the urban network and stop-related stationary vehicle operation during dwell. These states are introduced as an urban-operational decomposition within an ISO 14083-aligned accounting structure; they are not presented as reporting categories prescribed by ISO 14083 itself.

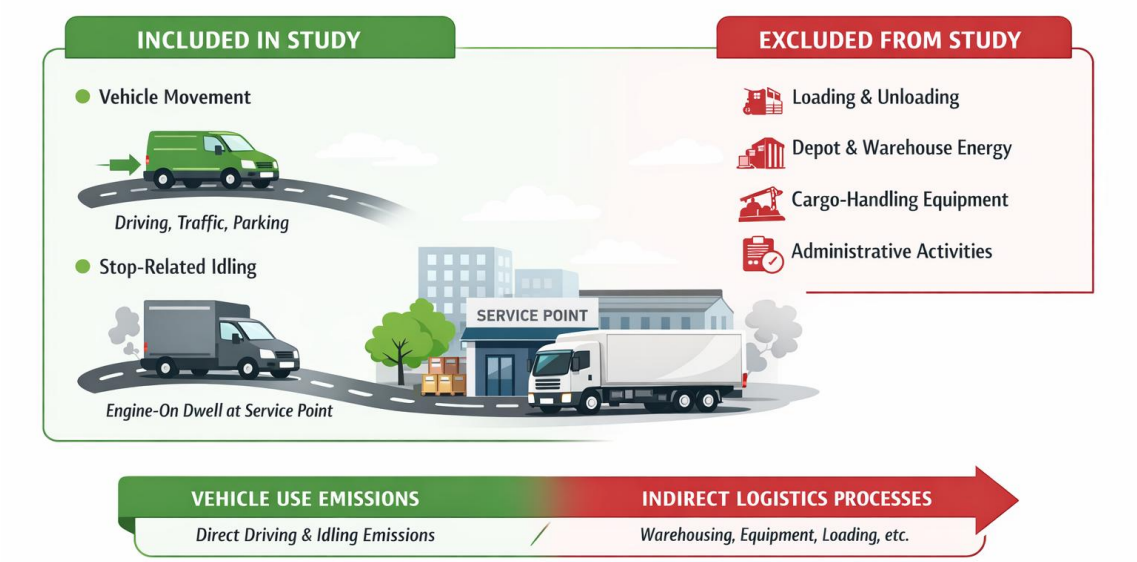


Figure 1: System boundary of the proposed framework, including direct vehicle movement and transport-related engine-on dwell, while excluding indirect logistics processes.

Movement includes all driven activity directly attributable to the round, including circulation between stops, congestion-related delays, curb approach, parking search, detours and empty repositioning. Stop-related operation includes only transport-related engine-on idling during stationary dwell at or immediately adjacent to the delivery point. Engine-off parked dwell is treated as zero-emission stop time within the defined boundary, because non-propulsion auxiliary loads are excluded. Auxiliary equipment and other non-propulsion energy uses, including power take-off (PTO) equipment, refrigerated units, cargo-handling devices and other stationary auxiliary loads, remain outside the boundary.

Several adjacent operations are recognised as relevant to the wider urban freight system but are excluded from the present study. These include loading and unloading, cargo-handling equipment, depot and warehouse energy use, cross-docking, packaging activities, administrative support and infrastructure-related processes. Engine restart and pull-out acceleration after a stop are also acknowledged as real operational effects, but they are excluded because their robust estimation would add disproportionate complexity relative to the compact purpose of the paper. The framework therefore reports vehicle use emissions and vehicle energy provision emissions within a narrow and transparent scope intended to support comparability across urban van and truck rounds.

3 Methodology for urban vehicle operations

3.1 Overall methodology and core principles

The framework quantifies greenhouse gas emissions from the road operation of light commercial vehicles (LCVs) and heavy goods vehicles (HGVs) performing urban delivery rounds within a defined city boundary. The unit of analysis is the urban delivery round, defined as the sequence of service stops undertaken by a single vehicle between departure from and return to the depot, or between two predefined route boundaries.

Two mutually exclusive operational states are distinguished. Movement covers all driven segments between stops and therefore includes circulation, congestion-related delay, curb approach, parking search, detours and empty repositioning directly linked to the round. Stop-related operation covers stationary dwell at or immediately adjacent to the service point and includes engine-on idling for combustion vehicles only. Within the defined boundary, stationary stop emissions for battery-electric vehicles are set to zero because auxiliary stationary energy demand is excluded. Loading and unloading, cargo-handling equipment, depot and warehouse energy use, cross-docking, packaging and infrastructure-related processes are likewise excluded, while engine restart and pull-out acceleration are recognised but omitted in order to preserve methodological simplicity and avoid the introduction of uncertain minor terms that would require a separate modelling layer.

For each operational state p, total greenhouse gas emissions are calculated as the sum of vehicle use emissions and vehicle energy provision emissions:

$$G_{T,p,r} = G_{TO,p,r} + G_{TEP,p,r} \quad (1)$$

where $G_{TO,p,r}$ denotes vehicle use emissions and $G_{TEP,p,r}$ denotes vehicle energy provision emissions for operational state $p$in round r. In the present study, these correspond to tank-to-wheel (TTW) and well-to-tank (WTT) components, respectively.

$$G_{T,r} = G_{T,mov,r} + G_{T,stop,r} \quad (2)$$

For battery-electric vehicles, vehicle use emissions at point of use are zero, whereas vehicle energy provision emissions are derived from electricity consumed during movement. Under the present boundary, stationary dwell is assigned zero emissions.

The framework distinguishes two vehicle categories throughout, namely LCVs, with gross vehicle weight up to 3.5 t, and HGVs, with gross vehicle weight above 3.5 t. This distinction is retained because the two categories differ materially in payload capacity, energy-consumption profile and urban duty cycle. A three-level data hierarchy is applied for each operational state. Tier 1 uses measured fuel or electricity consumption. Tier 2 estimates energy use from telematics, on-board diagnostics or representative operating profiles. Tier 3 uses default total emission intensities only when more specific operational data are unavailable. The selected tier is recorded for transparency and auditability. Missing data are therefore treated through a controlled substitution logic: measured consumption is used where available; otherwise, energy use is modelled from observed operational activity; only in the absence of such activity data are default average intensities used.

3.2 Procedure formulations

3.2.1 Movement on the urban network

The movement procedure covers all driven segments directly attributable to the urban delivery round, including laden and unladen travel between service stops. Where fuel or electricity consumption by energy carrier k is available, movement-related emissions are calculated as:

$$G_{TO,mov,r} = \sum_k \left(Q_{mov,r,k} \times EF_{TTW,k}\right) \quad (3)$$

$$G_{TEP,mov,r} = \sum_k \left(Q_{mov,r,k} \times EF_{WTT,k}\right) \quad (4)$$

$$G_{T,mov,r} = G_{TO,mov,r} + G_{TEP,mov,r} \quad (5)$$

where $Q_{mov,r,k}$ is the quantity of fuel or electricity consumed during movement and $EF_{TTW,k}$ and $EF_{WTT,k}$ are the corresponding vehicle use and vehicle energy provision emission factors.

When direct consumption data are unavailable, energy use is estimated from movement activity:

$$Q_{mov,r,k} = d_{drv,r} \times q_{mov,c,k} \quad (6)$$

where $d_{drv,r}$ is the actual driven distance of round r and $q_{mov,c,k}$ is the representative urban energy-use rate for vehicle class c. This parameter may be derived from telematics, test-cycle data or operator records and may, where data permit, be differentiated by congestion level, payload state or road type. For the calculation of transport activity, a distinction is maintained between actual driven distance and activity distance. Transport activity is defined at shipment level as:

$$TA_r = \sum_j \left(m_j \times d_{act,j}\right) \quad (7)$$

where $m_j$is the mass of shipment j and $d_{act,j}$ is the corresponding activity distance. In ISO-aligned implementation, activity distance should be determined through a documented method appropriate to collection and delivery rounds. In such rounds, activity distance should be determined using shortest feasible distance (SFD) or great-circle distance (GCD) and, where activity distance is not directly available, a documented distance adjustment factor (DAF) may be applied. Where shipment-level data are unavailable, an aggregate approximation may be used strictly as a fallback:

$$TA_r \approx \bar{m}_r \times d_{act,r}^{laden} \quad (8)$$

where m_r is the average payload during laden movement and $d_{act,r}^{laden}$ is the laden activity distance of the round. This approximation is intended for data-poor cases and should not be treated as equivalent to shipment-level allocation. For this reason, the framework is strongest at round, vehicle and fleet reporting level, whereas tonne-kilometre results derived from aggregate approximations should be interpreted with caution.

3.2.2 Stop-related vehicle operation

The stop-related operation procedure covers vehicle emissions generated during stationary dwell at or immediately adjacent to the delivery point. It includes only transport-related engine-on idling during dwell and excludes emissions associated with the physical handling of goods, as well as auxiliary equipment and other non-propulsion stationary energy uses, including PTO equipment, refrigerated units and other service-related loads. Restart effects and pull-out acceleration are acknowledged but omitted in order to reduce complexity. Where stop-related energy use is measurable, emissions are calculated as:

$$G_{TO,stop,r} = \sum_k \left(Q_{stop,r,k} \times EF_{TTW,k}\right) \quad (9)$$

$$G_{TEP,stop,r} = \sum_k \left(Q_{stop,r,k} \times EF_{WTT,k}\right) \quad (10)$$

$$G_{T,stop,r} = G_{TO,stop,r} + G_{TEP,stop,r} \quad (11)$$

where $Q_{stop,r,k}$ is the quantity of fuel consumed during stop-related operation.

When direct measurements are unavailable, stop-related fuel consumption for combustion vehicles is estimated from total engine-on dwell time.

$$Q_{stop,r,k} = FC_{idle,c,k} \times t_{stop,r} \quad (12)$$

where $FC_{idle,c,k}$ is the representative idle fuel-consumption rate for vehicle class c and energy carrier k, and $t_{stop,r}$ is the cumulative engine-on stop time across the round. Engine-off parked dwell is excluded from Eq. (12). Congestion queues, parking searches and curb approach remain part of movement rather than stop-related operation. Within the present boundary, auxiliary electricity demand during dwell is excluded and stationary stop emissions for battery-electric vehicles are therefore set to zero.

4 Data hierarchy and default factors

The framework applies a three-level data hierarchy in order to preserve consistency under different levels of data availability. Tier 1 uses measured fuel or electricity consumption at vehicle, trip or round level and is preferred whenever available. Tier 2 estimates energy use from operational variables, such as driven distance, stop duration and idle fuel-consumption rate, and therefore supports the movement and stop-related procedures defined above. Tier 3 applies default total emission intensities per tonne-kilometre when detailed operational data are

unavailable. The same methodological structure is therefore maintained across LCV and HGV operations, while the level of detail varies with data maturity rather than with a change in accounting logic (ISO, 2023; Smart Freight Centre, 2025).

Tier 3 is retained strictly as a fallback and should not be treated as methodologically equivalent to Tier 1 or Tier 2. Default intensities should therefore be used only as approximate fallback values and, where possible, selected or adjusted so that they remain consistent with the reduced boundary of direct vehicle operation adopted in the present study. For transparency, implementation should report the selected system boundary, the excluded processes, the activity-distance method, the source and reference year of the emission factors, the basis used for electricity factors and the tier applied in each calculation. Where results are produced from Tier 2 or Tier 3 inputs, this should be stated explicitly so that measured outputs can be distinguished from modelled or default estimates (ISO, 2023; Smart Freight Centre, 2025).

Table 1: Example implementation factors and parameters for the proposed urban freight framework

Category	Parameter(s)	Unit	Example implementation value(s)	Application in framework
Diesel emission factors	$EF_{TTW,diesel}$, $EF_{WTT,diesel}$	kg CO_2e/	100% mineral diesel: EF_{TTW}=2.66155, EF_{WTT}=0.62409; standard blended diesel: EF_{TTW}=2.51279, EF_{WTT}=0.61101	Vehicle use emissions and vehicle energy provision emissions for diesel operation; only the factor set matching the actual fue specification should be use
Gasoline emission factors	$EF_{TTW,gasoline}$, $EF_{WTT,gasoline}$	kg CO_2e/	100% mineral petrol/gasoline: EF_{TTW}=2.35372, EF_{WTT}=0.60664; standard blended petrol/gasoline: EF_{TTW}=2.08440, EF_{WTT}=0.58094	Vehicle use emissions and vehicle energy provision emissions for gasoline operation; only the factor set matching the actual fue specification should be use
Electricity emission factors	$EF_{TTW,el}$, $EF_{WTT,el}$	kg CO_2e/kW h	$EF_{TTW,el}$=0; $EF_{WTT,el}$ = country-specific electricity factor based on one documented accounting basis	Vehicle use emissions and vehicle energy provision emissions for battery-electric vehicles
Tier 2 movement parameters	$q_{mov,LCV}$, $q_{mov,HGV}$	l/km or kWh/km	study-specific	Tier 2 estimation of movement-related energy use for LCVs and HGVs
Tier 2 idling parameters	$FC_{idle,LCV}$, $FC_{idle,HGV}$	l/h	LCV: study-specific or operator-specific; HGV: study-specific with 1.7-3.2 l/h used only as an indicative screening range if no better value exist	Tier 2 estimation of stop-related engine-on dwell for combustion vehicles
Tier 3 total urban intensities	$g^{urban}_{T,van}$, $g^{urban}_{T,truck,3.5-7.5}$, $g^{urban}_{T,MGV,7.5-20}$, $g^{urban}_{T,HGV,>20}$	kg CO_2e/tkn	0.840 for van <3.5 t GVW; 0.335 for urban truck 3.5–7.5 t GVW; 0.210 for MGV 7.5–20 t GVW; 0.125 for HGV >20 t GVW	Tier 3 total urban fallback intensities; class-specific values are preferable to a single generic HGV constant

Table 1 presents example implementation values and calibration parameters used to operationalize the framework. These entries should not be interpreted as ISO 14083 or GLEC

standard constants; rather, they are implementable reference values drawn from selected external datasets or from study-specific calibration and should therefore be replaced, regionalized or recalibrated whenever operator-specific, vehicle-specific or country-specific data are available. In the present formulation, liquid-fuel emission factors are illustrated using DESNZ reporting factors, electricity should be reported on one clearly documented country-specific basis, the zero TTW factor for battery-electric vehicles follows point-of-use accounting logic, indicative idling values are included only as screening inputs and Tier 3 urban fallback intensities follow the starting-point values of GLEC v3.1 (Department for Energy Security and Net Zero, 2024a; European Environment Agency, 2024; ISO, 2023; Smart Freight Centre, 2025; U.S. Department of Energy, 2015; U.S. Environmental Protection Agency, 2016)

The values in Table 1 should be treated as implementation defaults rather than fixed constants. In application, they should be updated to the study year, energy carrier and geographic context, especially in the case of electricity, where country-specific factors are more appropriate than generic averages. Tier 3 intensity values are particularly sensitive to load factor, empty running, shipment profile and urban operating conditions and should therefore be interpreted cautiously. In this way, the framework remains stable in structure, while the numerical inputs reflect the operational and regional conditions of the case under study.

5 Discussion and Conclusions

The proposed framework was developed as a compact and operationally bounded method for quantifying CO_2e emissions from routine urban freight rounds performed by LCVs and HGVs. Its main contribution lies in translating the accounting logic of ISO 14083 and the implementation orientation of the GLEC Framework into a form suited to everyday city logistics, while avoiding the broader complexity of full supply-chain assessment. The principal conclusion is that a narrow methodological focus on vehicle movement and stop-related stationary vehicle operation can provide a transparent and auditable basis for urban freight emissions reporting, provided that the distinction between vehicle use emissions and vehicle energy provision emissions is retained and that data-hierarchy rules are applied consistently (ISO, 2023; Smart Freight Centre, 2025). This conclusion is important because recent comparative work on freight emissions tools indicates that the most robust approaches are those aligned with ISO 14083, while recent urban freight research continues to show that city-logistics interventions and sustainability outcomes remain highly context-dependent and difficult to compare across cities and operators (Olivari et al., 2025; Russo & Comi, 2016; Galambos et al., 2024). Compared with broad calculation tools such as EcoTransIT or EcoLogistics, the proposed framework is narrower in scope but more explicit about the urban delivery-round boundary, the separation between movement and engine-on dwell, and the documentation of the data tier used for each calculation.

A second conclusion concerns the value of the selected system boundary. By restricting the method to direct vehicle operations on the urban network, the framework responds to a practical problem repeatedly identified in the literature, namely that urban freight data are scarce, heterogeneous, often outdated and expensive to obtain, which makes comprehensive emissions accounting difficult in practice (Dablanc & Adoue, 2025). At the same time, the framework does not reduce urban freight to distance alone, because it explicitly incorporates stationary dwell-related vehicle operation as a distinct emissions source. This is methodologically relevant, since recent work on commercial-vehicle parking and dwell behaviour has shown that parking search, curb access and delivery-stop conditions are central features of urban freight activity rather than peripheral disturbances (Ghizzawi et al., 2024). In this sense, the framework contributes to the literature by treating repeated movement and stopping patterns as the core

operational structure of urban delivery rounds rather than as a residual element of route execution, while offering a more realistic basis for urban emissions estimation than approaches that rely exclusively on aggregate tonne-kilometre intensities or network-distance proxies (Pirra et al., 2022; Royo et al., 2025).

The framework nevertheless has important limitations. Results based on Tier 1 data should therefore be interpreted as more robust and auditable, whereas Tier 2 and especially Tier 3 results should be treated as estimates whose uncertainty depends on assumptions about energy-use rates, dwell duration, load factor, activity distance and emission-factor selection. The deliberate exclusion of loading and unloading, cargo-handling equipment, depot energy use, warehousing and packaging means that the method does not represent the full carbon footprint of urban logistics operations and should therefore be interpreted as a framework for direct road-vehicle emissions rather than as a complete logistics sustainability assessment. In addition, although the distinction between movement and stop-related operation improves transparency, the quantification of stop emissions remains sensitive to assumptions regarding idling and dwell duration. These aspects are difficult to observe consistently across fleets and cities and, in the absence of detailed telematics, may require simplified proxies. The accuracy of the framework is also strongly influenced by methodological choices already highlighted by ISO 14083, such as the use of primary rather than modelled or default data, the selection and regionalization of emission factors and the treatment of activity distance in collection and delivery rounds through a documented and consistent method (ISO, 2023). Finally, while the distinction between LCV and HGV improves comparability, it still aggregates a wide range of vehicle technologies, payload profiles and operational practices into broad classes, which is consistent with recent reviews showing continuing heterogeneity in indicator selection and assessment methods across city-logistics research (Aung et al., 2025).

Because the present paper develops a methodological framework rather than an empirical emissions inventory, validation and robustness testing are left for future work. Future research should proceed in four directions. The first is empirical validation through real-world pilot application using operator telematics, fuel records and electricity data, so that the performance of Tier 1, Tier 2 and Tier 3 calculations can be compared on the same urban rounds. Where real fleet data are not available, simulation-based applications could be used as an intermediate validation step. The second is sensitivity analysis of the main assumptions, especially load factor, activity distance, dwell duration, idle fuel-consumption rate and electricity emission factors. The third is methodological refinement of dwell-related operation, particularly through better integration of curbside data, stop-resolved dwell behaviour and vehicle-specific idling parameters. The fourth is gradual boundary expansion toward adjacent urban logistics processes such as loading and unloading, depot charging and other interface activities. Such developments would also support scenario testing for electrification, low-emission zones, PI-enabled consolidation, shared hubs and alternative urban delivery designs, which recent literature increasingly treats as central to the decarbonization of road freight and urban logistics (Alarcón et al., 2023; Royo et al., 2025; Schmidt et al., 2025).

Acknowledgement

This work was supported by the European Commission under the “IKIGAI” project (Grant Agreement No. 101202912). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.

References

- Alarcón F.E., A. Mac Cawley, E. Sauma (2023): Electric mobility toward sustainable cities and road-freight logistics: A systematic review and future research directions. Journal of Cleaner Production, v430, 138959, https://doi.org/10.1016/j.jclepro.2023.138959
- Aung T.Z., T. Prapaspongsa, A. Saraceni, D. Cotta, N. Surinkul, S.H. Gheewala (2025): Review of sustainability assessments in city logistics from holistic perspective: critical insights, gaps and future directions. Transportation Research Interdisciplinary Perspectives, v33, 101603, https://doi.org/10.1016/j.trip.2025.101603
- European Environment Agency (2024): Greenhouse gas emission intensity of electricity generation, country level. European Environment Agency, published 25 October 2024, modified 27 June 2025, https://www.eea.europa.eu/en/analysis/indicators/greenhouse-gas-emission-intensity-of-1-1751032678/greenhouse-gas-emission-intensity-of-electricity-generation-country-level, 2026/03/19.
- Dablanc L., F. Adoue (2025): New ways of collecting urban freight traffic data and applications for urban freight policies and research. Case Studies on Transport Policy, v19, 101315, https://doi.org/10.1016/j.cstp.2024.101315
- Department for Energy Security and Net Zero (2024a): Greenhouse gas reporting: conversion factors 2024. GOV.UK, published 8 July 2024, updated 30 October 2024, https://www.gov.uk/government/publications/greenhouse-gas-reporting-conversion-factors-2024, 2026/03/19.
- Galambos K.J., A.B. Palomino-Hernández, V.C. Hemmelmayr, B. Turan (2024): Sustainability initiatives in urban freight transportation in Europe. Transportation Research Interdisciplinary Perspectives, v23, 101013, https://doi.org/10.1016/j.trip.2023.101013
- Ghizzawi F., A. Galal, M.J. Roorda (2024): Modelling parking behaviour of commercial vehicles: a scoping review. Transport Reviews, v44, no4, 743-765, https://doi.org/10.1080/01441647.2024.2305202
- ISO (2023): Greenhouse gases — Quantification and reporting of greenhouse gas emissions arising from transport chain operations. ISO 14083:2023, 1st edition, International Organization for Standardization, Geneva, Switzerland.
- Olivari E., C. Caballini, X. Lluch (2025): How to calculate GHG emissions in freight transport? A review of the main existing online tools. Case Studies on Transport Policy, v19, 101343, https://doi.org/10.1016/j.cstp.2024.101343
- Pirra M., A. Carboni, F. Deflorio (2022): Freight delivery services in urban areas: Monitoring accessibility from vehicle traces and road network modelling. Research in Transportation Business & Management, v45, 100680, https://doi.org/10.1016/j.rtbm.2021.100680
- Royo B., Y. Zhang, A. Lewis, P. Dehdari (2025): CO2 emissions in urban freight transport: Developing and testing the EcoLogistics tool. Case Studies on Transport Policy, v20, 101415, https://doi.org/10.1016/j.cstp.2025.101415
- Russo F., A. Comi (2016): Urban Freight Transport Planning towards Green Goals: Synthetic Environmental Evidence from Tested Results. Sustainability, v8, no4, 381, https://doi.org/10.3390/su8040381
- Schmidt V., F. Aschauer, A. Scott, M. Gerschberger (2025): Does a detailed freight emissions calculation really matter? Transportation Research Part D: Transport and Environment, v148, 104761, https://doi.org/10.1016/j.trd.2025.104761
- Smart Freight Centre (2025): Global Logistics Emissions Council Framework for Logistics Emissions Accounting and Reporting, v3.1 edition, revised and updated February 2025, Smart Freight Centre, Amsterdam, The Netherlands.
- U.S. Department of Energy (2015): Fact #861 February 23, 2015: Idle Fuel Consumption for Selected Gasoline and Diesel Vehicles. Department of Energy, 23 February 2015, https://www.energy.gov/cmei/vehicles/fact-861-february-23-2015-idle-fuel-consumption-selected-gasoline-and-diesel-vehicles, 2026/03/19.
- U.S. EPA (2016): SmartWay Transport Partnership A Glance at Clean Freight Strategies: Idle Reduction, June 2016, https://19january2017snapshot.epa.gov/sites/production/files/2016-06/documents/420f16025.pdf, 2026/03/19.

Look Around Cyber-Physical Internet Port Operations

Yi You and George Q. Huang

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026
Bordeaux, FRANCE

The Hong Kong Polytechnic University
香港理工大學

Introduction

Multimodal transport refers to the delivery of goods to their destination through at least two different modes of transport. Ports serve as critical transfer nodes connecting domestic and international transport networks. As shown in Figure 1, a port typically connects multiple transport modes, including road, rail, air, and sea transport. Among them, this paper focuses on transshipment operations between sea and rail transport. In a multimodal port, sea transport can use both direct and indirect transshipment when connecting with rail transport, whereas transshipment between sea transport and other modes depend almost entirely on indirect transshipment.

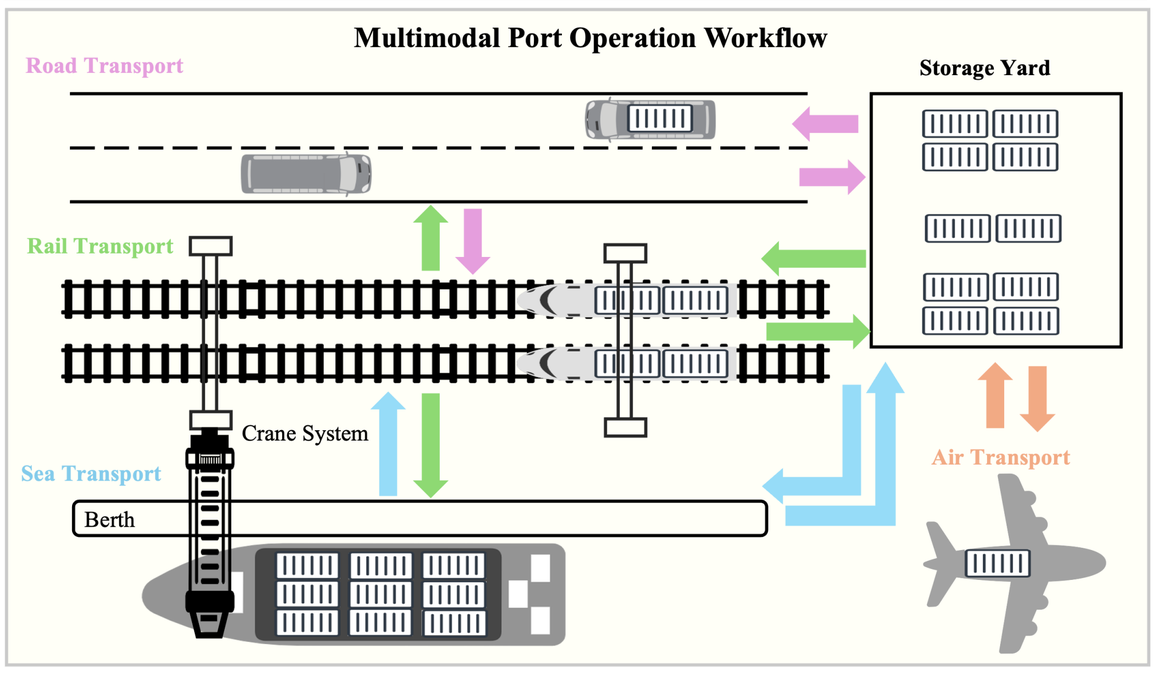


Figure 1 Multimodal port operations

The research objectives of this study are twofold:
(1) To develop a rolling-horizon dynamic decision-making framework that covers the full port operation process from sea transport to rail transport.
(2) To formulate a mathematical model and design a dedicated solution algorithm for a new sea–rail transshipment mode in which trains serve as mobile temporary yards.

Methodology

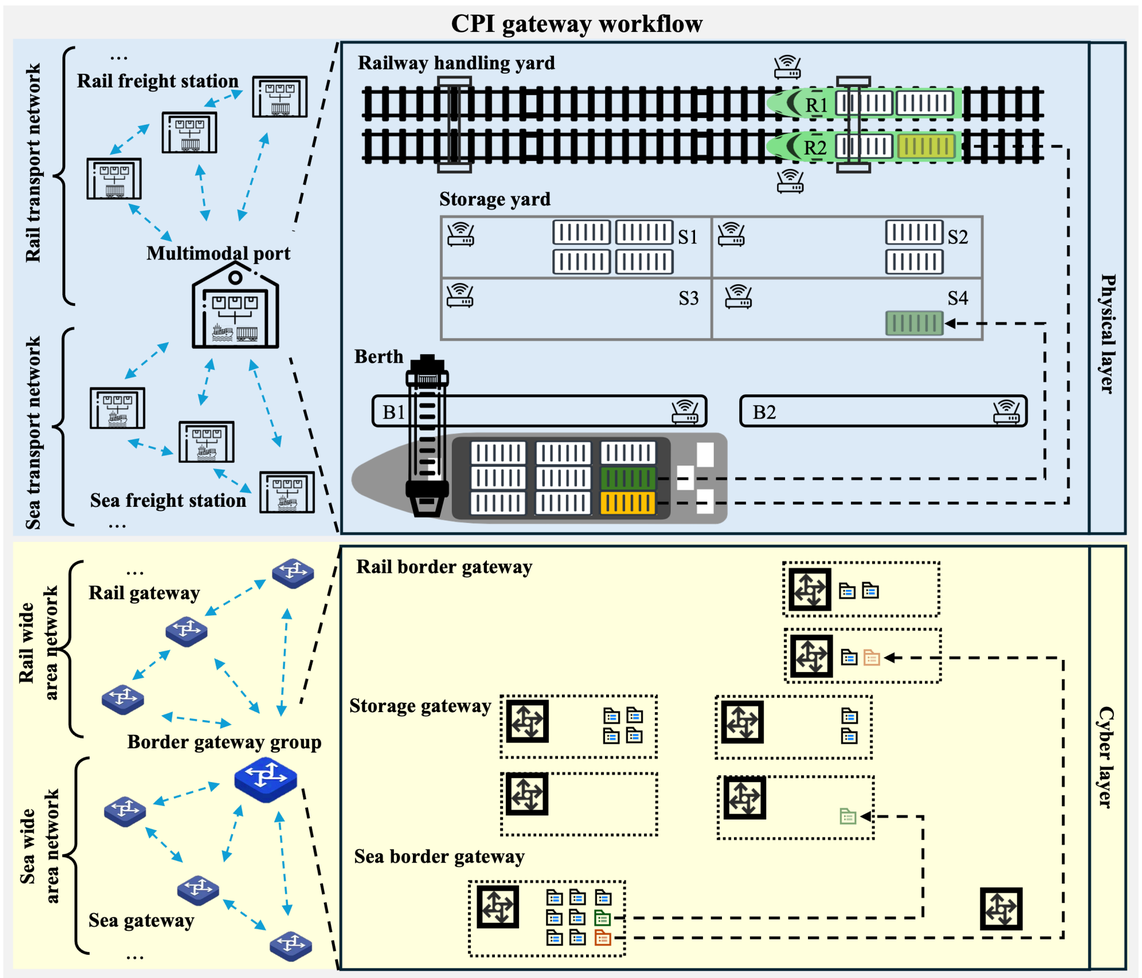


Figure 2 Deployment of CPI gateways in multimodal port

Methodology

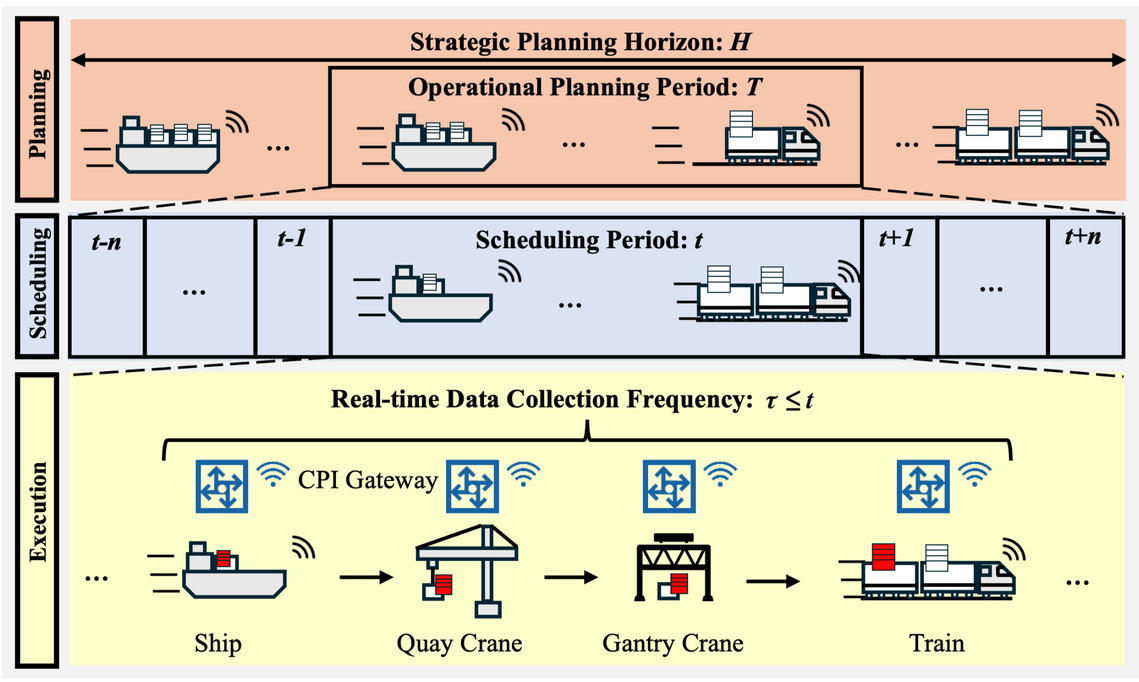


Figure 3 Look Around dynamic decision-making framework

1. To ensure real-time data accessibility and interoperability across heterogeneous transport networks, and establish subsequent dynamic decision-making frameworks, it is imperative to deploy smart devices such as CPI gateways in critical zones of smart ports, as shown in Figure 2.
2. As shown in Figure 3, the CPI Look Around rolling decision framework organizes decision-making information across layers and over time. The transshipment process consists of three stages: Planning, Scheduling, and Execution. In the Execution stage, the data collection frequency is denoted as r, with $r \leq t$, enabling real-time decision-making.
3. To facilitate understanding of container transitions from sea transport to rail transport, we illustrate the spatiotemporal positions of containers within an Operational Planning period T, as shown in Figure 4.

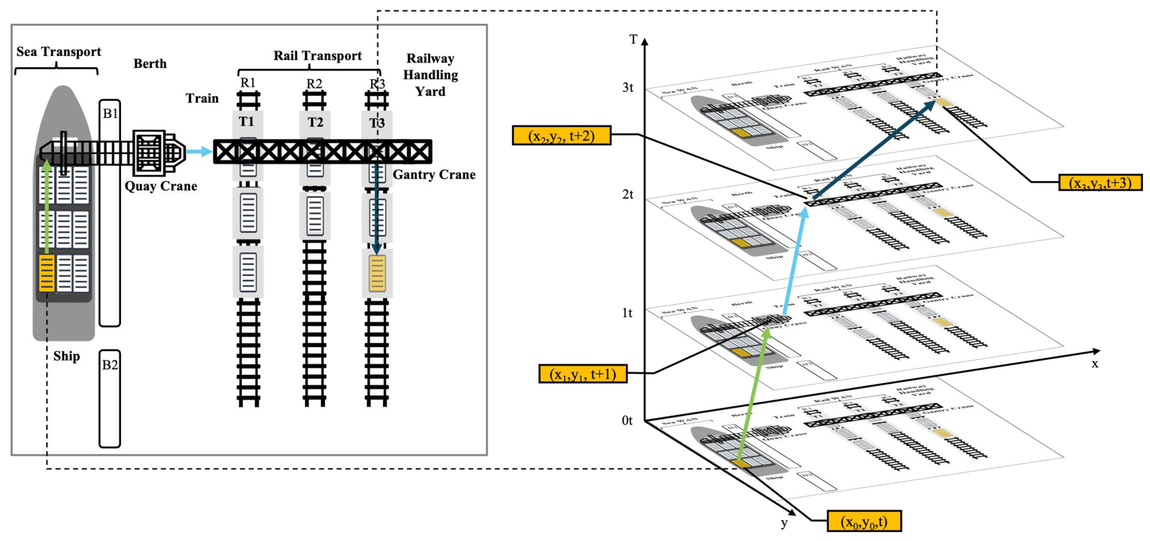


Figure 4 Direct transshipment considering spatiotemporal dimensions

Conclusion

We evaluate the framework through case studies built within the CPI five-layer setting. Both small-scale and large-scale scenarios are tested. We compare our solution algorithm with non-black-box baselines used in practice and in structured optimization, including fixed-rule manual strategies and dynamic programming method. Results show that our algorithm reduces waiting on both ship and train sides. It also cuts yard usage and unnecessary container reshuffling. Resource utilization improves for cranes, tracks, and train capacity. The advantage becomes clearer when operations are busy and constraints are tight. In addition, the rolling decision cycle meets real-time response needs, which supports continuous execution rather than offline replanning.

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE
Fraunhofer IML

End-to-End Supply Chain Management in Plant Construction: A Simulation-Based, Data Space-Ready Approach

Lucas Schreiber[1], Susanne Klöcker[1], Till Lemmer[1], Michael Dominik Görtz[1] and Julia Wingenfeld[1]

1. Fraunhofer Institute for Material Flow and Logistics
Joseph-von Fraunhofer Str. 2-4
44227 Dortmund

Corresponding author: Lucas Schreiber
lucas.schreiber@iml.fraunhofer.de

Keywords: *Digital Supply Chain Twin, Physical Internet, Data Space, Plant Construction Logistics, AI-Driven Risk Management, Supply Chain Control Tower*

Abstract: *Plant construction relies on temporary, globally dispersed supply chains that face high disruption risk and are characterized by a low level of digitalization. This paper presents an integrated logistics assistance system that combines fuzzy multi-criteria supplier evaluation, a discrete-event simulation-based digital supply chain twin, and AI-driven risk management within a data space-ready architecture. Near-real-time edge data continuously updates the digital twin, enabling scenario-based impact analysis of delays and disruptions. Large language models detect external risks, which are quantified via simulation to generate and rank mitigation options. The system operationalizes core Physical Internet principles for project-based, cross-organizational logistics networks.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan: ☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories), ☐ Transportation Equipment, ☐ PI Networks, ☒ System of Logistics Networks, ☒ Vertical Supply Consolidation, ☒ Horizontal Supply Chain Alignment, ☒ Logistics/Commercial Data Platform, ☒ Access and Adoption, ☒ Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

1 Introduction

Supply chains in the plant construction and engineering, procurement, and construction (EPC) sector are fundamentally distinct from those in serial manufacturing. They are inherently temporary, project-specific, and globally dispersed, and therefore require heterogeneous actors – general contractors, specialized suppliers, logistics service providers, and packaging companies – to collaborate in networks that are dissolved upon project completion. These networks must handle oversized and heavy-lift cargo, synchronize multimodal delivery sequences to capacity-constrained construction sites, and absorb disruptions caused by geopolitical instability, weather events, and cascading delays across interdependent transport legs. Despite the critical economic significance of this sector, it still exhibits low levels of digitalization with respect to integrated supply chain management. Existing approaches are limited to isolated aspects such as transport management or procurement. However, no solution currently available provides an end-to-end digital supply chain twin (DSCT) with simulation capabilities, near-real-time status integration, and sovereign cross-organizational data exchange

tailored to the specific requirements of temporary construction supply networks. This paper addresses the identified research gap by presenting an integrated logistics assistance system. For the first time, a discrete-event simulation-based DSCT is combined with multi-criteria supplier evaluation and AI-driven risk management within a federated, data space-ready architecture. Furthermore, the implications of the proposed approach are examined with respect to the implementation of Physical Internet principles in project-based industries.

2 Towards a Digital Supply Chain Twin for Plant Construction

The proposed system architecture is structured as three interdependent functional modules that collectively enable holistic, end-to-end supply chain orchestration (see Figure 1).

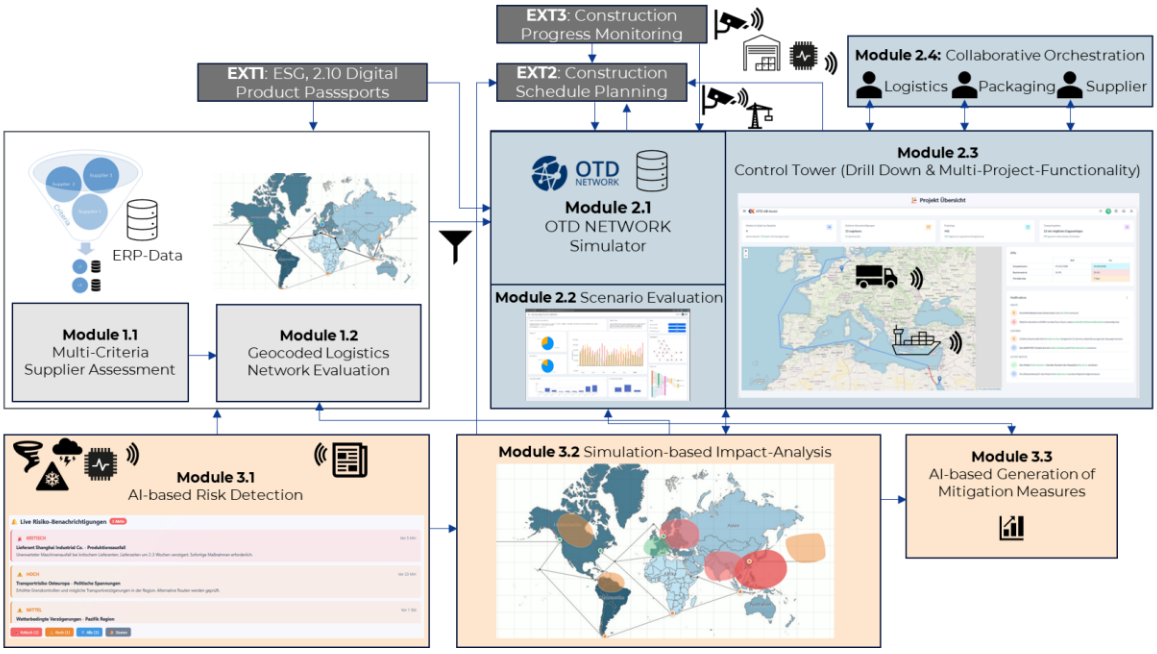


Figure 1: Big Picture of the Digital Supply Chain Twin for Plant Construction

The simulation-capable digital supply chain twin (Module 2) represents the central contribution, while supplier evaluation and AI-driven risk management provide essential input layers that continuously parameterize and enrich the digital twin. Built upon a discrete-event simulation engine (OTD NETWORK), the digital twin provides a geocoded, model-based representation of the entire logistics network - from suppliers through multimodal transport chains integrating heavy haulage, sea freight, air freight, and inland distribution to module yards or the construction site itself. Construction site entities are modeled explicitly to derive material requirements dynamically. The simulation environment enables scenario-based evaluation of cascading effects across the supply chain: the impact of delays on individual transport lanes, load carrier shortages, port diversions, and replenishment can be assessed quantitatively through what-if analyses. A scenario management layer allows the systematic comparison of alternative network configurations and the derivation of individual KPIs for multi-project portfolios. Near-real-time status data is processed within a cloud-edge based computing architecture. Edge devices, such as smart camera systems that assess unloading bay availability, and IoT sensors that detect damage risks during transport, perform decentralized data preprocessing and feed relevant supply chain events into the digital twin. The integration of edge-derived operational data and discrete-event simulation constitutes a novel contribution, enabling the transition from a purely planning-based approach to a continuously updated, state-aware representation of the supply chain. The digital twin is complemented by a supplier

evaluation module that employs multi-criteria decision-making methods to support the initial configuration of the logistics network (Module 1). Beyond conventional supplier assessment criteria, this module evaluates the complexity of the planned logistics network by combining geocoded supplier data with transport optimization services early in the procurement phase, thereby providing a network-aware foundation for the subsequent instantiation of the digital twin.

A risk management module (Module 3) leverages large language models to continuously monitor and interpret heterogeneous public data sources like newsfeeds, geopolitical event databases, meteorological services, and risk platforms for the automated detection, classification, and contextualization of supply chain risks. After identified risks have been linked to the corresponding entities in the supply chain model, their impacts are quantified through simulation-based analysis to assess disruption propagation across the network. Subsequently, the module generates prioritized, actionable recommendations to mitigate disruptions, including alternative routing, supplier substitution, and buffer capacity adjustments. These recommendations are proactively rolled out to downstream supply chain entities that are expected to be affected by the disruption.

A defining architectural characteristic of the proposed system is its integration with a federated data ecosystem based on data space trust frameworks, and policy-based access control. Achieving data space readiness addresses a fundamental challenge in plant construction logistics: the need for sovereign, temporary instantiations of cross-organizational data-sharing relationships that mirror the project-based lifecycle of the underlying supply chains. Each participating organization retains data sovereignty while enabling coordinated decision-making through standardized connectors that ensure interoperability with adjacent data ecosystems. This architecture represents a departure from centralized platform approaches and aligns with the decentralized, open, and interoperable network paradigm that underpins the Physical Internet vision.

3 Related Work

The end-to-end approach for supply chain management presented in this work combines several use cases in modules that are gaining importance and are being researched across different industries. The focus area of this study is the plant construction sector, which serves as the main application domain for the proposed framework.

3.1.1 Supplier Selection in Plant Construction

Supplier selection is a widely studied problem and is typically formulated as a Multicriteria Decision Making (MCDM) problem. Sahoo et al. [1] provide a comprehensive overview of methods and their application in various sectors. The literature review highlights the need for future work to close the gap between theoretical models and data-driven supplier selection in real-world settings, which is a central objective of our research. The most common approach in literature is the Analytic Hierarchy Process (AHP), which exists in several variants. Fang et al. [2] use AHP to weigh different criteria according to their impact on supply chain resilience and digitalization. Supplier ranking is then performed using the Technique for Order of Preference by Similarity to Ideal Solution (TOPSIS). Both AHP and TOPSIS are integrated into a multi-objective programming model (MOPM) to enable Pareto optimization. Kao et al. [3] use a fuzzy AHP (FAHP) to determine the importance of criteria and sub-criteria based on expert opinion and literature in the garment industry. Subsequently, the Weighted Aggregated Sum Product Assessment (WASPAS) method ranks supplier alternatives according to the FAHP-derived

weights and scores. Ran et al. [4] employ the Analytic Network Process (ANP), which extends AHP by capturing interdependencies both between clusters of factors and within each cluster, to evaluate the interrelated importance of criteria and determine which suppliers satisfy them best. The integrated AHP-grey relational analysis (GRA) is used by Wang et al. [5]. Their results confirm that quality, cost, and flexibility are the most important attributes of resilient suppliers, resulting in the need for suppliers to be flexible and provide high-quality standards. However, their framework is highly data-intensive and data-dependent, and the selection of criteria remains challenging due to its subjectivity.

3.1.2 Simulation-based Digital Supply Chain Twin in Plant Construction

The use of digital twins (DTs) has attracted considerable attention in recent years. Originally developed for the aerospace and healthcare sectors, DTs were later adopted in supply chain management, and more recently, in the construction industry. Jefferson et al. [6] propose a DT for general supply chain operations to visualize the network and its metrics, detect violations using alerts, conduct what-if analyses, and perform optimizations. They suggest extending the control tower with downstream data to enable further optimization. These objectives are closely aligned with our goals. For the construction industry, Lee et al. [7] present a DT for route planning, integrating geographic information systems to improve logistics. Their DT is continuously updated with building information models (BIM) and real-time data from Internet of Things (IoT) sensors. Bakhshi et al. [8] describe multiple applications of DTs to address poor information flow, low traceability, and the absence of real-time information. They emphasize that DTs can also be used to analyze and optimize complex aspects of construction projects, such as design alternatives or construction schedules. For future research, they recommend applying DTs across the entire supply chain rather than focusing on individual stages only. Ivanov et al. [9] highlight risk identification and mitigation as core principle of an effective DT. Their approach begins with a risk analysis that draws on external and internal data to generate multiple disruption scenarios. Simulation-based testing with example disruptions then supports the optimization of the supply chain design and model parameters. Finally, a business intelligence (BI) tool evaluates performance, disruption impacts, and affected KPIs.

3.1.3 AI-based Risk Management in Plant Construction

There exists a wide variety of approaches for the identification of supply chain risks. Most recent studies rely on machine learning, such as in the system proposed by Huang et al. [10]. While only a limited number of approaches employ large language models (LLMs) for risk identification, it is evident that data quality is crucial for obtaining useful results in all cases. Sadeek et al. [11] use news media and Twitter data as input to a Latent Dirichlet Allocation (LDA) model to identify risk topics, then test these using a nonparametric significance test. By adding an ordered probit model, they successfully link risks to sentiment. For future work, the authors suggest embedding a correlation analysis within LDA or employing more advanced models such as Bidirectional Encoder Representations from Transformers (BERT) or Generative Pre-trained Transformer (GPT) models. Wang et al. [12] present an integrated framework that combines natural language processing methods with anomaly detection via graph-based fusion. Graph neural networks model correlations, BERT-based encoders extract semantics from news, and LSTM autoencoders detect temporal anomalies in operational metrics. Together, these components identify disruptions in both external signals and internal operations. Existing research on decision-making in the context of supply chain risks and disruptions largely remains at a conceptual level, focusing on requirements and general needs rather than on concrete, implementable solutions. Baryannis et al. [13] review risk management techniques in supply chains, including multi-criteria decision analysis, mathematical optimization, and AI-based methods such as Petri nets, multi-agent systems, and rule-based

models. Manu et al. [14] and Sanusi et al. [15] do not present a fully autonomous mitigation system but rather use AI as a decision-support tool. Sanusi et al. use AI to propose and simulate mitigation options based on risk scores combining probability, impact, and criticality. Manu et al. use AI more operationally to generate, update, and compare alternative management and risk-response options over the project lifecycle. In the end for both approaches, experts choose from the different options.

4 Implementation Concept for the Digital Supply Chain Twin in Plant Construction

4.1 Supplier Selection

The supplier evaluation module implements a two-stage process that combines multi-criteria supplier assessment with a geocoded logistics network evaluation.

4.1.1 Module 1.1: Multi-Criteria Supplier Assessment

The first stage provides a parameterizable system for the definition, weighting, and scoring of supplier evaluation criteria. The module supports an open criteria catalog that accommodates both quantitative criteria (e.g., material costs, delivery lead times, production capacities, volumes, and weight specifications of products) and qualitative criteria (e.g., sustainability certifications according to ISO 20400:2017, quality management standards per DIN EN ISO 9001:2015, reliability assessments, and regulatory compliance including the European and German Supply Chain Due Diligence Act). The criteria catalog is fully configurable on a per-project basis, enabling users to define custom evaluation dimensions and adjust weighting schemes according to project-specific requirements.

The algorithmic core of the evaluation employs a fuzzy logic layer for the quantification of natural-language assessment scales. Qualitative criteria are captured through linguistic variables (e.g., "very good", "satisfactory", "poor") and mapped to fuzzy membership functions, enabling the systematic treatment of inherently imprecise expert judgments. Criteria weights are determined through pairwise comparison matrices following a fuzzy AHP approach. The final supplier ranking is computed using a TOPSIS algorithm that determines each supplier's relative closeness to the ideal solution based on the weighted, normalized decision matrix. The module outputs a ranked list of suppliers with detailed scoring breakdowns, visualized through configurable dashboards that support drill-down into individual criteria and their contribution to the overall score. The entire evaluation logic is exposed through a parameterizable REST API, enabling both interactive use via the web-based frontend and programmatic integration with external procurement systems.

4.1.2 Module 1.2: Geocoded Logistics Network Evaluation

For the subset of suppliers that pass the initial multi-criteria assessment, a second evaluation stage assesses the logistical complexity arising from each supplier's geographic location within the context of the planned project network. The module constructs a geocoded network representation, where supplier locations, intermediate logistics nodes (ports, transshipment hubs, module yards), and the construction site are placed on a geographic model using ISO 19115:2014-compliant metadata for geographic information.

Geocoded supplier data is transmitted through a dedicated API to a transport optimization service that computes multimodal routing options, accounting for the specific constraints of plant construction logistics: heavy-lift and oversized cargo requirements, modal restrictions (sea and air freight, rail, road and inland waterway), customs and regulatory barriers, as well as the availability of specialized transport equipment. The service returns a set of transport optimization scenarios with trade-off solutions encompassing logistical complexity indices,

estimated transport costs, sustainability metrics (i.e. CO_2 emissions per route), service-level implications, and alternative routing options including consolidation strategies.
These logistics-derived metrics are fed back into the multi-criteria framework as additional quantitative criteria, enriching the initial supplier ranking with network-aware information. The output of this stage is an initial supply chain network configuration: A geocoded project network with selected suppliers, logistics service providers, and defined distribution channels, that serves as the foundation for the subsequent instantiation of the digital twin. Figure 2 shows the conceptual overview of the two stages.

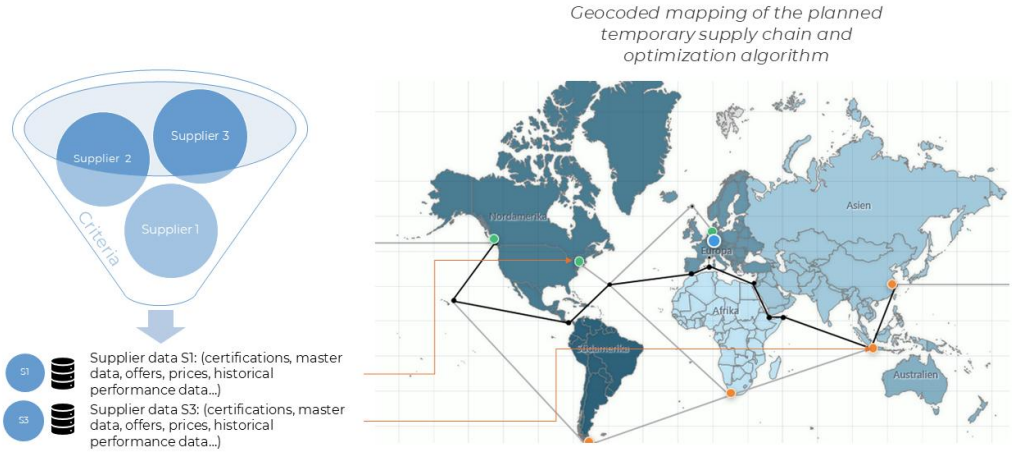


Figure 2: Module 1 - Supplier Selection Concept

4.2 Supply Chain and Logistics Management

The supply chain and logistics management component implements the simulation-capable digital supply chain twin at the core of the assistance system. Built upon the OTD NETWORK discrete-event simulation platform, extended with domain-specific building blocks for plant construction supply chains, this component transforms the geocoded network configuration from Module 1 into a state-aware decision-support environment. Its architecture consists of four modules.

4.2.1 Module 2.1: Supply Chain Modeling and Real-Time State Integration

The first module generates a simulation-capable digital representation of the project-specific supply chain. Its object-orientated model and discrete-event simulation logic cover supplier nodes with production capacities and schedules; transport relations with mode-specific parameters for heavy haulage, sea freight, air freight, and inland transport; logistics nodes such as ports, transshipment hubs, module yards, and staging areas; and the construction site with its work trades and progress states. Material supply requirements are derived dynamically from the construction progress, enabling a pull-based replenishment logic. Entities are interconnected through configurable control logics that trigger flows based on temporal conditions, utilization thresholds, or events.
The model is instantiated from the geocoded network generated in Module 1.2 and enriched with operational data from participating organizations via Eclipse Dataspace Component-based connectors, with each organization retaining full sovereignty over access policies. Near-real-time state data is ingested via a cloud-edge based architecture: edge devices, such as smart camera systems at unloading bays, IoT sensors on containers, and weather stations, perform decentralized preprocessing and only transmit supply-chain-relevant events.

4.2.2 Module 2.2: Supply Chain Scenario Evaluation

This module provides the structured definition, versioning, and simulation-based evaluation of supply chain scenarios. A scenario management layer ensures reproducibility of planning assumptions across the project lifecycle. The evaluation core performs discrete-event

simulation runs based on the model from Module 2.1 to quantify impacts on logistics costs, delivery dates, supply security, CO_2 emissions, and further configurable KPIs. A scenario wizard supports the automated, sequential execution of multiple predefined runs, addressing questions such as the effect of transport delays, load carrier shortages, port closures, hub diversions, staging capacity constraints, and replenishment optimization. Results are aggregated, persisted, and prepared as comparative KPI tables and geographic overlays. This module also serves as the computational backbone for Module 3.2, where risk-derived disruption scenarios are evaluated by the simulation engine.

4.2.3 Module 2.3: Supply Chain Control Tower

The Control Tower provides interactive, role-based dashboards with integrated geographic visualizations. Central KPIs, like throughput, costs, inventory levels, CO_2 emissions, and on-time delivery performance, are displayed through charts, tables, and a geocoded network map, showing real-time transport movements and node states.

Drill-down functions allow navigation from aggregated multi-project portfolio views through project-level overviews to individual transport segments or shipments. The module integrates event feeds from all other components: risk alerts from Module 3.1 appear as severity-classified notifications, scenario results from Module 2.2 are accessible on demand, and edge-derived state changes are displayed alongside simulation results. Supplementary materials such as photographic documentation or contractual references can be linked for context.

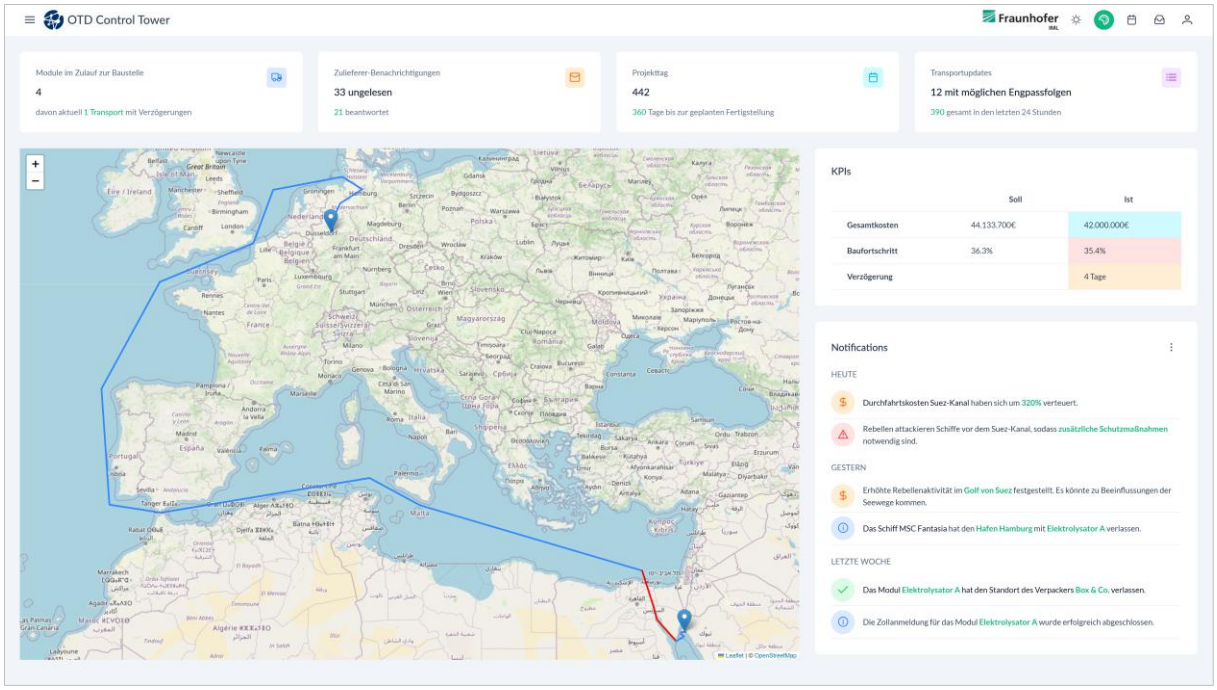


Figure 3: Mock-Up of the control tower with integrated project-specific real-time information

4.2.4 Module 2.4: Collaborative Supply Chain Orchestration

This module enables multi-stakeholder use of the digital twin through role-based access management. Each stakeholder category, like general contractors, logistics providers, packaging companies, and suppliers, receives differentiated permissions within the data space framework. A logistics provider can, for instance, update the transport status and participate in scenario evaluations for assigned transport legs without accessing supplier cost structures or competing offers.

The module implements coordinated disruption response workflows: when Module 3.3 generates mitigation recommendations, affected stakeholders jointly review and approve

countermeasures in a shared decision environment. Cascading notifications proactively alert downstream partners who are not yet physically affected but will be within a foreseeable timeframe, enabling anticipatory planning adjustments and thereby operationalizing the Physical Internet principle of decentralized actors collaborating through standardized interfaces to achieve system-level resilience.

4.3 AI-driven Risk Management

The risk management component implements a three-stage pipeline for continuous identification, quantification, and mitigation of supply chain risks. It tightly integrates the digital twin for geocoded contextualization with the discrete-event simulation engine for impact quantification and AI training for mitigation measure generation, presents results in the Control Tower, and triggers proactive, cascading notifications via the collaboration layer.

4.3.1 Module 3.1: AI-based Risk Detection

This module employs large language models (LLMs) for the continuous monitoring of heterogeneous public data streams: general and industry-specific newsfeeds, geopolitical event databases, meteorological services, risk platforms, social media, and data generated by edge devices such as on-site weather stations. The LLM pipeline extracts risk-relevant information from unstructured text, classifies risks by type (geopolitical, meteorological, logistic, economic), and assesses probability and severity.

Identified risks are contextualized with respect to the current supply chain model: through spatial intersection with the geocoded network representation, each risk is mapped to potentially affected supply chain entities, such as suppliers, transport corridors, ports, or construction sites. The output is a prioritized, source-attributed list of risks linked to specific supply chain entities.

4.3.2 Module 3.2: Simulation-Based Impact Analysis

Identified risks are modeled as disruption scenarios for simulation: port closures are represented as capacity reductions on affected routes; geopolitical events result in route unavailability or increased processing times; and weather events impose temporal restrictions on affected transport legs. The parameterization is based on historical disruption data combined with severity estimates from Module 3.1.

Once defined, the scenarios are applied to the current supply chain model using the discrete-event simulation engine from Module 2.1, thereby propagating cascading effects of disruptions across the network. Specifically, horizontal cascading effects, such as disruptions to one transport leg that lead to delays or capacity conflicts in downstream segments, are effectively captured by the simulation. This capability is of particular importance for the sequential, multi-stage supply chains inherent to the plant construction industry. Outputs of the module include quantified impacts on timelines, costs, and construction progress; affected supply chain elements; a geographic risk heatmap; and comparisons across best-case, worst-case, and baseline scenarios. Results are visualized through the Control Tower (Module 2.3).

4.3.3 Module 3.3: AI-based Generation of Mitigation Measures

An AI agent, trained using the simulation-based impact analysis from Module 3.2, generates mitigation options for disruptions prioritised according to their expected KPI improvements: alternative routing across different transport corridors and modes, backup supplier activation (based on Modules 1.1, 1.2), buffer capacity adjustments, schedule modifications, or shipment consolidation to circumvent bottlenecks. Each proposed alternative is validated in simulation relative to the disrupted baseline scenario and ranked accordingly. The final recommendations include the proposed action, expected risk reduction, cost and timeline implications, required resources, and involved stakeholders, and are presented in the Control Tower (Module 2.3) for expert review and approval.

A distinctive feature of this module is the proactive, cascading notification mechanism: downstream entities not yet physically affected by a disruption are alerted in advance through the collaboration layer (Module 2.4), enabling anticipatory planning adjustments – a concrete operationalization of the Physical Internet principle of distributed, anticipatory logistics intelligence. The system also follows a human-in-the-loop paradigm: recommendations are reviewed and approved by experts before execution, reflecting the contractual and legal complexity inherent to plant construction decisions.

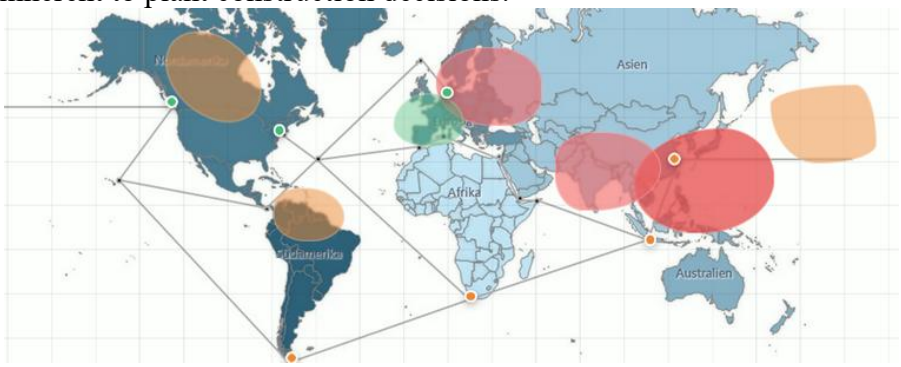


Figure 4: Control Tower visualization of the global risk heatmap including node-level risk assessments

5 Conclusion

This paper has presented an integrated logistics assistance system for end-to-end supply chain management in plant construction. The system combines multi-criteria supplier evaluation with logistics network assessment, a simulation-capable digital supply chain twin with near-real-time state integration, and AI-driven risk management – within a unified, data space-ready architecture. The central contribution is the discrete-event simulation-based digital twin that, through cloud-edge integration, transitions from a static planning instrument to a continuously updated, state-aware decision-support system. The integration of supplier evaluation and risk management creates a closed-loop cycle: supplier evaluation configures the network, the digital twin monitors its operations, and risk management detects disruptions whose impacts are quantified through simulation and fed back as alternative scenarios. From a Physical Internet perspective, the federated data ecosystem architecture operationalizes the PI principle of open, interconnected logistics networks in a domain of temporary, heterogeneous constellations of actors – demonstrating that PI-enabling infrastructure can be realized even in non-persistent network configurations. The collaborative Control Tower implements distributed logistics intelligence across organizational boundaries, while the integration of AI-driven risk detection and mitigation through simulation-based evaluation bridges high-tech capabilities with practical decision support. The modular, standards-based architecture ensures transferability to other project-based industries such as infrastructure construction, shipbuilding, or large-scale event logistics. Several directions for future research emerge. First, comprehensive empirical evaluation of the system's impact on decision quality and logistics costs is required as industrial validation with consortium partners progresses. Second, human-AI interaction dynamics – how experienced professionals adopt, modify, or override generated recommendations – present an important research avenue. Third, the long-term extension toward prescriptive analytics, where pre-approved mitigation measures are executed autonomously within defined policy boundaries, would further advance the PI vision of self-organizing logistics networks.

Acknowledgment. This work is conducted within the Construct-X project, funded by the German Federal Ministry for Economic Affairs and Climate Action (BMWK) under the IPCEI-CIS framework for Cloud-Edge Infrastructure and Services.

References

[1] S. K. Sahoo, S. S. Goswami and R. Halder, "Supplier selection in the age of industry 4.0: a review on MCDM applications and trends," *Decision making advances,* p. 32–47, 2024.

[2] F. Jiaqi, W. Zhou and L. Xiong, "Multi-criteria decision making approach for supplier selection and order allocation in a digital supply chain resilience," *Annals of Operations Research,* pp. 1-37, 2024.

[3] J.-C. Kao, C.-N. Wang, V. T. Nguyen and S. T. Husain, "A Fuzzy MCDM Model of Supplier Selection in Supply Chain Management," *Intelligent Automation & Soft Computing,* 2022.

[4] Y. Ran, R. Yang, N. Sangkhiew and others, "ANALYTIC NETWORK PROCESS FOR GREEN SUPPLIER SELECTION: A CASE OF A CONSTRUCTION PROJECT," Silpakorn University, 2024.

[5] T.-K. Wang, Q. Zhang, H.-Y. Chong and X. Wang, "Integrated supplier selection framework in a resilient construction supply chain: An approach via analytic hierarchy process (AHP) and grey relational analysis (GRA)," *Sustainability,* p. 289, 2017.

[6] S. Jefferson and R. Musa, "Supply Chain Digital Twin Application: Control Tower," in *Optimizing Supply Chains Through Digital 2026*, p. 43–57.

[7] D. Lee and S. Lee, "Digital twin for supply chain coordination in modular construction," *Applied Sciences,* 2021.

[8] S. Bakhshi, A. Ghaffarianhoseini, A. Ghaffarianhoseini, M. Najafi, F. Rahimian, C. Park and D. Lee, "Digital twin applications for overcoming construction supply chain challenges," *Automation in Construction,* p. 105679, 2024.

[9] D. Ivanov and A. Dolgui, "A digital supply chain twin for managing the disruption risks and resilience in the era of Industry 4.0," *Production Planning & Control,* p. 775–788, 2021.

[10] S. Huang and others, "Ai-driven early warning systems for supply chain risk detection: A machine learning approach," *Academic Journal of Computing & Information Science,* p. 92–107, 2025.

[11] S. N. Sadeek and S. Hanaoka, "Assessment of text-generated supply chain risks considering news and social media during disruptive events," *Social network analysis and mining,* p. 96, 2023.

[12] J. Wang, "Multimodal Deep Learning Approach for Early Warning of Supply Chain Disruptions Using NLP and Anomaly Detection," *Artificial Intelligence and Machine Learning Review,* p. 98–110, 2024.

[13] G. Baryannis, S. Dani, S. Validi and G. Antoniou, "Decision Support Systems and Artificial Intelligence in Supply Chain Risk Management," in *Zsidisin, Henke (Hg.) 2019 – Revisiting Supply Chain Risk*, 2019, p. 53–71.

[14] B. A. Manu, "Leveraging Artificial Intelligence for optimized project management and risk mitigation in construction industry," *World Journal of Advanced Research and Reviews,* p. 2924–2940, 2024.

[15] A. N. Sanusi, O. F. Bayeroju and Z. Q. S. Nwokediegwu, "Framework for applying artificial intelligence to construction cost prediction and risk mitigation," *Journal of Frontiers in Multidisciplinary Research,* p. 93–101, 2020.

Optimizing Physical Internet Container Circulation under Demand and Return Uncertainty: A Case Study of Global Rubber Shipment

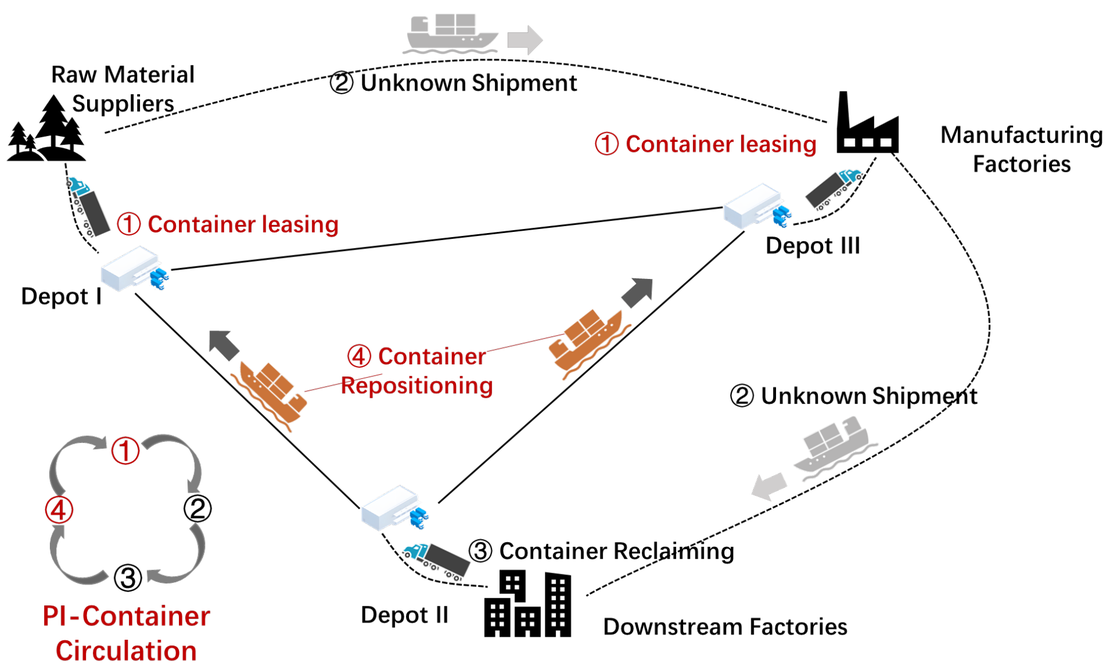
Yunyu Zhang, Song Huang, Fangchun Peng, Dong Yang, Saijun Shao, Ming Li, Zhiyuan Ouyang *

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026
Bordeaux, FRANCE

The Hong Kong Polytechnic University 香港理工大學

INTRODUCTION

PI container circulation includes the whole process from issuing, shipment, reclaiming to repositioning. In this process, issuing and repositioning are governed by a PI container leasing company while shipment and reclaiming are controlled by the leasing customers. The limited information sharing between the company and customers makes both leasing and reclaiming requests uncertain.



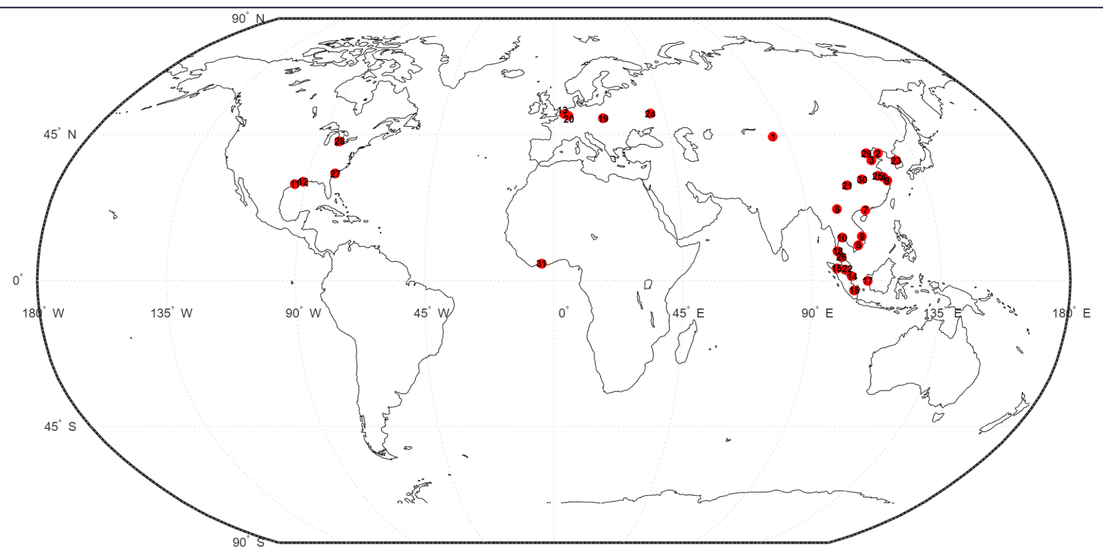
PI-container leasing company: ①④

Customers: ②③

Unknown leasing demand & container returns

- When should PI-container be repositioned, in what quantities, and between which depots?
- Should received leasing demands be fulfilled immediately, or processed at a later time?

METHODOLOGY



PI-Container Circulation Problem (PI-CCP)

Model Input: Logistics network N (Shipment time & cost)
Received leasing demand I_t
Received reclaiming request $\Delta v_{t,n}^{inv}$

Decision Variables: Order processing decision $z_{t,i}$
Container repositioning decision $x_{t,n,n'}$

Decision Period: $t = [0, \tau, 2\tau, \dots, T]$

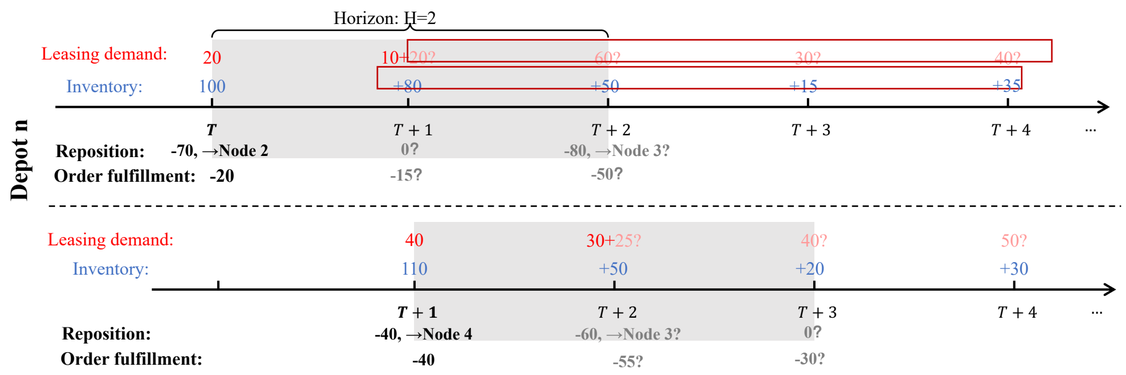
Optimal Objective:
Minimize repositioning cost + order processing delay cost

Modelling: Markov Decision Process

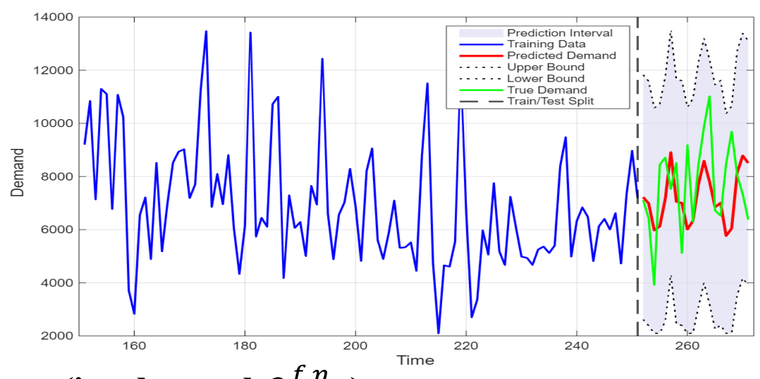
Transition function: $v_{t+\tau,n}^{inv} = v_{t,n}^{inv} - \sum_{n' \in N} x_{t,n,n'} - \sum_{i \in I'_{t,n}} q_i + \sum_{j \in J'_{t+\tau,n}} \hat{q}_j + \Delta v_{t+\tau,n}^{inv}$

METHODOLOGY

Rolling Horizon Optimization



RNN/Croston TSB (Prediction) + Robust Optimization



Uncertainty set (i.e. demand $Q_{t+k\tau}^{f,n}$):
(1)forecasted volume; (2)maximum deviations;

$$[Q_{t+k\tau}^{f,n} - \hat{Q}_{t+k\tau}^{f,n}, Q_{t+k\tau}^{f,n} + \hat{Q}_{t+k\tau}^{f,n}]$$

The scaled deviation: $\vartheta_n = (\tilde{Q}_{t+k\tau}^{n} - Q_{t+k\tau}^{f,n})/\hat{Q}_{t+k\tau}^{f,n}$

Robust budget: $\sum_{n \in N} |\vartheta_n| \leq \kappa_1$;

RESULTS

	PI-CCP	BM	Historical Operation
Total system cost (10^3 \$)	8,014 (-31%)	12,955 (+11%)	11,654
Distance (10^6 km·con)	2,012 (-25 %)	2,525 (+6%)	2,678
Volume (10^3 con)	975 (+56%)	680 (+9%)	624
Average processing cycle (w)	0.96	0.07	2.33
Overdue fulfillment rate	0.87%	0.80%	-
Maximum overdue lateness (w)	3	8	-
Multi-leg reposition rate	12.41%	9.73%	7.69%

- PI-CCP achieves a **24.87%** reduction in repositioning cost with an overdue rate of only 0.87%.
- Customer-held PI-container account for an average of 76.67% (<80%) and a maximum of 85.75% (<90%) of the total fleet size.

CONCLUSION

In this study, we propose a new problem for the entire container leasing, reclaiming and repositioning management, which considers the uncertainties in demand and container reclaiming. Numerical studies and a real-world case study demonstrate:

- The proposed model can achieve a 25% reduction in repositioning cost.
- The empirical consistency between testing instances and the real-world case study validates the 80% (average) and 90% (peak) thresholds as reliable indicators for fleet sizing.
- PI-CCP demonstrates clear advantages in managing long-distance repositioning and long-horizon optimization.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE
KEDGE BUSINESS SCHOOL

Modularization and Containerized Consolidation for Multimodal Three Tier City Logistics

Najib Errami [1,2], S.Ehsan Hashemi-Petroodi [1], Jakob Puchinger [3,2], Walid Klibi [1,4]
1. The Center of Excellence in Supply Chain (CESIT), Kedge Business School, France
2. Laboratoire Génie Industriel, CentraleSupélec Paris, France
3. EM Normandie Business School, Paris, France
4. Physical Internet Center, Supply Chain & Logistics Institute, Atlanta, United States
Corresponding author: Najib.errami@kedgebs.com

Abstract: *Within the European MODALSHIFT project on multimodal transport optimization, we propose a hyperconnected solution for three-tier city logistics under the Physical Internet (PI) framework. The goal is to promote a greener, cost-effective, and collaborative last-mile delivery via diverse modes, including public transport (trams, trains, metro), bikes, electric vans, etc. Structured around a three tier hub system (gateway, local, access) with regional demand clustering, it optimizes PI-container consolidations, enabling multiple logistics service providers to share urban journey portions across these modes and leveraging modularization for efficient detection and handling. Methodologically, this yields an original modeling of integrated networks-of-networks involving multiple actors and transport modes, addressed through advanced optimization of service networks and vehicle routing. Integrating these at the tactical-operational level via a model-based orchestrator supports public-private collaboration in a real-world use case. At the conference, we will present numerical results comparing it to classic networks, demonstrating expected ecological (emission reductions), economic (cost savings), and societal (convenience, collaboration) gains.*

Keywords: *Physical Internet, Collaboration, Modularity, City Logistics, Service network design, Vehicle routing problems, Containerized Consolidation.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☒ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☒ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☐ *In-Person presentation*

1 Introduction

Since 2020, last-mile deliveries have been significantly impacted by the COVID-19 pandemic, leading to a surge in e-commerce. The World Economic Forum (WEF) 2024 warns that this increase could soon overwhelm urban networks without urgent intervention, further exacerbating negative externalities. As the EV transformation will take time, the WEF suggests strategic, tactical, and operational decisions for public and private actors to mitigate the negative impacts of last-mile deliveries. Among these decisions, the WEF emphasizes collaboration between public and private actors, but also among private actors through resource sharing, showing that collaboration can significantly reduce the overall costs and the negative externalities of last-mile deliveries. According to WEF. (2024) the collaboration between

private actors contributes to a more resilient supply chain, improves goods consolidation, and as decision-makers, we should mathematically demonstrate the interest and benefits of this collaboration. In game theory, it comes from the coalition game. The first step is to decide how to form the coalitions, and the second step is to decide how to share the benefits of the coalition. To reduce congestion and improve the real time traffic, numerous projects have been launched worldwide, optimizing delivery times (such as off-hour deliveries in New York City) and spatially exploiting alternative modes: pedestrian deliveries in London by Amazon, river deliveries along the Seine by IKEA, cargo bikes by La Poste, and many others. Moreover, the physical internet naturally proposes some solutions to help the decision-makers to handle the complexity of the network deliveries. Kim et al., (2021) proposed a definition of the region modularization consisting of dividing the region into smaller and synchronized regions. Each truck can be assigned to a smaller route, and hierarchical modularization, which consists of dividing the delivery operations into multiple tiers. These concepts are implicitly used to design some mathematical model inside the city logistics problems. Crainic et al., (2023) exploit the physical concepts of the Internet within the context of city logistics and provide a comprehensive description of the notion of hyperconnected city logistics and the links to existing multimodal city logistics problems at two levels. Our study is motivated from a European project MODALSHIFT (Multimodal Optimisation leveraging Data Acquisition from Local Stakeholders towards a Holistic Improvement of Freight and people Transport). The project aims to develop a tool integrating several techniques (e.g., digital twins, simulation, learning and optimization) to improving the freight and/or people transportation through three main industrial use cases existing in the project. The focus of the project is on the improvement of the real-time traffic, congestion, and operations inter-modal stations within the transportation systems in three use cases Varna, Trieste and Madrid. This study specifically deals with the latter (Figure 1).



Figure 1: Showing the intermodal station Estación sur in Madrid

2 Literature review

This section reviews the literature around the 2TM-CL(Two tier Multimodal City Logistics) problem, the collaboration between private actors and the collaboration between public and private actors. We will show how the 2TM-CL problem is modeled in the literature, and how the collaboration is integrated into such models. The section highlights our main contributions on the collaboration between public and private actors in city logistics. Two-echelon and multi-tier city logistics problems aim to coordinate freight transportation between large distribution

centers and urban customers through intermediate facilities such as satellites or urban hubs. These problems typically involve strategic and operational decisions, including service network design, freight consolidation, and last-mile vehicle routing. In terms of modeling, Crainic et al., (2009) were among the first to introduce the concepts underlying the two-echelon vehicle routing problem and the two-tier city logistics problem, highlighting the differences between traditional single-tier distribution systems and two-tier urban logistics structures.

Because this problem is highly complex and difficult to solve, Fontaine et al., (2021) proposed an approximate modeling approach in which the first echelon is formulated as a service network design problem and the second echelon is simplified as an assignment problem. Since the second tier was not explicitly modeled through an arc-based formulation in this work, Gückel et al., (2026) later proposed an explicit representation of the two-tier city logistics problem by considering only external-zone demand to customers and formulating the second tier as a vehicle routing problem with release and due date (VRPRD). Moreover, they integrated collaboration between different logistics service providers within a centralized optimization framework.

Ghilas et al., (2016a) presented a model to optimize the delivery of goods using public transport, based on pickup and delivery operations along scheduled lines, and solved it using metaheuristic methods. This study was later extended by Ghilas et al., (2016b), where the demand is considered stochastic and the problem is solved using a sample average approximation approach. Recently, Mandal and Archetti, (2025) extended the two-echelon vehicle routing with time windows (2E-VRPTW) by adding an intermediate tier corresponding to pickup and delivery decisions within the public transportation network. They formulate the problem as the Three Tier Delivery Problem on Public Transportation (3T-DPPT). To solve the problem, they propose a matheuristic that solves the tiers sequentially. Because of the difficulty of the problem, Delle Donne et al., (2025) later proposed a decomposition approach based on column generation to solve the 3T-DPPT. Their results outperform the matheuristic proposed by Mandal and Archetti, (2025).

Collaboration in logistics has gained increasing attention to improve efficiency, within centralized collaborative models, several studies have introduced additional constraints to ensure fair collaboration between actors. For instance, Mancini et al., (2021) incorporated workload-balance constraints to guarantee that each carrier serves at least a minimum number of customers, as well as an individual-rationality constraint ensuring that each player obtains a better profit within the coalition than when operating independently. This requires solving each participant's standalone problem before evaluating the collaborative solution.

Despite the recent advances in the literature, several aspects of three-tier containerized multimodal city logistics remain insufficiently studied. In this work, we propose a centralized model that considers collaboration between a logistics service provider and public authorities to use public transport for freight movements. We build a mathematical model that aims to optimize the design and coordination of freight flows from the regional level to last-mile urban delivery. Overall, the proposed framework integrates hierarchical and regional modularization principles by structuring freight flows across coordinated logistics layers and fixed urban hubs, while also enabling collaboration between public and private actors using shared π-containers and a public orchestrator coordinating consolidation and distribution decisions. This integrated approach contributes to improving coordination, consolidation efficiency, and sustainability in urban freight systems. The main contributions of this study are as follows:

- We propose a new model for the three-tier containerized multimodal city logistics (3TCM-CL) problem that integrates Physical Internet modularization concepts and compare the results when we have open routes in the last tier (OMDVRPTW) and when the city freighters should return at departure depot (MDVRPTW). This approach reduces the complexity of first tier while introducing packing decisions for π-containers in the second tier.
- We introduce new fairness constraints to model collaboration in city logistics through the sharing and consolidation of π-containers among actors.

3 Problem statement

We study a three-tier containerized multi-modal city logistics problem 3TCM-CLP. The current model presented concerns only one Logistic Service Provider (LSP), and we will show in section 5 how to extend this type of the model for a collaboration between private actors.

As shown in the Figure 2 the network is divided into a set Z of tiers (three tiers in our case), and each tier is defined by the set of S_z hubs. The demand points are located into the last tier. As we get closer to the city center, the hubs become smaller and more numerous. In the first tier, the hubs can be represented like a city distribution center or fulfillment center, in the second tier, an intermodal station and in the last tier a rendez-vous point for example.
In the physical internet, these hubs are called gateways, local, and access. In the first tier, we determine the service network and routing schedules assuming fixed local hubs, which provides arrival times for long-distance freight services. In the second tier, we optimize the consolidation and packing of shipments into modular π-containers, managed by a public orchestrator who decides how to group freight and dispatch it to rendezvous points or access hubs. In the third tier, we optimize the routes and schedules of city freighters through a multi-depot vehicle routing problem with time windows and synchronization constraints with upstream tiers.

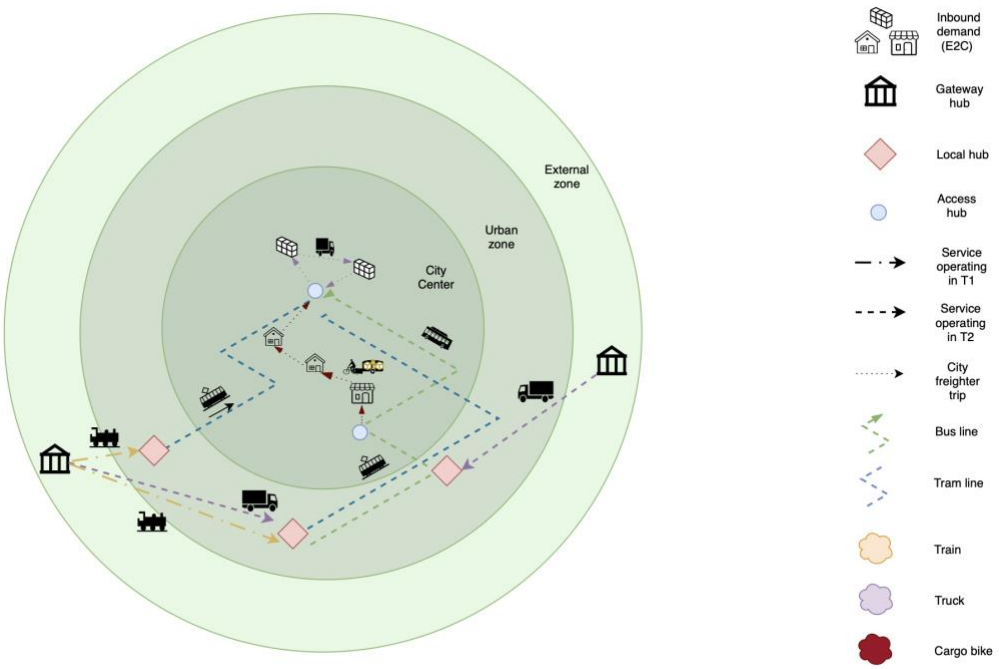


Figure 2: 3TCM-CLP

We consider that the set of origin-destination demands D can come from the firms, organizations, private actors, parcel lockers or shops, and we treat small packages that can be consolidated in CDC, re-consolidated, and containerized in the local hub by assumption.
We assume only the inbound demands located in the city center. Each demand $d \in D$ is addressed by a volume v_d and a time window $a_d\, b_d$ corresponding to the time at which the parcels can be delivered. We assume that all demands can be consolidated while the trip respects the time windows of demands.

In the first tier, the services are defined by a vehicle type $t_r \in T$ and a transportation mode $m_r \in M$, with $r \in R_1$ the services of first tier. For example, the transportation mode can be a truck, and the vehicle type can be a small or a large truck. As we assume that the demand is unsplittable, we treat only the assignment of demands into the services and hubs for the first tier. These services start from the first-tier hubs S_1 and go to the second-tier hubs. The service cost c_r, a service time w_r and the arrival time at each delivery point as τ_{sr} are associated to the service. The cost includes both the monetary and environmental components of using the service, while the service time represents the estimated loading and unloading duration at each delivery point. The cost of delivering demand d using the service r between the first and second tier is represented by f_{dsr}. These costs can be estimated by a distance-based cost, or a time-based cost, or a more complex cost function including the negative externalities of the delivery.

Although first-tier hubs have large capacity, services can be limited by vehicle capacity u_r, fleet size by vehicle type $a(s,t)$ at each hub, the maximum number of vehicles that we can transfer to each hub $a(p,s)$,the maximum number of vehicles that we can transfer to each hub for each mode $\bar{a}(p,m,s)$, and the maximum cargo volume that can be handled to each local hub $u_s(p)$.To centralize service information, we define subsets of services depending on node, time period, mode, type, and compatibility with demand. For example, $R_1(p,s)$ is the set of first-tier services operating at local hub s in period p, and $R_1(d,s)$ is the services compatible with demand d and visiting local hub s. A demand is compatible with a service if the release date of the demand is less than the departure time of the service, and in the optimized model the arrival time of the last-mile service should be within customer time windows. In the same logic, we also define the subsets $R_1(p,s,t)$ and $R_1(p,s,m)$ for the services of type t and mode m operating at local hub in period p.

In the second tier, after all demands reach the local hub, we decide how to containerize parcels and dispatch the pi-containers to access hubs by using public transport R_2. They will transport the π-containers and dropped off them into some potential access hub or rendez-vous points $G_{r_2} \in G$ the set of public transport stops. The public transport can be buses, trams, metro, or even river transport. As the service in the first tier, the public transport is defined by a capacity Q_r and a schedule defined by the arrival times at each stop τ_{rs}, with $r \in R_2$ and $s \in S_2$. For the public transport, we also define a binary parameter β_{ruv} equal to 1 if public service r goes from stop u to v, and 0 otherwise. If the public transport is used, we should also consider the service time at each node given by w_r, which is the estimated time needed to pick up or drop off a container at each stop. In our model, no cost is associated to the second tier while the time-windows in the last mile delivery are respected.

In the last tier, city freighters K deliver parcels from access hubs S_3 to customers D. Many options are possible for the containerization, to decrease the service time, we assume that if two demands are consolidated in the same pi-container, the city freighters should deliver the two

demands in the same trip from the access hub. It minimizes the number of exchanges in access hubs but is less flexible. The city freighters are defined by a capacity Q_k that is expressed in number of smallest pi-containers, and a service time w_k that is the estimated time needed to unload the freight. The cost in the last tier represented by c_{ijk} is given by the time spent by all city freighters to deliver parcels. We minimize the total cost of city-freighter trips depending on containerization decisions in the second tier and the travel times between demands given by t_{ij}. The model is defined through three groups of decision variables, each group corresponding to one tier of the network. In the first tier, the binary variable ξ_r indicates whether service $r \in R_1$ is selected, while x_{dsr} assigns demand d to local hub s through service r.

In the second tier, the binary variable $\gamma_{d\pi}$ indicates whether demand d is packed into container π, and $y_{\pi sr}$ assigns container π to an access hub s through public service $r \in R_2$. These variables capture the consolidation and transfer decisions inside the public transportation network, where freight is grouped into modular units before being dispatched to the last-mile.

In the third tier, z_{ijk} indicates whether city freighter $k \in K$ uses arc (i, j), $\lambda_{\pi k}$ assigns container π to city freighter k, and ω_{dk} links demand d to the freighter that serves it. The variable ϕ_{ds} identifies the access hub associated with demand d, and T_{ik} represents the arrival time of freighter k at node i. Together, these variables define the operational routing decisions of the last mile, while ensuring consistency with the containerization and synchronization decisions made in the previous tiers.

3.1.1 Approximation by regional modularization constraints

We introduce an approximated version of the 3TCM-CL (see Figure 3) considering the modularization constraints, so called, that should be used to cut the search space. The difference is the by adding these constraints, two demands from different zone Z cannot be containerized together.

$$\gamma_{d\pi} + \gamma_{d'\pi} \leq 1, \forall (d, d') : d \in Z, d' \notin Z. \quad (1)$$

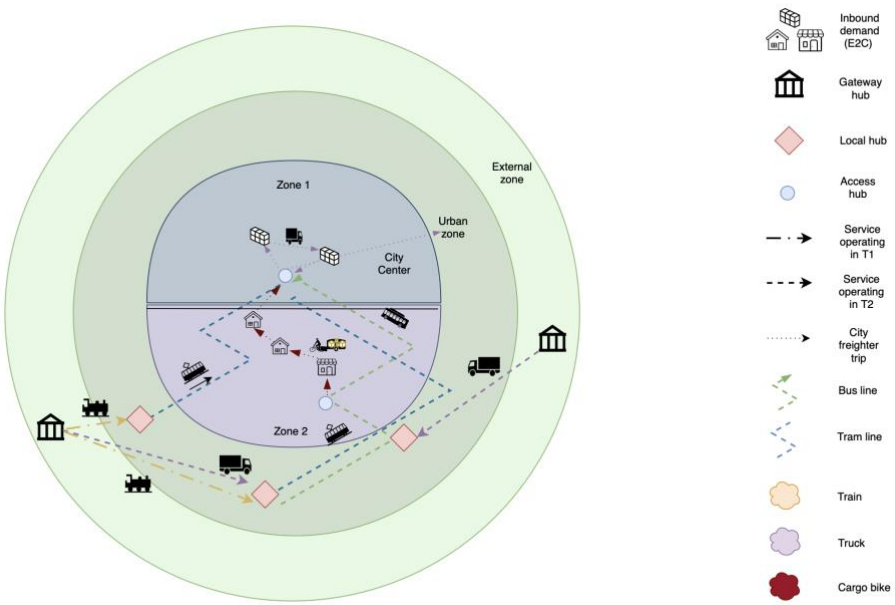


Figure 3: Illustrating hierarchical modularity inside the 3TCM-CL

3.1.2 Extend the collaboration between private actors

We note $\boldsymbol{N}$ the set of LSPs and $\boldsymbol{S(n)}$ the sub coalition of n LSPs that collaborate. To ensure group rationality in the collaboration, we extend constraints introduced by Mancini et al., (2021) by considering all sub coalitions instead of each carrier (constraint 2), it ensures stability concept of coalition game theory. In another words, it ensures that the solution provided by the grand coalition has a lower cost than any sub coalition. We also introduce a new fairness constraint (constraint 3) that ensures each logistics service provider (LSP) meets a minimum workload threshold $\lceil \boldsymbol{B_n} * \boldsymbol{\alpha} \rceil$. $\boldsymbol{B_n}$ is an input of the centralized model. It is the number of sub-trips or containers used by the LSP to deliver all demands. $\boldsymbol{\alpha} \in [\mathbf{0}, \mathbf{1}]$ is set by the user, indicating the flexibility of the new solution to be different than the individual solution, more $\boldsymbol{\alpha}$ is close to one, more the solution resembles to the solution without collaboration. This requires solving each subcoalition's $\boldsymbol{S(n)}$ standalone problem first to establish its baseline. This difficulty illustrates the complexity of solving collaborative problems when there are many actors.

$$\sum_{r\in\mathcal{R}(n)} c_r * y_r \sum_{d\in\mathcal{D}(n)} \sum_{s\in\mathcal{S}} \sum_{r\in\mathcal{R}(n)} f_{ds_r} * x_{dsr} + \sum_{k\in\mathcal{K}(n)} \sum_{(i,j)\in\mathcal{A}_3} c_{ij} * z_{ijk} \leq C_{S(n)}, \forall S(n) \in \mathcal{N}(2)$$

$$\sum_{k\in\mathcal{K}(n)} \sum_{\pi\in\Pi} \lambda_{\pi k} \geq \lceil \boldsymbol{B_n} * \boldsymbol{\alpha} \rceil \quad \forall n \in N. (3)$$

4 Numerical results

The data generated come from the work of Delle Donne et al., (2025). We adapt the data by creating first-tier services based on public transport terminal stations. Based on these data, we generate new points randomly with different coordinates and different volumes around the access hub. We try to solve small instances of 14 customers 1 gateway hub, 2 local hub, 2 access hub, 21 public services and 10 first tier services. The first-tier services are based on trucks, and the city freighters have a capacity five times less than a truck. The models are implemented using the Python package gurobpy (Gurobi 13) and run on CPU Intel Xeon Gold 6230 20C @ 2.1GHz with a time limit of 3 hours (10800 seconds). The mathematical models used during these tests will be presented during the conference.

4.1.1 Modeling individual LSP problem

The results for OMDVRPTW inside the last tier are consistently slightly better in terms of cost than those of the closed-route scheme, although proving optimality is more challenging. In 10 instances, 6 are solved compared with 9 out of 10 instances for the closed-route approach. Interestingly, even when the solution gap is larger, the obtained solution is always at least as good as, and often better than, the classical scheme.

Table 1: Comparison of MDVRPTW with OMDVRPTW on small instances for the last tier

	Total cost (obj value)	Optimality Gap(%)	CPU Time(s)

Instance	OMDVRPTW	MDVRPTW	OMDVRPTW	MDVRPTW	OMDVRPTW	MDVRPTW
0	975.81	977.62	3.21	0.00	10800.04	1267.42
1	1000.48	1000.48	0.00	0.00	399.28	10694.72
2	996.80	996.80	1.40	2.25	10800.10	10800.08
3	867.61	867.61	1.03	0.00	10800.13	6887.38
4	1603.77	1605.56	0.00	0.00	66.15	48.18
5	1003.55	1004.71	3.20	0.00	10800.06	817.95
6	935.78	935.78	0.00	0.00	845.50	729.22
7	1048.61	1076.28	0.00	0.00	70.27	581.15
8	932.29	939.53	0.00	0.00	738.13	582.85
9	990.97	1018.62	0.00	0.00	9176.38	276.84
Average	**1035.57**	1042.30	(6/10 optimal) 0.88	(9/10 optimal) **0.22**	5449.60	**3268.58**

4.1.2 Approximation by zone containerization constraints

To test the zone containerization constraints, we have first generated the zones using the machine learning method Support Vector Machine, which is used to classify the data into two classes. Each class is then used as a zone. By adding the constraints, we solve one instance more than the classic model, but we can see that it is not always faster than the first model, and the approximation model is closer to the optimal value. The model has been tested only with 2 zones, and it would be interesting to extend the tests to include different machine learning methods and more than 2 zones.

Table 2: Showing the impact of zone containerization constraints (ZCC) in optimality using OMDVRPTW in the last tier.

	Total cost (obj value)		Optimality Gap (%)		CPU Time(s)	
Instance	Basic	ZCC	Basic	ZCC	Basic	ZCC
0	975.81	977.62	3.21	4.80	10800.04	10800.05
1	1000.48	1000.48	0.00	0.00	399.28	5525.80
2	996.80	996.80	1.40	7.58	10800.10	10800.04
3	867.61	867.61	1.03	0.00	10800.13	843.31
4	1603.77	1603.77	0.00	0.00	66.15	19.45
5	1003.55	1003.55	3.20	12.12	10800.06	10800.06

6	935.78	935.78	0.00	0.00	845.50	985.82
7	1048.61	1071.72	0.00	0.00	70.27	4213.42
8	932.29	932.29	0.00	0.00	738.13	1424.64
9	990.97	1015.54	0.00	0.00	9176.38	377.05
Average	**1035.57**	1040.52	(6/10 optimal) **0.88**	(**7/10** optimal) 2.45	5449.60	**4578.96**

4.1.3 Modeling collaborative problem

For these experiments, we put the α value to 0.3 for the fairness constraints and consider 2 LSP, we assign randomly to each demand a LSP and include the stability constraints. We assume that the LSP are willing to share the containerized trip only. Then, the cost corresponds to the total time of freighters in the last mile. The results show that the LSP has an interest in collaborating in most instances and gains can go up to 24% in some cases.

Table 3: Showing the positive impact of the collaboration into the last mile costs in OMDVRPTW.

	LSP1			LSP2		
Instance	Individually cost	Collaborative cost	Gains (%)	Individually cost	Collaborative cost	Gains (%)
0	239.83	197.596	18	191.22	189.938	1
1	319.80	283.59	11	168.51	128.613	24
2	220.64	203.024	8	220.56	205.507	7
3	231.31	210.345	9	71.00	68.987	3
4	237.42	225.102	5	206.12	202.122	2
5	282.07	266.712	5	162.10	148.567	8
6	186.19	157.041	16	205.37	190.469	7
7	286.69	286.563	0	176.41	173.77	1
8	247.85	247.743	0	108.21	96.276	11
9	228.58	208.538	9	198.218	194.16	2
Average	248.04	**228.63**	8	170.77	**159.84**	7

Conclusion

In this study, we present some results about our new 3TCM-CL problem. We show the difficulty of solving it by using a commercial solver. The next step would be to propose an advanced methodology to solve it, like column generation or the Branch & Cut & Price method. Even though zone constraints can help the user solve more instances, the results depend on the classification method used. With the growing interest of private actors, we demonstrate their willingness to collaborate in a city logistics context. Inside the collaboration, the next step will be to test more instances by varying the alpha parameter in the goal to provide fairness solutions. In the last tier, we show the advantages of solving an open routing problem compared to the classical multi-depot routing problem. The difficulty in implementing the model lies in the synchronization between container packing and operating routes, but our model can be simplified further by considering a service network design in the last tier. If the city freighters used to take the same routes every day, it makes sense to delve deeper into the problem. Because in the context of city logistics, many unforeseen factors can impact real-time traffic, we plan to simplify the model and include some uncertainties in it to get closer to reality.

References

- T.G Crainic, N. Ricciardi, G. Storchi., (2009). Models for Evaluating and Planning City Logistics Systems, *Transportation Science ,* 43(4):432–454
- T.G Crainic, W. Klibi, B. Montreuil, (2023). Hyperconnected city logistics: a conceptual framework. In Handbook on City Logistics and Urban Freight, (pp. 398-421). *Edward Elgar Publishing.*
- D. Donne, S. Alberto, C. Archetti., (2025). Integrating public transport in sustainable last-mile delivery: Column generation approaches, *European Journal of Operational Research* 324(1):75-91
- P. Fontaine, T.G Crainic, O. Jabali, W. Rei., (2021). Scheduled service network design with resource management for two-tier multimodal city logistics *European Journal of Operational Research* 294(2): 558-570
- J. Gückel and T.G Crainic and P. Fontaine., (2026). A two-step large neighborhood search for a collaborative two-tier city logistics system, *European Journal of Operational Research* 331(1):139-155
- V. Ghilas, E. Demir, T.V Woensel., (2016a). An adaptive large neighborhood search heuristic for the Pickup and Delivery Problem with Time Windows and Scheduled Lines *Computers & Operations Research* 72:12-30
- V. Ghilas, E. Demir, T.V Woensel., (2016b). A scenario-based planning for the pickup and delivery problem with time windows, scheduled lines and stochastic demands, *Transportation Research Part B: Methodological* 91:34-51
- N. Kim, B. Montreuil, W. Klibi., (2021). Modularization of Delivery and Transportation, *Proceedings of the International Physical Internet Conference*
- S. Mancini, M. Gansterer , R.F Hartl., (2021). The collaborative consistent vehicle routing problem with workload balance, *European Journal of Operational Research* 293(3): 955-965
- M. Minakshi, C. Archetti., (2025). Decomposition matheuristics for last mile delivery using public transportation systems, *Soft Comput 29*, 1511–1539 (2025)
- WEF (2024)., Transforming Urban Logistics: Sustainable and Efficient Last-Mile Delivery in Cities, *WEF Report*

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE

Constraint-Based Spatial-Temporal Alignment for Traceable Modular Integrated Construction Logistics in a Cyber-Physical Internet Framework

Pengjun Yue[1], Zhao Zhiheng[1, 2] and George Q. Huang[1, 2]
1 The Department of Industrial and Systems Engineering, The Hong Kong Polytechnic University, Hong Kong, China
2 Research Institute for Advance Manufacturing, The Hong Kong Polytechnic University, Hong Kong, China
Corresponding author: zhiheng.zhao@polyu.edu.hk

Abstract: *Modular Integrated Construction (MiC) improves construction productivity, site safety, and space utilization, but its logistics management relies on reliable spatial-temporal traceability of modular units. Raw traceability data collected from sensing and localization systems are often affected by noise, non-line-of-sight interference, and drift, which may lead to semantic inconsistencies with production, storage, transportation, and installation logic. Conventional filtering methods can smooth noisy trajectories, but they cannot explicitly correct observations that violate logistics process rules. To overcome these challenges, this paper proposes a MiC-Cyber-Physical Internet (CPI) spatial-temporal alignment constraint framework. Specifically, this study formulates MiC traceability as a cyber-physical alignment problem, introduces CPI-based logical constraints to normalize raw traceability observations, and validates the framework on a cyber-physical laboratory platform covering factory, prefabrication-yard, and on-site scenarios. Experimental results show that the proposed method improves position accuracy, zone consistency, and post-assembly stability compared with raw observations and conventional filtering. These findings indicate that CPI-based logical alignment can enhance the reliability of traceability data for MiC logistics management.*

Keywords: *Modular Integrated Construction (MiC); Cyber-Physical Internet (CPI); Spatial-temporal traceability; Spatial-temporal alignment.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☐ *PI Networks,* ☐ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☒ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

1 Introduction

Modular Integrated Construction (MiC) improves productivity, quality control, and construction speed, but it also turns construction into a logistics-intensive process (Wang et al., 2026). In dense urban contexts such as Hong Kong and Singapore, limited and costly on-site buffer space often makes direct on-site storage impractical. Prefabrication yards therefore act as intermediate nodes for off-site storage, dispatch coordination, and just-in-time delivery.

Reliable spatial-temporal traceability is essential in MiC logistics, allowing stakeholders to know the current position of each module, its process stage, and whether its movement aligns

with the planned workflow (Li et al., 2017). This requirement can be framed within the Physical Internet (PI) paradigm, which promotes modular, standardized, and interoperable logistics units across factory, yard, and site nodes (Pan et al., 2021). Building on PI, the Cyber-Physical Internet (CPI) integrates real-time sensing, computation, and digital control to enhance traceability reliability and enforce logical constraints on module movement, effectively aligning physical and cyber representations for MiC operations (Wu et al., 2025).

In practice, sensing and localization data are often corrupted by noise, non-line-of-sight interference, signal blockage, and drift, especially in factory–yard–site environments (Nkrow et al., 2024). The resulting trajectories may contain not only numerical errors but also semantically unreasonable states, such as wrong zone jumps or invalid stage transitions.

Conventional filtering methods such as moving average and Kalman filtering can smooth trajectories, but they do not explicitly consider whether the observations are logically consistent with the MiC process (Ranjan et al., 2024).

To address this limitation, this paper proposes a constraint-based spatio-temporal alignment framework for traceable MiC logistics in a CPI framework. The framework links physical modules with digital entities and uses CPI-based logical constraints, including zone rules, stage transitions, temporal continuity, and post-assembly stability, to normalize raw traceability data. Experiments show that the proposed method improves MiC traceability compared with raw observations and conventional filtering.

The main contributions of this paper are as follows:

(1) It formulates MiC traceability as a cyber-physical alignment problem rather than a pure localization correction problem;

(2) It develops a CPI-based logical constraint mechanism for normalizing raw traceability observations from heterogeneous sensing systems; and

(3) It validates the proposed framework in a laboratory platform covering MiC factory, yard, and on-site scenarios, demonstrating its effectiveness in improving traceability consistency.

The remainder of this paper is organized as follows. Section 2 reviews related work. Section 3 presents the system overview. Section 4 introduces the proposed methodology. Section 5 reports the experiment and validation results. Section 6 concludes the paper.

2 Related Work

2.1 MiC logistics and traceability

MiC has been widely recognized as a logistics-intensive construction mode in which modules are manufactured off-site, transferred through intermediate logistics nodes, and assembled on-site. Prior studies show that successful MiC implementation depends heavily on supply chain coordination, planning, and logistics synchronization rather than solely on design or manufacturing efficiency. Reviews by Hussein et al. (2021), Lim et al. (2022), Zaalouk et al. (2023), and Arshad and Zayed (2022) all highlight the importance of logistics and supply chain management in MiC.

However, most existing studies focus on planning and managerial coordination. Comparatively less attention has been paid to continuous module-level tracking across the factory–yard–site chain and to converting raw sensing data into trustworthy spatial-temporal traceability information for logistics decision-making.

2.2 Digital twin, IoT, and CPI-enabled MiC management

To improve information visibility in prefabricated construction, BIM, IoT, and digital twin technologies have increasingly been integrated into logistics and production management. Zhong et al. (2015), Zhou et al. (2021), and Yang et al. (2024) show that these technologies can enhance information exchange, visibility, and traceability in MiC-related supply chains.

At a broader level, digital twins and cyber-physical systems enable real-time synchronization between physical entities and their cyber representations (Tao et al., 2019). In logistics, Wu et al. (2025) formalized the CPI and Yin et al. (2024) developed a CPI-based platform for off-site construction supply chains.

Nevertheless, most CPI and digital twin-related studies emphasize system integration and visibility, rather than how cyber-layer process logic can be used to evaluate and correct noisy module-level trajectories.

2.3 Localization technologies and the limitation of signal-level filtering

UWB and related sensing technologies have been extensively studied for real-time construction tracking because they can provide location data for workers, materials, and equipment in dynamic jobsite environments. Prior work confirms both their applicability and their sensitivity to deployment conditions (Cheng et al., 2011).

Filtering and sensor-fusion approaches can suppress outliers and smooth trajectories. For example, Dong et al. (2023) proposed an improved robust adaptive cubature Kalman filter for UWB localization under NLOS errors.

For MiC logistics, however, trustworthy traceability should mean not only smoother coordinates but also logically consistent states across time, space, and process stages. The current literature therefore reveals a gap between localization correction and logistics-consistent spatial-temporal alignment.

3 Overview of the CATS-MiC System

The proposed constraint-based alignment and traceability system for MiC logistics (CATS-MiC) supports trustworthy spatial-temporal traceability across the factory, prefabrication yard, and construction site. As illustrated in Figure 1, the system links the physical layer, cyber layer, and alignment module in a closed cyber-physical loop.

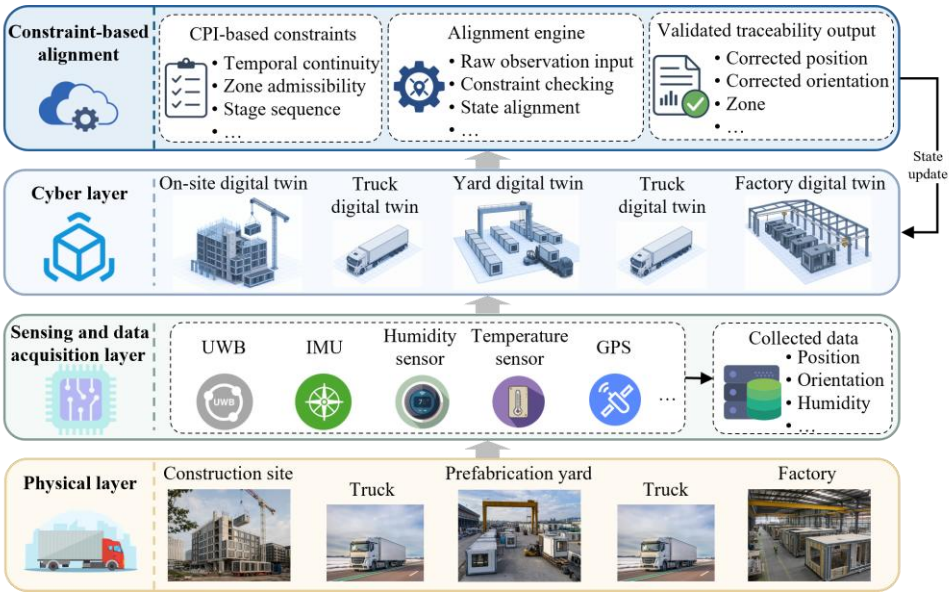


Figure 1 Architecture of CATS-MiC System

3.1 Physical environment

The physical layer represents the logistics flow of MiC modules across the factory, prefabrication yard, and construction site. In the present implementation, embedded sensing modules provide real-time observations of module position and orientation. Through these sensing devices, the physical layer continuously generates observations of each module's position and orientation during movement, transfer, storage, and placement.

However, the collected data are not always accurate. In realistic logistics environments, sensing results may be affected by non-line-of-sight interference, signal blockage, multipath effects, drift accumulation, and temporary instability. As a result, raw observations may deviate from the actual module state, leading not only to numerical errors but also to traceability records that are inconsistent with the real logistics process.

3.2 Cyber environment

The cyber layer is the digital counterpart of the MiC logistics system. Through one-to-one physical-to-cyber mapping, it maintains module identities, digital twin scenes, functional zones, and expected logistics flow rules.

The cyber layer therefore interprets not only where a module is located, but also whether the current state is operationally meaningful. It acts as a constraint source for judging whether sensed observations are credible.

More importantly, the cyber layer encodes the expected logistics flow under the CPI framework. In MiC logistics, module movement is not arbitrary; it follows a reasonable process sequence and spatial transition logic. A module should move through admissible zones and stages in a feasible order, rather than appearing in downstream locations or states without passing through the required upstream process. Therefore, the cyber layer acts not only as a visualization environment, but also as a constraint source for judging whether sensed observations are credible.

3.3 System workflow

The workflow begins with real-time data acquisition from the physical layer. Observations are mapped to corresponding digital entities and interpreted as candidate zone-stage states in the cyber layer.

The system then performs constraint-based alignment by checking zone admissibility, temporal continuity, stage transitions, and post-assembly stability. Invalid observations are adjusted toward the most plausible state under the current cyber-physical context.

Finally, CATS-MiC outputs validated traceability information, including corrected spatio-temporal data, interpreted zone, and process-consistent logistics status.

Finally, the system outputs validated traceability information, including corrected spatio-temporal data, interpreted zone, and process-consistent logistics status. Through this workflow, CATS-MiC produces traceability results that are not only smoother but also logically aligned with the MiC logistics process.

4 Methodology

4.1 Problem definition

In MiC logistics, a modular unit moves across multiple logistics contexts, such as the factory, prefabrication yard, and construction site. During this process, heterogeneous sensing devices

generate real-time observations related to the module's traceability state, including position, motion, orientation, or other state-related information. Let the raw observation of module m at time t be denoted as:

$$o_t^m \tag{1}$$

Over time, these observations form a raw traceability sequence

$$O^m = o_1^m, o_2^m, \dots, o_T^m \tag{2}$$

Due to sensing noise, occlusion, drift, and environmental interference, the raw observations may deviate from the actual module state. More importantly, some observations may be inconsistent with the MiC logistics process, for example by indicating an unreasonable zone jump, an invalid process transition, or an implausible motion state.

Therefore, the objective of this study is to transform raw observations into aligned traceability states that are not only close to the sensed data but also consistent with the CPI logic of MiC logistics. The aligned state of module $m$at time $t$is defined as

$$s_t^m = (p_t^m, z_t^m, q_t^m, a_t^m, c_t^m) \tag{3}$$

where p_t^m denotes the aligned spatial position, z_t^m the interpreted functional zone, q_t^m the logistics stage, $a_t^m$the optional motion or state attribute, and $c_t^m$the confidence score. The task of the proposed method is to estimate the aligned state sequence

$$\hat{S}^m = \hat{s}_1^m, \hat{s}_2^m, \dots, \hat{s}_T^m \tag{4}$$

4.2 CPI-based constraint modeling

The key idea of the proposed method is that module traceability should follow a reasonable cyber-physical flow rather than relying solely on raw sensing observations. To achieve this, the cyber layer provides structured context information, including module identity, zone configuration, process plans, storage and dispatch rules, site demand, and accessibility conditions. Based on this context, a set of CPI-based logical constraints is constructed.

(1) Zone admissibility constraint

A module can only appear in zones that are permissible under its current logistics stage. If z_t^mdenotes the interpreted zone and $Z(q_t^m)$denotes the admissible zone set of stage q_t^m, then the aligned state should satisfy

$$z_t^m \in Z(q_t^m) \tag{5}$$

(2) Stage transition constraint

The logistics stage of a module should follow the predefined process sequence. Let q_t^mdenote the current stage and $\mathcal{T}$denote the set of legal transitions. Then

$$(q_{t-1}^m, q_t^m) \in \mathcal{T} \tag{6}$$

This constraint prevents logically impossible stage jumps.

(3) Temporal consistency constraint

Module states should evolve in a temporally consistent manner. A module should not appear in a downstream stage before the corresponding upstream operation is completed, nor should it violate expected delivery or dispatch timing windows.

(4) Path continuity constraint

The movement of a module should remain spatially continuous. Large discontinuities between consecutive timestamps are treated as unlikely unless supported by the cyber context.

(5) Resource and accessibility constraint

Some module states are only feasible when corresponding resources, destination readiness, or physical accessibility conditions are satisfied. This constraint is used to avoid traceability interpretations that are inconsistent with the actual logistics environment.

(6) Motion plausibility constraint

When motion- or orientation-related observations are available, the state evolution should remain consistent with feasible handling behavior. Abrupt or excessive changes may indicate sensing error or logical inconsistency.

Accordingly, the overall CPI-based constraint set is written as

$$\mathcal{C} = C_{zone}, C_{stage}, C_{time}, C_{path}, C_{resource}, C_{motion} \tag{7}$$

4.3 Constraint-based spatial-temporal alignment

Based on the above state representation and constraint set, a constraint-based alignment process is developed to normalize raw traceability observations.

First, each raw observation is mapped to the cyber layer according to module identity and the digital twin environment, generating an initial candidate interpretation of the observed state.

Second, the system constructs a candidate state set around the initial interpretation. Each candidate corresponds to a plausible zone-stage-position combination under the current cyber context.

Third, each candidate state is evaluated using a joint scoring function that balances observation consistency, temporal smoothness, and logical feasibility:

$$J(s_t^m) = \alpha D(s_t^m, o_t^m) + \beta H(s_t^m, \hat{s}_{t-1}^m) + \gamma \Phi(s_t^m) \tag{8}$$

where:

- $D(s_t^m, o_t^m)$ measures the deviation from the raw observation;
- $H(s_t^m, \hat{s}_{t-1}^m)$ measures temporal continuity with the previous aligned state;
- $\Phi(s_t^m)$ is the CPI-based constraint violation penalty;
- α, β, and $\gamma$are weighting parameters.

The final aligned state is selected as

$$\hat{s}_t^m = arg \min_{s \in \Omega_t^m} J(s) \tag{9}$$

In this way, the method does not simply smooth the sensed trajectory, but instead selects the most plausible state under both sensing evidence and cyber-physical process logic.

Finally, a confidence score is assigned to the aligned result according to the final alignment cost or violation level. A lower penalty indicates stronger consistency between the physical observation and the cyber-layer logic, and thus a higher confidence in the traceability result.

Through this process, the proposed method converts raw sensing observations into validated traceability states that are geometrically plausible, semantically interpretable, and consistent with the expected MiC logistics flow in the CPI framework.

5. Experiment and Validation

5.1 Experimental setup

In dense urban MiC projects, modules are typically produced off-site and delivered to construction sites either directly or via a prefabrication yard for temporary storage and dispatch coordination. A representative Greater Bay Area scenario involves module production in Nansha, transfer through a prefabrication yard in Yuen Long, and final delivery to construction sites in Hong Kong, as illustrated in Figure 2.

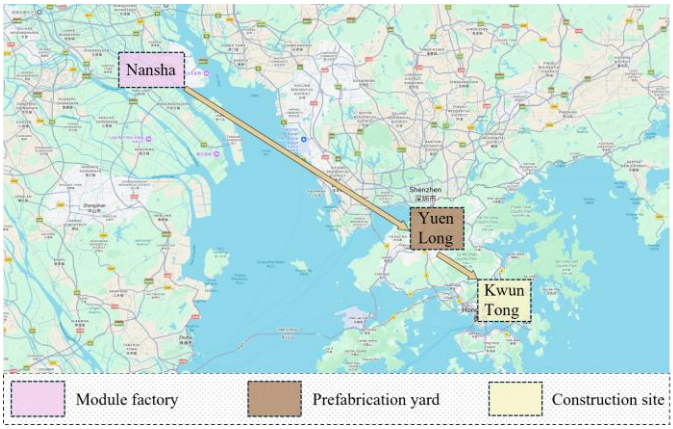


Figure 2. Application scenarios for MiC in the Greater Bay Area

To emulate this process, a cyber-physical experimental platform with three physical nodes—construction site, prefabrication yard, and factory—was developed. As shown in Figure 3, three adjacent tables represent these nodes physically, where the left, middle, and right tables correspond to the construction site, prefabrication yard, and factory, respectively.

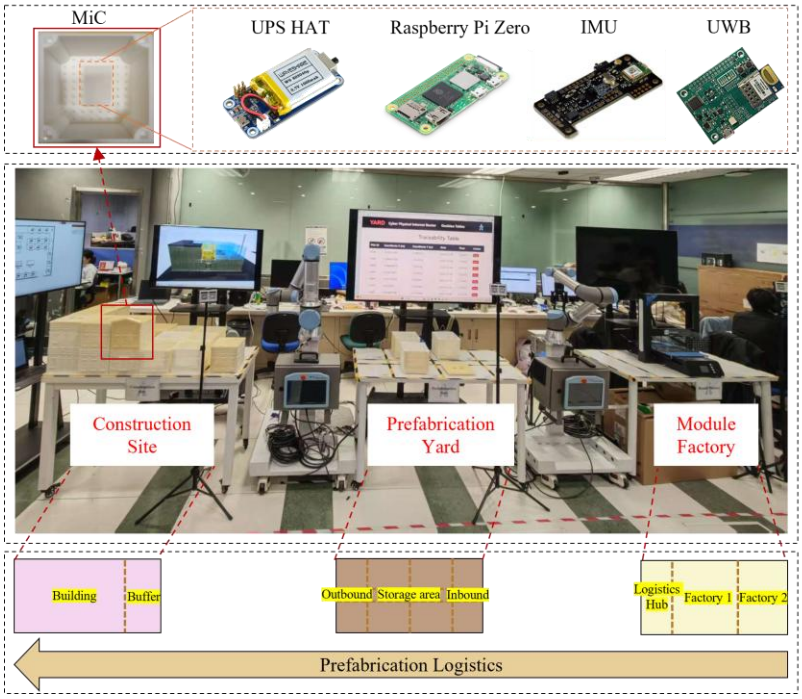


Figure 3. Cyber-physical experimental platform and its sensors

On the cyber side, a digital twin model of the experimental scene was developed using the three.js engine. The model reproduces the spatial layout and node configuration of the yard and construction site and provides the cyber representation used for state alignment.

To simulate MiC modules, 3D-printed modular units were adopted. Each module was equipped with a data acquisition unit integrating a Raspberry Pi Zero, an IMU, a UWB tag, and a UPS HAT power module. This embedded unit collected real-time position and orientation data and transmitted them to the cyber side for subsequent traceability interpretation and alignment, as shown in Figure 3.

Based on this setup, the proposed framework was validated by tracking module states across the factory, prefabrication yard, and construction site and by comparing raw observations with logically aligned traceability results.

5.2 Validation scenarios and data collection

To keep the validation concise while still demonstrating the effectiveness of the proposed method, two experiments were designed: a nominal process trial and a disturbance trial.

In the nominal process trial, the module followed a predefined logistics path: yard-storage → yard-outbound → on-site-buffer → on-site-building.

The module was kept stationary for several seconds in each zone, and after reaching the final on-site-building position, it remained stationary for an extended period to emulate the post-assembly state.

In the disturbance trial, the same path was repeated, but positioning disturbance was intentionally introduced during movement or after final placement. As illustrated in Figure 4, these disturbance-induced anomalies may include position drift, through-modeling, boundary crossing, and suspended states. The purpose of this trial was to examine whether the proposed CPI-based constraints could suppress semantically unreasonable traceability outputs under sensing instability.

For each trial, the raw sensing data were continuously recorded and processed in three ways: raw observations, conventional filtering, and the proposed CPI-based alignment method. The physical reference states of the four functional zones were defined in advance, enabling direct comparison between the recorded traceability outputs and the expected logistics states.

The dataset comprises continuous position observations for each module at a 10 Hz sampling rate, covering all four functional zones. Each zone includes multiple hold periods, with start and end timestamps, reference positions, and quality metrics. In total, the dataset contains 1,500 valid samples per module, enabling quantitative evaluation of position error, zone accuracy, stage consistency, logical violations, and post-assembly drift.

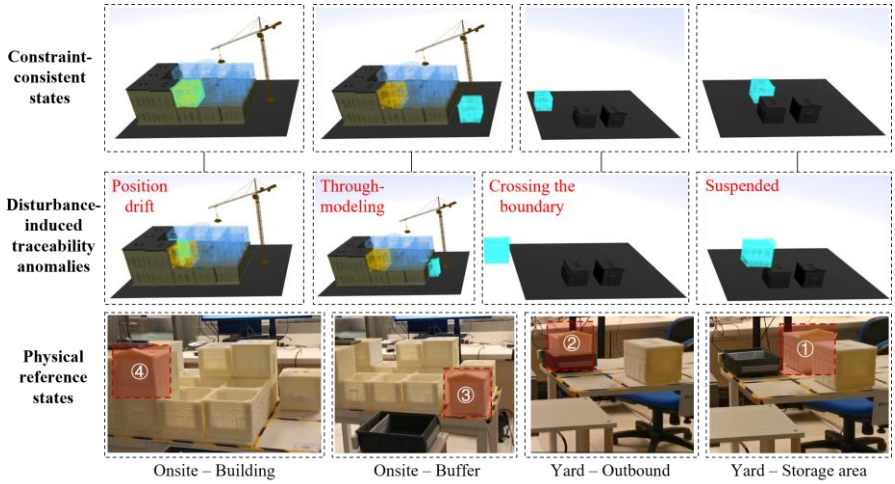


Figure 4: Nominal states, anomalies, and physical reference states used in the validation

5.3 Results and Analysis

The proposed method was evaluated by comparing raw observations, conventional filtering, and CPI-based alignment in terms of position error, zone accuracy, stage consistency, logical violation rate, and drift error. The results are summarized in Table 1.

Table 1: Overall comparison of traceability performance under different methods

Method	Position Error (cm)	Zone Accuracy (%)	Stage Consistency (%)	Logical Violation Rate (%)	Drift Error (cm)
Raw observations	10.819	81.653	97.315	18.347	5.434
Conventional filtering	10.575	82.056	100.0	17.944	5.158
Proposed method	4.706	100.0	100.0	0.0	0.0

As shown in Table 1, the proposed method achieved the best overall performance among the three methods. Compared with raw observations, the position error was reduced from 10.819 cm to 4.706 cm, while the zone accuracy improved from 81.653% to 100.0%. The logical violation rate was reduced to 0.0%, and the drift error after final placement decreased from 5.434 cm to 0.0 cm.

Compared with conventional filtering, the proposed method showed a clear advantage in semantic consistency rather than signal smoothing alone. In particular, the CPI-based constraints effectively eliminated invalid zone interpretation and suppressed post-placement drift, enabling the module to remain in the correct target state after final placement. These results demonstrate that the proposed framework improves both the geometric accuracy and the reliability of MiC traceability data. Beyond module-level calibration, the proposed method contributes to PI-enabled logistics by providing a node-level framework that can enhance coordination, standardization, and traceability across interconnected factory–yard–site networks.

6. Conclusion

This paper proposed a constraint-based spatial-temporal alignment framework for traceable MiC logistics in a CPI framework. By introducing CPI-based logical constraints, including temporal continuity, zone admissibility, stage sequence, and post-assembly stability, the proposed method aligns raw traceability observations with the expected logistics flow of MiC modules. The experimental results show that the framework improves traceability quality compared with raw observations and conventional filtering, particularly in terms of zone recognition, stage consistency, and stability after assembly.

This study was validated on a laboratory-scale cyber-physical platform and focused mainly on the yard and on-site digital twin environments. Future work will extend the framework to more detailed logistics states, larger-scale and real project scenarios, and richer sensing configurations. Further research can also investigate how the proposed framework supports real-time logistics monitoring and decision-making in CPI-enabled MiC systems.

References

Arshad, H. & Zayed, T. 2022. Critical influencing factors of supply chain management for modular integrated construction. *Automation in construction,* 144**,** 104612.

Cheng, T., Venugopal, M., Teizer, J. & Vela, P. 2011. Performance evaluation of ultra wideband technology for construction resource location tracking in harsh environments. *Automation in Construction,* 20**,** 1173-1184.

Dong, J., Lian, Z., Xu, J. & Yue, Z. 2023. UWB localization based on improved robust adaptive cubature Kalman filter. *Sensors,* 23**,** 2669.

Hussein, M., Eltoukhy, A. E., Karam, A., Shaban, I. A. & Zayed, T. 2021. Modelling in off-site construction supply chain management: A review and future directions for sustainable modular integrated construction. *Journal of cleaner production,* 310**,** 127503.

Li, C. Z., Zhong, R. Y., Xue, F., Xu, G., Chen, K., Huang, G. G. & Shen, G. Q. 2017. Integrating RFID and BIM technologies for mitigating risks and improving schedule performance of prefabricated house construction. *Journal of cleaner production,* 165**,** 1048-1062.

Lim, Y.-W., Ling, P. C., Tan, C. S., Chong, H.-Y. & Thurairajah, A. 2022. Planning and coordination of modular construction. *Automation in Construction,* 141**,** 104455.

Nkrow, R. E., Silva, B., Boshoff, D., Hancke, G., Gidlund, M. & Abu-Mahfouz, A. 2024. NLOS identification and mitigation for time-based indoor localization systems: Survey and future research directions. *ACM Computing Surveys,* 56**,** 1-41.

Pan, S., Trentesaux, D., Mcfarlane, D., Montreuil, B., Ballot, E. & Huang, G. Q. 2021. Digital interoperability in logistics and supply chain management: state-of-the-art and research avenues towards Physical Internet. *Computers in industry,* 128**,** 103435.

Ranjan, R., Shin, D., Jung, Y., Kim, S., Yun, J.-H., Kim, C.-H., Lee, S. & Kye, J. 2024. Comparative analysis of integrated filtering methods using UWB localization in indoor environment. *Sensors,* 24**,** 1052.

Tao, F., Qi, Q., Wang, L. & Nee, A. 2019. Digital twins and cyber–physical systems toward smart manufacturing and industry 4.0: Correlation and comparison. *Engineering,* 5**,** 653-661.

Wang, H., Zhang, H. & Liu, P. 2026. Adopting modular integrated construction (MiC) in road pavement construction: Concepts, applications, and challenges. *Automation in Construction,* 181**,** 106620.

Wu, H., Li, M., Yu, C., Ouyang, Z., Lai, K.-H., Zhao, Z., Pan, S., Wang, S., Zhong, R. Y. & Kuo, Y.-H. 2025. Towards cyber-physical internet: A systematic review, fundamental model and future perspectives. *Transportation Research Part E: Logistics and Transportation Review,* 197**,** 104051.

Yang, Y., Li, M., Yu, C. & Zhong, R. Y. 2024. Digital twin-enabled visibility and traceability for building materials in on-site fit-out construction. *Automation in Construction,* 166**,** 105640.

Yin, L., Li, T., Li, G., Besklubova, S., Zhong, R. Y. & Zhao, S. 2024. Cyber-physical internet based intelligent operation platform for off-site construction supply chain management. *Enterprise Information Systems,* 18**,** 2346743.

Zaalouk, A., Moon, S. & Han, S. 2023. Operations planning and scheduling in off-site construction supply chain management: Scope definition and future directions. *Automation in construction,* 153**,** 104952.

Zhong, R. Y., Peng, Y., Fang, J., Xu, G., Xue, F., Zou, W. & Huang, G. Q. 2015. Towards Physical Internet-enabled prefabricated housing construction in Hong Kong. *IFAC-PapersOnLine,* 48**,** 1079-1086.

Zhou, J. X., Shen, G. Q., Yoon, S. H. & Jin, X. 2021. Customization of on-site assembly services by integrating the internet of things and BIM technologies in modular integrated construction. *Automation in Construction,* 126**,** 103663.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Open Parcel Lockers as Physical Internet Nodes: A Framework for Accelerated and Sustainable Last-Mile Delivery

Ola Qasseer[1], Péter Bajor[1,2] and Norina Szander[1,2]
1. KTI Institute for Transport Sciences, Budapest, Hungary
2. Dennis Gabor University, Budapest, Hungary
Corresponding author: ola.qasseer@kti.hu

Abstract: *The rapid growth of e-commerce has intensified pressure on urban logistics and last-mile delivery, driving higher operational costs, traffic congestion, and environmental concerns. This paper examines Open Parcel Lockers (OPLs) as critical nodes within the Physical Internet (PI) framework to improve the efficiency and sustainability of last-mile delivery. Using the accessibility-based 4S framework (Slipper, Sneaker, Snack/Shop, Stop), we analyze the optimal distribution and governance of OPL networks. We base our framework on spatial observations of the parcel locker (PL) network in Budapest and compare it with the capital's urban structure. Our research shows that OPLs can serve as shared micro-hubs, improving delivery reliability and reducing redundant trips, thereby contributing to sustainable urban logistics. Ultimately, this study offers insights into the strategic deployment of OPLs to balance the competing demands of cost efficiency and accessibility in last-mile delivery, advancing the conversation on collaborative business models and shared logistics infrastructure.*

Keywords: *Open Parcel Lockers (OPL); Physical Internet (PI); Last-mile delivery; City logistics; Shared logistics infrastructure*

Physical Internet (PI) Roadmap Fitness: *☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories), ☐ Transportation Equipment, ☐ PI Networks, ☐ System of Logistics Networks, ☐ Vertical Supply Consolidation, ☐ Horizontal Supply Chain Alignment, ☐ Logistics/Commercial Data Platform, ☒ Access and Adoption, ☒ Governance.*

Targeted Delivery Mode-s: *☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

1 Introduction

The growth of e-commerce and on-demand delivery services has steadily increased pressure on urban logistics and last-mile delivery in cities. Rising operational costs, traffic congestion, failed home deliveries, environmental effects, and other issues are clear results of this growth (Quak, 2008). This is particularly true in a metropolitan environment such as Budapest, where the urban structure, the condition of the infrastructure, and logistic needs interact closely to shape transport solutions. However, there is no universal solution to urban logistics problems. Every city must assess based on its spatial urban structure, social aspects, and economic environment to determine which logistics approach offers a realistic and sustainable solution.

Several practices have been implemented in urban areas to reduce the pressure of urban logistics and decrease emissions. Last-mile solutions with cargo bikes, especially in “zero-emission zones” represent one of these options (Gruber et al., 2014). Electric light commercial vehicles are also applicable, as they reduce air and noise pollution while complying with emission limits

(European Alternative Fuels Observatory, 2025). However, the use of vehicles still creates the challenges often faced in urban freight transport, such as congestion and parking problems.

Considering that the first and last kilometers of delivery are the most expensive parts of the supply chain (Allen et al., 2012), implementing parcel lockers and other PUDO (Pick-Up & Drop-Off) points has emerged as a useful way to reduce last-mile delivery costs and manage the complexity of home delivery. The study of Edwards et al. (2010) showed that implementing these collection points could help reduce CO_2 emissions associated with failed deliveries. These lockers or collection points can serve as the final nodes in the physical internet network, offering shorter delivery routes and lower emissions than door-to-door delivery, provided that consumers access them on foot or via public transport (Carotenuto et al., 2018). However, the effective implementation of parcel locker systems requires careful consideration of their governance models and citywide distribution to ensure an optimal solution for both users and shipping companies. The core research problem addressed in this paper is the lack of a clear framework that simultaneously considers the optimal placement of parcel lockers in urban contexts and the governance mechanisms needed to coordinate their use among multiple stakeholders through the supply chain. In response to this gap, we are developing the 4S framework to capture the various dimensions of parcel locker placement decisions and to address the contradiction posed by the location preferences. We interpret these dimensions in open parcel locker (OPL) networks to mitigate last-mile delivery challenges and support the development of Physical Internet (PI) as a globally interconnected and standardized logistics system in which modular, shared, and interoperable infrastructure supports seamless freight movement, analogous to the digital Internet. A network of community PLs can be viewed as a multi-stakeholder ecosystem: the needs of multiple service providers and users must be coordinated; otherwise, the system could become unstable. Our research begins by identifying the main obstacles to determining the optimal locker locations and the governance models needed to organize their use among different stakeholders. It presents the current literature on the operation and location selection of parcel lockers. Then, we present our methodology through a spatial analysis of the current parcel locker network in Budapest, conducted as part of a collaborative research project within the framework of the Economic Innovation Partnership Program (EIPP), in cooperation with Budapest University of Technology and Economics (BME). We also propose a new analytical framework for the location and governance decisions of OPL networks as the final nodes of the physical internet. The final parts of the paper show the main results of our observations and the final recommendations for developing and implementing this kind of network.

2 Open Parcel Lockers: A Matter of Location and Governance

OPLs represent a promising solution for reducing last-mile delivery costs and managing the complexity of home delivery. They could be used by multiple service providers and governed to ensure optimal use by different stakeholders (Sina Mohri et al., 2024), leveraging the concept of the Physical Internet (PI). However, it remains unclear whether these lockers should be distributed in a centralized or decentralized manner. From an operational perspective, to attract consumers, the parcel locker network must be centralized (concentrated) because, in a price-sensitive market, this enables service providers to offer lower prices by offering shorter routes and delivering larger quantities of parcels simultaneously. Conversely, the parcel locker (PL) network must be evenly distributed across the municipality by population and density, well dispersed, and within walking distance of target consumers' homes, as consumers are willing to travel this distance (Ma et al., 2024). Although a distributed network diminishes the advantages of concentration - such as delivering 30 or more parcels to a single location - and may increase route length, service time, and labor requirements, these losses can be partially

offset in a well-designed (open) parcel locker network. Location decisions, as always, depend heavily on capacity and utilization rates. Without sufficient demand, high enough parcel flow will not result in the return on investment of installed locker capacities. Therefore, the trade-off between centralized and decentralized configurations cannot be assessed purely in spatial terms but must also account for demographic and (effective) demand patterns, locker capacity (decision on the proportions of inbound-outbound racks, sizing decisions), and utilization dynamics (overall frequency, seasonal patterns). A well-designed OPL network may partially mitigate these trade-offs by enabling shared resource use across multiple service providers, thereby increasing utilization rates and improving both operational efficiency and accessibility. Compared with home delivery, such systems can reduce unnecessary driving, congestion, eliminate parking search time, and avoid delays associated with customer interactions (e.g., waiting, answering the door, and completing handovers), ultimately improving environmental and operational performance.

The literature on PL systems shows that a key feature is their role as consolidation points within the urban delivery network. Instead of couriers delivering each package to a separate address, they deposit dozens of packages at once into a locker, where recipients can pick them up at their convenience after receiving a notification. Data collected by courier companies confirms that switching to lockers can reduce delivery time and address scattered customer locations, especially for multi-stop and multi-package routes (Ranjbari et al., 2023). Delivery time and ultimately operational costs are highly influenced by the locations of parcel lockers. Yang et al. (2020) investigated locker location decision models based on several factors related to construction, operation, and transport costs. Another study emphasized that the optimal location is connected to people's place of residence and work. Therefore, taking into consideration the function of an area can lead to somewhat informed location decisions. It is also an important aspect to consider retail locations and commuting hubs (Zhang et al., 2023). However, the locations of parcel lockers and their operating models should balance minimizing operational costs with maximizing customer attraction benefits (Dayal et al., 2025). A study from Poland showed that users prefer parcel lockers when the locations are distributed close to their commuting patterns between work and home (Lemke et al., 2016). Another experience, also in Poland, showed that the benefits of lockers disappear in poor locations with low traffic (Iwan et al., 2016). The need for customers to travel longer distances and the limited capacity of PLs could pose significant challenges to their use. In many cases, we see that individual providers install their PLs in the same locations, particularly right next to each other (closed model), thereby significantly oversizing the service infrastructure (especially during off-peak periods) at a single point within a given service area. Similarly, especially in large cities, these locations are often accessible only by car, making them inconvenient for those who rely on public transport (PT), bicycles, or walking.

On the other hand, the regulatory framework that encourages the spread of the “Locker first” model through tax incentives, permits for the use of public spaces, or infrastructure support is essential to making PL systems an integral part of sustainable urban logistics (Allen et al., 2012). From this perspective, the advantages of operating an OPL network can be summarized by its role in reducing the costs of duplicate development and the overburdening of public spaces. However, in a study from Germany, coordination challenges such as access rights, adequate availability, and maintenance responsibilities are the main obstacles hindering the operation of these networks. This is especially evident with the dominance of a main actor on the PL delivery option and the difficulties related to securing suitable locations, negotiating access conditions, and navigating administrative procedures in cooperation with municipalities (Thiele, 2023).

3 Methodology

In this paper, we propose an analytical framework that examines both the spatial distribution and the governance aspects of OPL nodes based on the "4S terminology (Slipper, Sneaker, Snack/Shop, and Stop)" which was suggested by some of the last-mile delivery experts. This study adopts an empirical, spatial analysis approach combined with a conceptual framework development, rather than a formal optimization or simulation model. The aim is to provide a structured decision-support tool for evaluating parcel locker placement in urban contexts. The 4S dimensions are defined as follows:

- Slipper distance refers to the area that users can access within a 5-minute walk from their homes.
- Sneaker distance describes the reach users have within a 5-10-minute walk from their residences.
- Snack/Shop distance encompasses the range that individuals cover on foot or by bicycle daily with minimal effort.
- Stop distance indicates the extent users can travel to a junction (transport hub) by bike, car, or public transport as part of their commuting routine.

Building on available concept reports and literature that emphasize walkable "slipper-distance" access and on-the-way pickup at daily activity nodes (e.g., transit hubs) (Orhan & Wallace, 2025; Ozyavas et al., 2025), we refine the location logic into a 4S framework (Slipper–Sneaker–Snack/Shop–Stop). While prior published knowledge typically provides single-threshold proximity rules or general guidance on accessibility and activity-node placement (minutes of walking, points of changing transport mode), the 4S approach operationalizes these insights into four distinct, comparable siting contexts (very-close home-based access; broader walkable catchment area; routine shopping/activity anchors; and transit-stop nodes). This refinement improves the transparency and offers a structured, decision-ready typology for consistent site screening and comparison. We illustrate this framework through our previous study conducted in Budapest, using a multi-step spatial analytical procedure, in search of optimal locations for OPLs, and we discuss the advantages, disadvantages, and challenges associated with this choice. This is done through several steps. The first step involves a spatial analysis of the current distribution of PLs in Budapest. The second step presents the level of accessibility to PLs based on population density and the proximity of commercial areas, public institutions, and transport hubs, and then identifies places within a 5-minute walk of direct access to PLs on a 629*925-meter grid raster map. The third step analyzes the current parcel locations and compares them with the estimated annual parcel demand, which is calculated from the total national demand and normalized according to the population in Budapest's catchment area on the same raster map (Uglik & Bóna, in press). It then proposes a theoretical framework to optimize the network and choose the best locations, using the 4S terminology, by aligning them with Budapest's urban structure. The final step extends beyond spatial analysis by offering recommendations for governing the OPL network as a common good, providing planners and policymakers with practical guidance to support the implementation of the PI.

4 Results and Discussion

According to spatial observations, the current density of PLs in Hungary's major cities is generally considered satisfactory. However, operators face significant fluctuations in demand. There are periods (typically during holidays) when PLs are overloaded, and deliveries are delayed, while during quieter periods, their utilization is extremely low. Sometimes, only a quarter of the available storage compartments are occupied. This suggests utilization rates can drop to nearly 25% during off-peak periods, highlighting substantial inefficiencies in capacity

allocation and indicating a mismatch between spatial supply and temporal demand patterns. In Budapest, the network of PLs covers approximately 65% of the capital's administrative area. At least one locker can be found within approximately a 700-meter radius in most of these areas. This coverage is concentrated in the central parts of the city, while outer districts, especially XXII, XXIII, and XVII, exhibit a more dispersed parcel locker distribution. At the same time, several providers have installed lockers in close proximity to one another in these districts, indicating strong potential for OPL solutions to improve spatial balance and overall accessibility (as shown in Figure 1, where differently colored dots represent the locations of PLs for different companies with an approximate total number of 1279 lockers in 2024). This implies that approximately 35% of the administrative area lacks direct access to parcel lockers. Furthermore, coverage is uneven, with a significant share of lockers concentrated in central districts.

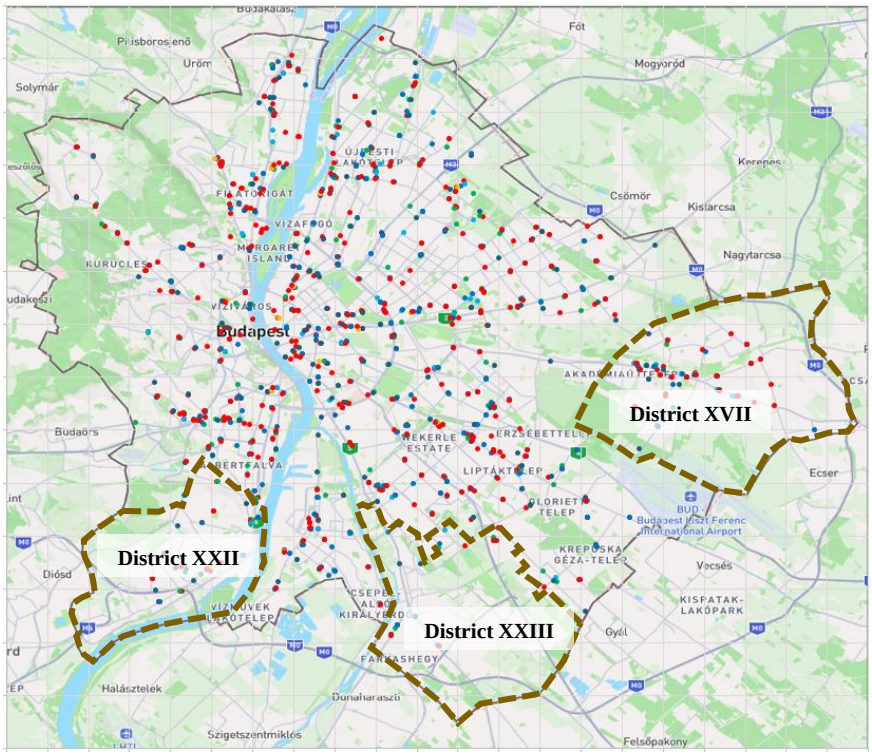


Figure 1: Distribution of Parcel Lockers in Budapest – Source: BME ALRT Engineering

Analyzing the level of accessibility to PLs showed that PLs can be categorized into three main categories:

- High performing locations with easy access: mainly the inner part of the city.
- Average performance locations: primarily located in residential areas or semi-peripheral locations.
- Poor performing locations with limited accessibility and unstructured layouts.

Figure 2(a) shows the distribution coverage of PLs among the proposed raster cells as follows:

- Red cells (A) represent the cells that include PLs and are also surrounded by cell(s) that include PLs.
- Orange cells (B) represent the cells that include PLs but are not surrounded by cells that include PLs.
- Yellow cells (C) represent the cells that do not include PLs but are surrounded by cell(s) that include PLs.

- Green cells (D) represent the cells that do not include PLs and are also not surrounded by PLs.

Figure 2(b) shows the number of compartments per cell that can be accessed from any location within the same cell. The places covered by these cells represent areas currently covered within 5 or 5-10 minutes (Slipper or Sneaker distances).

When looking at the annual estimated parcel traffic demand in Budapest (as shown in Figure 2(c)), several areas can be identified where, despite high population or service density, no PLs are available. As a result, Budapest exhibits four distinct types of zones in terms of the PLs (see Figure 2(b), for comparison):

- Locations with high accessibility and high demand.
- Locations with low accessibility and high potential for development and increasing accessibility.
- Locations with high accessibility and low utilization that need to be provided with new parcel lockers.
- Locations with low accessibility and low demand.

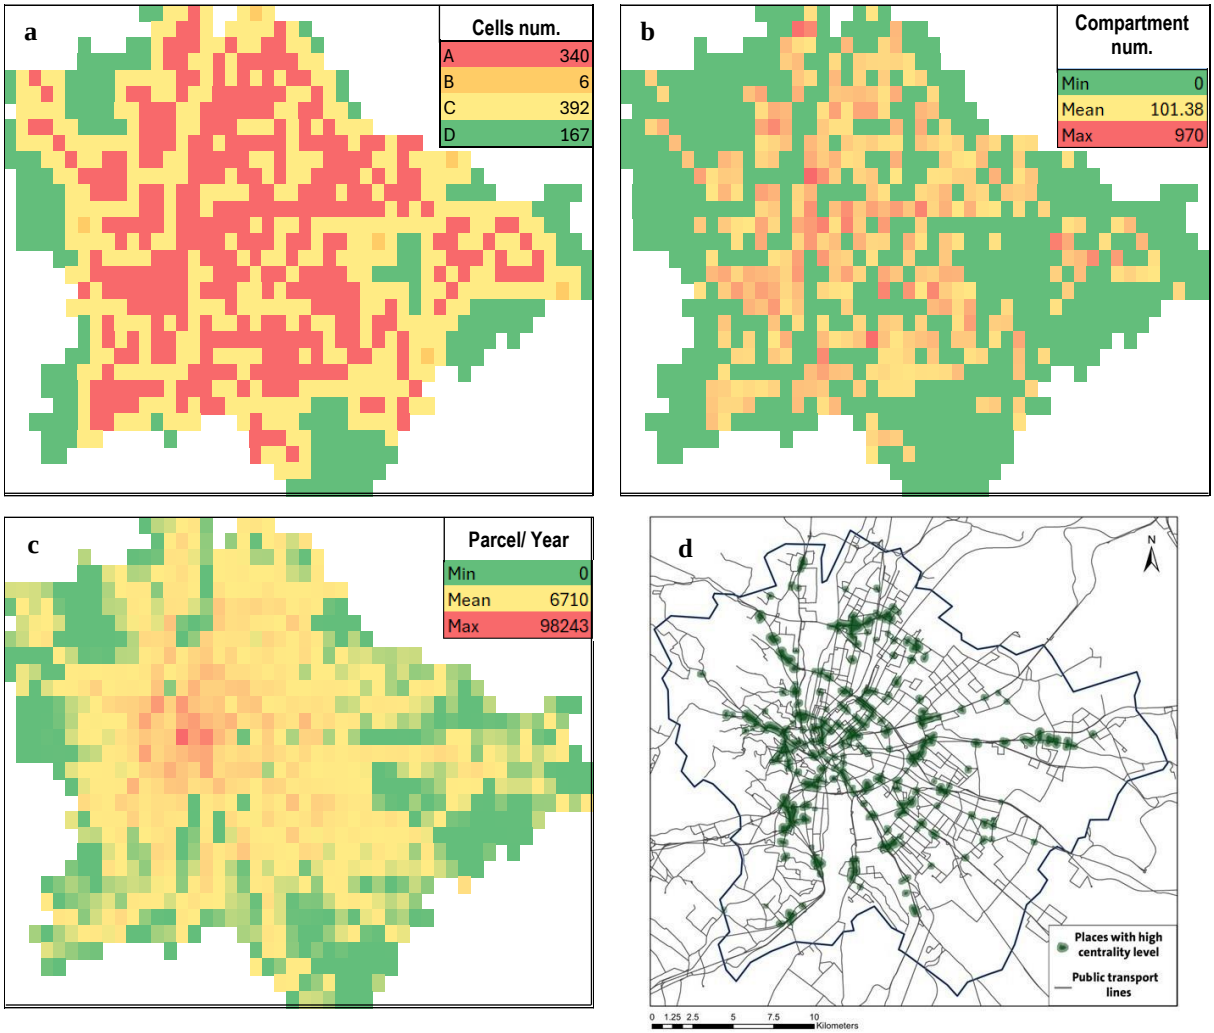


Figure 2: a) Parcel Lockers Accessibility according to Raster Categorization, b) Direct Accessible Locker Compartments per Raster Cells, c) Estimated Annual Parcel Demand, d) Possible Locations for Implementing Parcel Lockers within Snack/ Shop or Stop Distances – Source: a,b,c) BME ALRT , d) Qasseer & Munkácsy (2026)

Comparing Figure 2(b) and Figure 2(c) highlights the importance of implementing OPL networks to improve accessibility by providing the right number of locker compartments in underserved areas. On the other hand, it allows for the reorganization and redistribution of lockers to easily accessible locations without causing shortages or overcrowded compartments. The analysis shows that in some densely populated urban areas, providers install their lockers right next to each other, thereby spoiling the cityscape and causing wasted capacity and fluctuations. This also reinforces the viability of an urban freight strategy that employs an OPL network. As a result, some targeted measures can be identified in terms of development opportunities as follows:

- Installing new PLs in the areas with demonstrable demand but low capacity or no service.
- The use of integrated PLs accessible to every carrier, rather than the parallel presence of multiple providers.
- Including the residential areas located far from transport hubs, particularly in the southern and eastern parts of Budapest, in parallel with the “locker first” legal ecosystem.

From a managerial perspective, these findings translate into several actionable strategies. First, operators should prioritize demand-driven deployment by targeting high-demand, low-accessibility zones identified in the analysis. Second, transitioning toward shared (open) locker systems can significantly improve utilization rates by pooling demand across multiple carriers. Third, dynamic capacity management strategies can help mitigate peak-period congestion while reducing underutilization during off-peak periods. Finally, collaboration with urban planners is essential to integrate locker placement with land-use patterns and mobility hubs, ensuring long-term system efficiency.

The question about the optimal locations for these PLs has been addressed from several perspectives. The debate over centralized versus decentralized networks rests on the assumption that consumers prefer to collect their parcels near home; however, this assumption holds only for some recipients and shipments. Rephrasing the companies' arguments for deploying parcel lockers from a consumer perspective reveals an obvious contradiction: parcel lockers should be located both close to and far from consumers.

In our proposed framework, the parcel locker does not need to be near the residence but should be accessible to recipients who, depending on their daily travel routines–regardless of their intention to visit the parcel locker–change their location within the settlement frequently, mostly following predictable lifestyle patterns (going to work, school, sports, shopping, etc.). In this context, the 4S typology serves as a rule-of-thumb decision framework for classifying potential locker locations based on their accessibility profiles and integration with daily mobility patterns. Each category corresponds to a specific siting logic: Slipper and Sneaker locations prioritize proximity to residential areas and walking accessibility, while Snack/Shop and Stop locations are selected for their alignment with routine activities and transport flows. This enables a structured allocation of locker functions across the urban space rather than relying on uniform proximity thresholds. In the case of Budapest, as mentioned before, the current locations of the PLs accessible within a 5-10-minute walking distance (in Figure 2 (b)) represent the Slipper or Sneaker distances. Thinking of the other areas that need to be served to cover the annual parcel demand shown in Figure 2(c), and comparing it with Budapest’s urban structure, with its radial (ring-like) PT network, it can be seen from the perspective of serving those within Snack/Shop or Stop distances. Operationally, this means that areas identified as underserved in the spatial analysis can be matched with appropriate 4S categories: high-density residential gaps are assigned Slipper or Sneaker solutions, while areas with strong mobility

flows but lower residential density are better suited for Snack/Shop or Stop-based locker placement. Defining these locations can involve comparing parcel demand with the city's centralization process to identify the most powerful areas. This shall be done in future research with reference to a previous study that analyzed locations with high centrality levels according to PT networks (Qasseer & Munkácsy, 2026), as these places represent major transport hubs or centers that shape mobility patterns, and could be used to optimize locker capacities by making the pick-up process part of residents' travel habits. The framework we propose in this paper for selecting the best locations for OPLs and governing their operation is based on our spatial observations in Budapest. As shown in Figure 3, our framework establishes a conceptual link between PL accessibility levels and their functional roles in the urban environment. It shows how the four distance categories (Slipper, Sneaker, Snack/Shop, and Stop) correspond to different types of user access, ranging from home-based to transport-integrated access. These categories are directly affected by urban density, with higher densities supporting more centralized and easily accessible solutions (Slipper, Sneaker), while lower densities require controlled-dispersed networks (Snack/Shop, Stop) strategically integrated with daily activity locations and public transport nodes. In this way, dispersion is not left to random, market-driven placement decisions but is instead guided through planning strategies that aim to maximize accessibility and overall system benefits. When planning decisions about the best locations for the PL network are made, several aspects and detailed observations should be taken into consideration. These aspects contain analyzing urban density, mobility behavior, and land-use functions, emphasizing how daily activity patterns shape optimal locker placement. Overall, the framework demonstrates a structured approach to aligning PL distribution with spatial, behavioral, and functional urban characteristics. Importantly, the framework does not aim to compute a single optimal solution, but rather to guide decision-making by linking spatial characteristics (density, mobility behaviour, and land-use function) with appropriate locker typologies. In this way, the 4S approach acts as a heuristic model that translates urban structure and demand patterns into concrete location strategies.

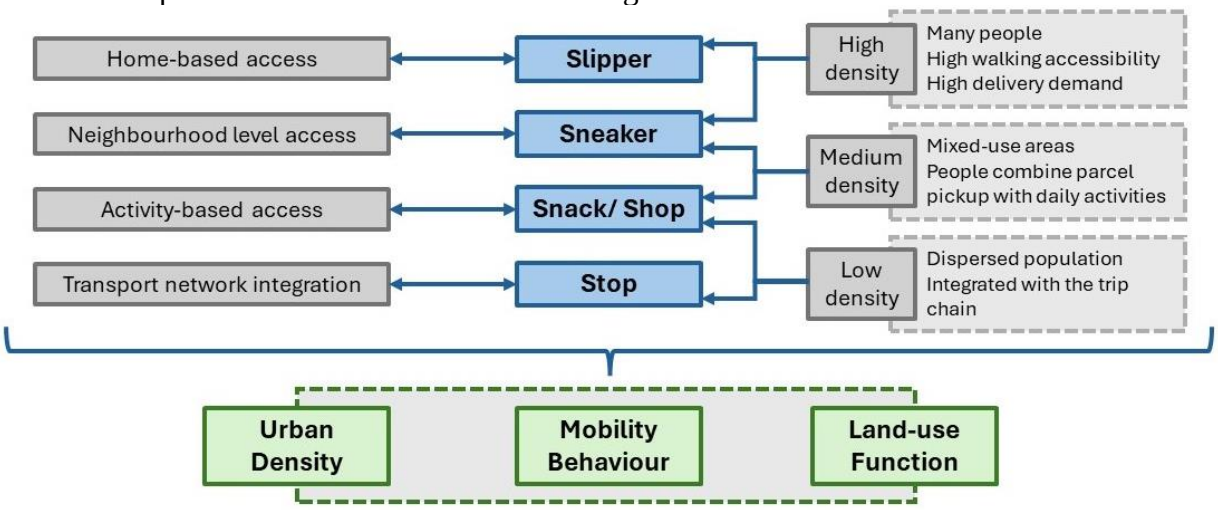


Figure 3: Suggested Framework for Choosing the Best Locations for Open Parcel Lockers – Source: Own elaboration

Table 1 shows the recommended distance for PLs by urban area type. However, without effective governance of this system, OPLs cannot work as intended to reduce trips and emissions and to provide a more efficient system aligned with the PI's notion (Qasseer et al., in press). If we think of these lockers as a common good, clear boundaries and regulations need to be established to coordinate and streamline their use by different stakeholders.

Table 1: Categorizing the 4S Dimensions for Choosing the Best Locations for Open Parcel Lockers – Source: Own elaboration

Area Type	Recommended Distance
Dense urban residential core area	Slipper distance (mainly), Sneaker distance
Residential area	Sneaker + Snack/Shop distance
Mixed-use area	Snack/Shop + Stop distance
Transport hub	Stop distance

Conclusion

Our paper highlighted practical implications for logistics operators, urban planners, and policymakers aiming to implement PI-compatible infrastructure. It also identified key considerations for selecting the locations of locker nodes as final points in the supply chain network, as well as the actions required for successful governance models. Future research may explore simulation-based evaluation of the proposed 4S framework using empirical parcel flow and mobility data to compare different OPL network configurations under varying demand and utilization scenarios. Such analyses could quantify the impacts of alternative placement and governance strategies on accessibility, delivery efficiency, locker utilization rates, and environmental performance.

Our proposed 4S dimensions in OPL deployment can support accelerated delivery services, sustainable, flexible local distribution, and decentralized pick-up/drop-off nodes. These dimensions capture key performance objectives, including delivery speed, parcel safety, customer convenience, network efficiency, and sustainability outcomes, which correspond to the findings of previous studies. When viewed through the lens of the 4S dimensions, a practical trade-off (“lockerland” vs. accessibility) is addressed: the 4S logic allows accessibility to be improved not only through dense deployment but also by prioritizing placements at routine activity points and mobility hubs, thereby reducing the risk of public-space overcrowding. The framework links the 4S dimensions to core PI principles, including openness, modularity, standardization, and resource sharing. It further highlights how OPL systems can accelerate and optimize last-mile delivery by enabling dynamic carrier collaboration, reducing failed delivery attempts, supporting multimodal logistics integration, and minimizing environmental externalities. It also bridges the gap between academic principles and decision support: the 4S framework provides a shared vocabulary for PL placement discussions among municipalities, operators, and service providers, because each location type can be linked to clear, measurable requirements (e.g., accessibility, capacity, and SLA targets).

Overall, the proposed framework highlights the importance of OPLs as essential components of a more efficient, customer-focused, and sustainable last-mile logistics ecosystem that aligns with the PI's goals and vision. Based on the analysis of the Budapest case, it is recommended to expand and optimize the PL network in a data-driven manner. This should be supported by IT development to enhance the digital interface by allowing flexible adjustments to pick-up time slots and optimizing compartment size in correlation with package dimensions. The software architecture for these lockers should enable integration with the logistics systems of individual service providers. Consider reducing the burden on the urban environment by placing the new PLs in locations that are easily accessible on foot or by bicycle. However, zone boundaries should be rethought and updated to make it easier to integrate logistics units. The

usage of public spaces needs to be more flexible, and partnership models need to be developed between local governments, logistics service providers, property owners, and transport planners.

References

- Allen, J., Browne, M., Woodburn, A., & Leonardi, J. (2012). The Role of Urban Consolidation Centres in Sustainable Freight Transport. *Transport Reviews*, *32*(4), 473–490. https://doi.org/10.1080/01441647.2012.688074
- Carotenuto, P., Gastaldi, M., Giordani, S., Rossi, R., Rabachin, A., & Salvatore, A. (2018). Comparison of various urban distribution systems supporting e-commerce. Point-to-point vs collection-point-based deliveries. *Transportation Research Procedia, EURO Mini Conference on 'Advances in Freight Transportation and Logistics'*, *30*, 188–196. https://doi.org/10.1016/j.trpro.2018.09.021
- Dayal, E., Unnikrishnan, A., Pani, A., & Lakshay. (2025). Optimal Location and Capacity Planning for Parcel Lockers: A Bi-Level Optimization Model Considering Planners' and Customers' Perspectives. *Transportation Research Record: Journal of the Transportation Research Board*. https://doi.org/10.1177/03611981241311564
- Edwards, J., McKinnon, A., Cherrett, T., McLeod, F., & Song, L. (2010). Carbon Dioxide Benefits of Using Collection–Delivery Points for Failed Home Deliveries in the United Kingdom. *Transportation Research Record*, *2191*(1). https://journals.sagepub.com/doi/10.3141/2191-17
- European Alternative Fuels Observatory. (2025). *Vehicles and fleet | European Alternative Fuels Observatory*. https://alternative-fuels-observatory.ec.europa.eu/transport-mode/road/european-union-eu27/vehicles-and-fleet
- Gruber, J., Kihm, A., & Lenz, B. (2014). A new vehicle for urban freight? An ex-ante evaluation of electric cargo bikes in courier services. *Research in Transportation Business & Management, Managing Freight in Urban Areas*, *11*, 53–62. https://doi.org/10.1016/j.rtbm.2014.03.004
- Iwan, S., Kijewska, K., & Lemke, J. (2016). Analysis of Parcel Lockers' Efficiency as the Last Mile Delivery Solution – The Results of the Research in Poland. *Transportation Research Procedia, Tenth International Conference on City Logistics 17-19 June 2015, Tenerife, Spain*, *12*, 644–655. https://doi.org/10.1016/j.trpro.2016.02.018
- Lemke, J., Iwan, S., & Korczak, J. (2016). Usability of the Parcel Lockers from the Customer Perspective – The Research in Polish Cities. *Transportation Research Procedia, The 2nd International Conference 'Green Cities - Green Logistics for Greener Cities', 2-3 March 2016, Szczecin, Poland*, *16*, 272–287. https://doi.org/10.1016/j.trpro.2016.11.027
- Ma, B., Teo, C.-C., & Wong, Y. D. (2024). Location analysis of parcel locker Network: Effects of spatial characteristics on operational performance. *Transportation Research Part E: Logistics and Transportation Review*, *192*, 103776. https://doi.org/10.1016/j.tre.2024.103776
- Orhan, C. C., & Wallace, S. W. (2025). E-commerce shipments in an X-minute city: Informing authorities on freight transport through parcel lockers. *Transportation Research Part A: Policy and Practice*, *199*, 104545. https://doi.org/10.1016/j.tra.2025.104545
- Ozyavas, P., Buijs, P., Ursavas, E., & Teunter, R. (2025). Designing a sustainable delivery network with parcel locker systems as collection and transfer points. *Omega*, *131*, 103199. https://doi.org/10.1016/j.omega.2024.103199
- Qasseer, O., Bajor, P., & Szander, N. (in press). From Chaos to Commons: Integrating Common-Pool Resources (CPR) into Physical Internet Models for Efficient Logistics. In D. Esztergar-Kiss, A. Munkácsy, & Á. Török (Eds), *Re-Generation in Transport:*

Proceedings of the 11th TRA Conference, 2026, Budapest, Hungary. Springer Nature Switzerland.

- Qasseer, O., & Munkácsy, A. (2026). Centrality through public transport networks: The case of Budapest and Vienna. *Journal of Urban Mobility*, *9*, 100198. https://doi.org/10.1016/j.urbmob.2026.100198
- Quak, H. (2008). *Sustainability of urban freight transport: Retail distribution and local regulations in cities*. Erasmus Research Institute of Management (ERIM). https://repub.eur.nl/pub/11990/EPS2008124LIS9058921543Quak.pdf
- Ranjbari, A., Diehl, C., Dalla Chiara, G., & Goodchild, A. (2023). Do parcel lockers reduce delivery times? Evidence from the field. *Transportation Research Part E: Logistics and Transportation Review*, *172*, 103070. https://doi.org/10.1016/j.tre.2023.103070
- Sina Mohri, S., Nassir, N., Thompson, R. G., & Ghaderi, H. (2024). Last-Mile logistics with on-premises parcel Lockers: Who are the real Beneficiaries? *Transportation Research Part E: Logistics and Transportation Review*, *183*, 103458. https://doi.org/10.1016/j.tre.2024.103458
- Thiele, S. (2023). *Kooperative Paketinfrastruktur für nachhaltige Zustellung: Anbieterneutrale Paketstationen*. WIK Wissenschaftliches Institut für Infrastruktur und Kommunikationsdienste GmbH. https://www.wik.org/fileadmin/user_upload/Unternehmen/Veroeffentlichungen/Diskus/2023/WIK_Diskussionsbeitrag_Nr_505.pdf
- Uglik, Z., & Bóna, K. (in press). Application of a Multi-Criteria Decision Support Model for the Consolidation and Sustainable Spatial Optimization of Parcel Locker Networks – A Case Study of Budapest. In D. Esztergár-Kiss, A. Munkácsy, & Á. Török (Eds), *Re-Generation in Transport: Proceedings of the 11th TRA Conference, 2026, Budapest, Hungary*. Springer Nature Switzerland.
- Yang, G., Huang, Y., Fu, Y., Huang, B., Sheng, S., Mao, L., Huang, S., Xu, Y., Le, J., Ouyang, Y., & Yin, Q. (2020). Parcel Locker Location Based on a Bilevel Programming Model. *Mathematical Problems in Engineering*, *2020*(1), 5263689. https://doi.org/10.1155/2020/5263689
- Zhang, S., V. Le, T., & Roorda, M. (2023). Optimizing Parcel Locker Locations in a City Crowd Logistics Network. *Transportation Research Record: Journal of the Transportation Research Board*, *2677*(12). https://doi.org/10.1177/03611981231167425

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE
Physical Internet Center

Proposed Framework for Modular Container Usage-Estimation based on Regional Contextualization

Priyali Bandla, Yinzhu Quan, Praveen Muthukrishnan, Benoit Montreuil
Physical Internet Center, Supply Chain and Logistics Institute,
H. Milton Stewart School of Industrial & Systems Engineering,
Georgia Institute of Technology, Atlanta, GA, USA
Corresponding author: pbandla3@gatech.edu

Abstract: *This paper proposes a data-grounded demand generation framework to support the design and evaluation of multi-tier hyper-connected freight networks. The framework generates modular container proportions by integrating socio-economic attributes, geographic context, commodity characteristics, leveraging an LLM-based agent, and seasonal patterns to produce operationally realistic freight demand streams. Grounded in observed freight flows and regional characteristics, it provides a scalable spatiotemporal distribution of modular containers of different sizes suitable for simulation and optimization models that assess consolidation strategies, network performance and sustainability outcomes. We provide an example for the sub-commodity Furniture freight across US counties to describe the framework.*

Keywords: *Physical Internet, Modular Containers, Agentic Models, Demand Generation, Large Language Models, Generative AI, Hyperconnected Systems*

Physical Internet (PI) Roadmap Fitness*: ☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories), ☒ Modular Containers, ☒ PI Networks, ☒ Large Language Models, ☒ Hyper-connected Networks*

Targeted Delivery Mode-s*: ☒ Paper, ☒ In-Person presentation*

1 Introduction

Freight logistic networks in the United States remain highly fragmented and are primarily operated by private stakeholders who preserve competitive advantage through territorial dominance or niche routing. This fragmentation is substantial, as reported by the American Trucking Association: 91% of carriers operate ten or fewer trucks even though trucks move approximately 72% of freight by weight annually nationwide. Because networks rely on proprietary and disaggregated data, shippers rarely share operational information, limiting system-wide coordination and efficiency. Truckload utilization remains low, with recent estimates showing that 58% of truckloads move partially empty, compared to 43% the previous year (American Trucking Association, 2025). As freight demand continues to grow, this under-utilization contributes to avoidable costs, emissions, waste (Mishra et al., 2014), congestion, delays, and environmental impacts. Although shipment consolidation yield system-wide benefits, particularly for small and medium-sized shippers, hesitation to share routes and logistics data remains a major barrier due to lack of protocols.

Physical Internet (PI, π)–inspired logistic networks offer a promising pathway to overcome these inefficiencies through consolidation, modularization, and cooperative routing. Standardized modular π-containers (Montreuil, 2011; Montreuil et al., 2014) enable seamless handling and transfers within hyperconnected logistic networks. However, designing and evaluating such systems is hindered by the absence of demand data expressed in modular container units, requiring translations into container equivalents to model π-enabled operations. Prior studies on designing PI-based large-scale regional (Garbers et al., 2025; Kim et al., 2025)

and metropolitan freight systems (Liu et al., 2025) have relied on simplifying assumptions to convert bulk tonnage flows into container flows, enabling tractable modeling and scalable optimization. While these abstractions supported high-level network design, they lacked the granularity needed to estimate hub sizing, throughput, and operational capacity requirements across tiers. As a result, determining realistic container volumes, staging needs, and processing capacities at continental, regional and gateway hubs remain challenging. These limitations stem in part from the coarse representation of shipment demand, which fails to capture variation in shipment size distributions and logistics behavior across industries and regions.

To address this limitation, a large language model (LLM)–based agent (Quan and Liu, 2024; Quan et al., 2025; Quan et al., 2026) is introduced as an interpretable refinement layer to enhance the estimation of county-level shipment demand. While the baseline model relies on aggregated freight statistics and county profiles, it lacks the ability to incorporate qualitative context and to explain deviations from expected patterns. The LLM-based agent addresses this gap by integrating structured data with contextual information that links manufacturing characteristics to shipment behavior. In particular, the agent can incorporate facility-level indicators and commodity-specific knowledge to assess whether estimated container volumes are consistent with local industrial capacity and product attributes, while also generating human-readable explanations of its inferences. This improves both the accuracy and interpretability of commodity-dependent demand generation.

Leveraging the interpretation provided by the agent, a region's characteristics and industry-specific behaviors are used to estimate shipment behaviors in units of modular handling and transport containers. The proposed framework can be generalized to a region broken up into geographic clusters, for example regional distribution centers to a national supply network (typically used for US logistics) or equivalently prefectures to country regions for East Asian countries, assuming the relevant data sources exist and are accessible. The scale of the framework is from a localized level where characteristics like housing, employment, and population data are easily traceable, to one-skip step in geography. In the case study we present, the estimation of modular container-usage behavior is localized to counties and generalized at for the US Southeast. We develop the approach for the furniture sub-industry within the larger manufacturing industry, where bulky and volume-constrained products provide a representative setting to evaluate container selection and shipment patterns. We generalize the proposed framework methodology, then provide an example through the case study, and present our results and future work.

2 Methodology

There are three main steps to the proposed estimation framework. First, the necessary data and tools are gathered to understand the characteristics at the local level. Since the estimation wants to more closely align with consumer characteristics, equal emphasis is placed on employment, housing, population, and industrial establishments in a geographical unit at the local level. These inputs are used to train and set up a large language model-based agent. The agent takes these vast resources and outputs a likelihood ranking on the modular handling containers most realistic to be used for the training industry at the generalized geographical area, we call a large region. The agent also provides a short reasoning for the ranking provided. The ranking is translated into the estimation model. Figure 1 summarizes the framework.

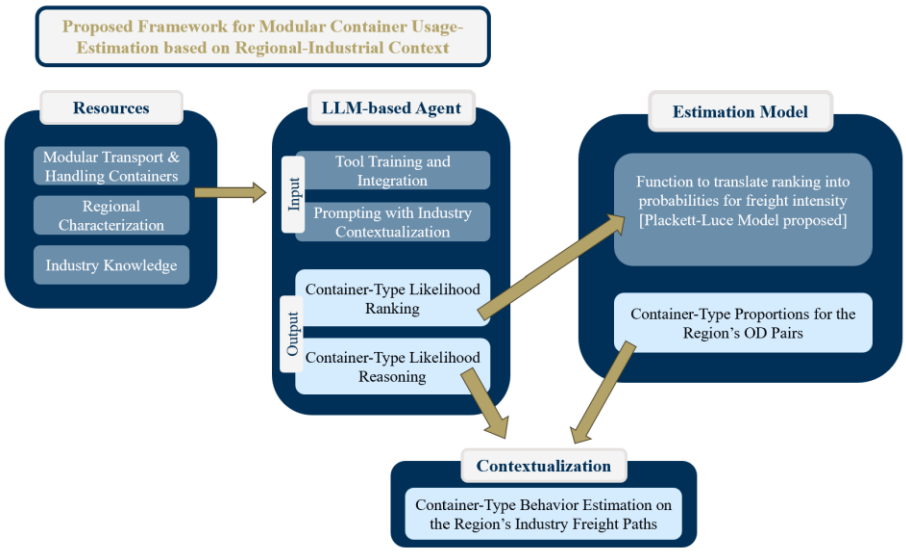


Figure 1: Proposed Framework for Modular Container Usage-Estimation based on Regional-Industrial Context

2.1 Modular Handling and Transport Containers

To standardize the representation of shipment scale, this study defines a set of modular container types that serve as abstractions of typical freight units. The container system consists of six categories (A–F), ranging from very small modular containers to full truckload (FTL) shipments. The six types are PI-modular handling and transport containers (Montreuil, 2011).

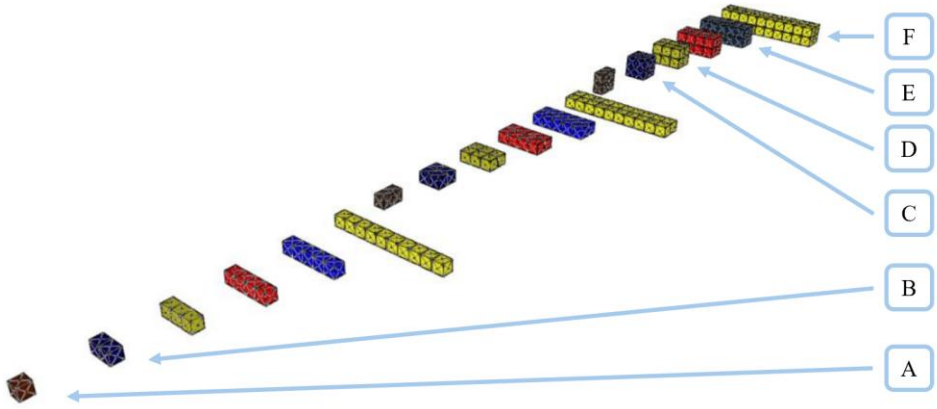


Figure 2: Modular handling and transport containers adapted from Montreuil et al. (2014)

Table 1: Container type dimensions

Container Type	X (m)	Y (m)	Z (m)	Volume (m^3)
A	1.2	1.2	1.2	1.73
B	2.4	1.2	1.2	3.46
C	6	1.2	1.2	8.64
D	6	2.4	1.2	17.28
E	6	2.4	2.4	34.56
F	12	2.4	2.4	69.12

Each container type is characterized by its approximate dimensions, typical use cases, representative cargo, and an estimated equivalence to full truckload capacity as shown in Table 1 adapted from Montreuil et al. (2014). We provide the dimensions of container A and E and

show that A-sized containers can build a container of size E in Figure 3. This design allows different shipment sizes to be expressed in a consistent and comparable framework.

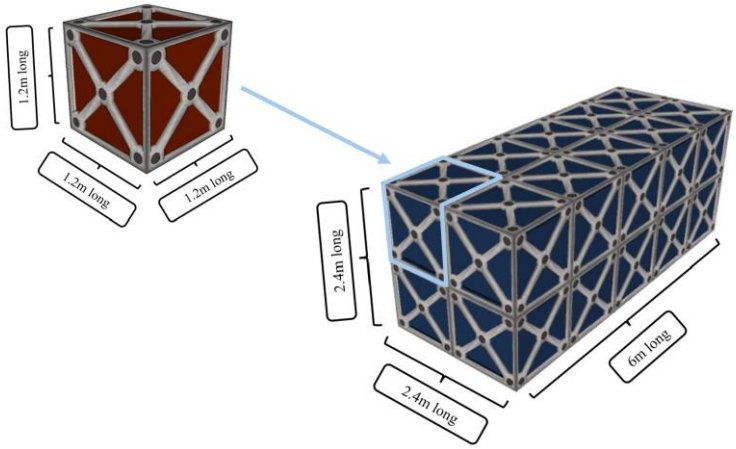


Figure 3: Container type A (red) and E (dark blue) adapted from Montreuil et al. (2014)

The modular container descriptions are not intended to represent specific real-world packaging standards, but rather to provide a simplified and scalable representation of freight movement. Future work will add size and weight constraints when determining the proportion of each container type based on industry norms as well as distinguishing between handling and transport containers. Smaller containers (A and B) correspond to compact, high-value shipments such as precision components, while medium-sized containers (C and D) capture pallet-scale or crate-based movements commonly observed in industrial supply chains. Larger containers (E and F) represent bulk shipments, and full truckload flows associated with large equipment or high-volume distribution. This structured abstraction enables the LLM-based agent to map inferred shipment characteristics to discrete container types, facilitating consistent reasoning and comparison across industries and regions.

2.2 LLM-based Agent

We develop an LLM-based agent to estimate the most likely freight container types used for shipments associated with a given manufacturing activity and county profile. The use of a large language model enables the integration of heterogeneous data sources and supports flexible, context-aware reasoning. In particular, the LLM is well suited for capturing semantic relationships, such as the connection between product attributes, including size and form, and shipment behavior, while providing interpretable and human-readable justifications rather than purely numerical outputs. By combining data retrieval, contextual understanding, and explainable inference, the agent infers shipment patterns and ranks container usage in a consistent and transparent manner. Domain-informed rules embedded in the reasoning process further guide this inference, such as prioritizing larger containers for bulky consumer goods and smaller containers for compact, high-value components. The agent then maps these inferred shipment characteristics to predefined container types and produces a ranked list of container usage likelihoods, ensuring that the results are both data-informed and interpretable.

2.3 Estimation Model

The ordinal ranking is transformed into a continuous score using an exponential decay function. The ranking defines the relative separation between container types. The lambda parameter controls the rate of decay, increasing the penalty on lower-ranked containers. The ranking-based scores are combined with a logarithmic transformation to define container-specific utility. Structural preference is captured through the ranking-based score while demand

intensity is incorporated by the logarithmic component. The log transformation stabilizes scale differences in OD-level commodity weights.

These utilities are converted into choice probabilities using the Plackett-Luce model yielding a probability distribution over container types, as the agent provides an ordered preference. Probabilities reflect relative likelihoods under both structural preferences and observed freight flows. The model maps utilities to probabilities within a consistent probabilistic framework. This provides stable and interpretable estimates of container selection probabilities, supporting downstream simulation and optimization.

To incorporate data-driven adjustments, weights are used to represent the empirical OD-level commodity distribution. Weights are calculated using freight intensity to shift container choice towards commodity-driven behavior reflecting shipment size, packaging requirements, and consolidation tendencies. By using weights, we prevent zero-probability outcomes, anchoring simulated behavior to the observed FAF distributions. In the future shipment tonnage will be incorporated to make the estimates more realistic.

2.4 Seasonality

For macro-level temporal variation, multiplicative scaling factors are applied to the utilities. To represent seasonal variation, three distinct behaviors- peak season with high demand, warm-up periods, the months immediately preceding peak periods, and low periods with reduced volume and more sporadic demand are included. Similarly, time-of-day regimes are used to influence the estimates, early morning, midday, and end-of-day. Early morning is characterized by smaller, more urgent shipments with lower levels of consolidation, while end-of-day reflects mixed shipments balancing residual demand. Time-of-day serves to add short-term variability to freight movement across counties reflecting responsiveness and demand-driven logistics. Additionally, high-priority shipments require faster delivery and involve less consolidation, potentially in smaller containers, although commodity characteristics may moderate this effect. Priority requests capture the trade-off between speed and consolidation efficiency. Seasonality, time-of-day, and priority introduce stochastic variability into container allocation proportions to meet the freight demand along county OD paths. It transforms static estimation into a context-aware model improving realism in downstream forecasting, optimization, and scenario analysis.

3 Case Study

The case study for the Southeast US region for the furniture manufacturing industry is provided. Due to the large number of counties in the US Southeast, we first made clusters of counties that share behaviors and characteristics. The clustering process and the data sources used are detailed below.

3.1 Data Sources

In order to create a consistent spatial description, we have used the standardized Federal Information Processing Standards (FIPS) county identification numbers, known as GEOID, and matted the data attributes to the counties.

3.1.1 Freight Data

The Freight Analysis Framework serves as the basis of freight movement to estimate the aggregated shipment flows between geographic regions. The database reports the total tonnage segmented by sub-commodities from a desired origin-destination (OD) pair. These flows are used to construct the county profile clusters (US DOT, 2017). We later apply the methodology

on these flows using a propagation factor to estimate the modular container flows between an origin-destination pair.

3.1.2 Socio-economic Data

We gather county-level socio-economic attributes from the United States Census Bureau's County Business Patterns. We pay attention to characteristics like population size, employment levels, income metrics, industry prevalence, and household data (US Census Bureau, 2024). Indicators such as total employment, number of establishments, average establishment size, and firm size distribution provide key signals for understanding production organization and shipment behavior. With these data points, we develop a county profile mapping used to group counties with similar economic structures, enabling efficient scaling and distributed patterns within a sub-commodity and season.

3.2 County Clustering

For each county, we created a feature vector that is constructed to balance its socio-economic and freight attributes. Historically, models have emphasized the population of a geographical region as the sole basis for distributing freight patterns (Asgari et al., 2025; Middela et al., 2024; Mishra et al., 2014). To better depict unique types of counties we use z-score standardization to improve downstream clustering.

3.2.1 County Profile Representation

We used k-means to group counties into profiles that help us apply seasonality and sub-commodity level patterns to influence freight flows. After statistical analysis we found that eight profiles balance the individuality of profiles and computational tractability. These profiles are interpreted as a summary of the feature vectors consisting of rural farmlands, urban consumption centers, metropolitan areas, or major logistics hubs among others. By clustering the 755 counties within the US Southeast we are able to reduce the high dimensional computation that would be required when disaggregating freight volume on time of day and seasonality. We provide the eight county profiles below. Instead of developing $755 \times 755 = 570{,}025$ relationships, we simplify $8 \times 8 = 64$ paths.

3.2.2 County Cluster Profiles and Descriptions

To capture regional variation in manufacturing structure, counties are grouped into predefined profiles based on their economic characteristics. Each profile represents counties with similar patterns in manufacturing activity, including establishment density, employment scale, and firm size distribution. A brief description of each county cluster is provided in Figure 3.

Profile	Structural Scale	Freight Intensity	Unique Characteristics	Manufacturing Role
0	Very Small	Very Low	~8.5K pop, ~1.1K employees, very low inbound/outbound. Slight outbound bias.	Peripheral feeder counties; small rural suppliers.
1	Small	High (relative to size)	Similar population to Profile 0 but ~5× freight volume. Freight-heavy for its size.	Specialized rural producers; small manufacturing/ag-processing nodes.
2	Lower-Medium	Balanced & Moderate	~21K pop, ~4.3K employees, inbound ≈ outbound (~1K). Freight proportional to size.	Local production + local consumption markets.
3	Medium	High	Similar pop to Profile 2 but ~5× freight. Strong two-way flows (~5K).	Regional industrial or redistribution nodes.
4	Upper-Medium	Moderate	~51K pop, ~13K employees, moderate freight (~2.2K). More residential scale than freight scale.	Regional consumption-heavy counties.
5	Large	Moderate (for size)	~153K pop, ~46K employees, freight ~5K. Large diversified economies.	Major regional economies; balanced manufacturing + consumption.
6	Medium-Large	Very High	~81K pop but ~15K inbound/outbound. Freight ~3× larger than Profile 5 despite smaller size.	Industrial freight powerhouses; manufacturing corridors.
7	Mega	Extremely High	~585K pop, ~257K employees, highest inbound/outbound (~25K).	National / multi-state logistics super-hubs.

Figure 3: County Profile Characteristics and Descriptions

This aggregation reduces noise while preserving meaningful differences across regions. For each profile, summary statistics are constructed using CBP data, including total employment, number of establishments, average establishment size, and firm size distributions. These indicators characterize whether manufacturing is fragmented or concentrated and serve as structured inputs for the LLM-based agent.

3.3 Prompting for the Furniture Industry

For the furniture case study, container types F (Full Truckload) and E are prioritized due to the typical size and dimensional characteristics of the products. The agent operates through a structured workflow that integrates tool-based data retrieval with contextual reasoning. Specifically, the agent is equipped with a set of external tools, including a container definition tool and a County Business Patterns (CBP) aggregation tool. Given a manufacturing type and a county profile identifier, the agent first invokes the container definition tool to retrieve standardized descriptions of container types, including shipment scale, typical use cases, and physical dimensions. It then calls the CBP-based tool to obtain economic indicators for the specified county profile, such as establishment counts, employment levels, payroll, and firm size distribution. These tools provide complementary information that jointly characterizes both shipment constraints and regional manufacturing structure.

The manufacturing activity in this case includes furniture, mattresses, lighting products, and related goods. Profile 7 is characterized by approximately 26,000 establishments and nearly 1 million employees, with a strong skew towards small firms- approximately 92%. This indicates a fragmented manufacturing base. However, the bulky nature of these products supports larger shipment units such as palletized loads and full truckloads. This creates a hybrid pattern of frequent shipments with relatively large physical volume. The sub-commodity of furniture is well suited for modular containerization, especially due to the growth in build-at-home and lean designs, a textbook opportunity for encapsulation (Montreuil, 2011).

Prompt for Agent (Condensed Version)

You are a freight container analysis agent responsible for estimating the most likely container types used for shipments associated with a given manufacturing activity and county profile. Your role is to retrieve and integrate structured information on container definitions and regional economic characteristics, and to interpret how manufacturing structure and product attributes influence shipment patterns and container usage.

The agent must identify and retrieve relevant indicators, including container specifications and County Business Patterns (CBP) statistics, and use them to infer shipment scale and logistics behavior. The analysis should be evidence-based and guided by domain-informed reasoning, where product characteristics such as size and form determine shipment constraints, and establishment structure provides signals on fragmentation and consolidation. The agent must produce a ranked list of container types (A–F) based on likelihood of use, along with concise and interpretable explanations, without relying on unsupported assumptions.

Figure 4: Condensed prompt for the LLM-based agent

3.4 LLM-based Agent Output

The ranking and contextualization given by the LLM agent are used for a sub-commodity on the OD path that is requested by the user. We receive the following output from the agent.

- Rank 1: F — Full Truckload container

Reason: Bulky consumer goods shipments (furniture, mattresses, large lighting fixtures) are often moved in full truckloads to retailers or distribution centers, maximizing efficiency for large-volume pallets or boxed items.

- Rank 2: E — (large equipment or bulk parts)
  Reason: Suitable for several pallet loads or single large items, matching the scale of large furniture components or multiple pallets of lighting/mattress components.

- Rank 3: D — (medium pallet shipments)
  Reason: Common pallet-scale shipments (1–2 pallets) for bulk items like larger lighting fixtures, mattress sets, or assembled components.

The ranking reflects the interaction between product characteristics and the underlying manufacturing structure of the county profile. The dominance of Full Truckload (FTL) as the top-ranked option is primarily driven by the bulky and low-density nature of the products, such as furniture, mattresses, and large lighting fixtures, which tend to fill transport capacity by volume and are therefore efficiently moved in full truckloads. The high ranking of the container type E further captures the prevalence of multi-pallet shipments and large individual items, which require substantial space but may not always justify a full truckload. Meanwhile, container D remains relevant for smaller-scale palletized movements, reflecting the fragmented structure of the manufacturing base, where many small establishments generate frequent but moderate-sized shipments. Overall, the ranking highlights a logistics pattern in which bulky product characteristics favor larger shipment units, while establishment structure contributes to the continued presence of pallet-scale transport.

3.5 LLM-based Agent Validation

We target four key ways to validate the use of the LLM Agent. First we run the agent multiple times using the same county profile and commodity input to evaluate the stability of the generated container rankings. Consistent rankings across the repeated runs indicate robustness and reliability of the inference process. We will compare the generated rankings and contextualization with assessments from logistics and manufacturing domain experts to have an expert-based validation. The experts can evaluate whether the inferred shipment behavior aligns with realistic freight operations, packaging practices, and consolidation behavior for the target industry. For example, bulky furniture products being associated with larger container types (E and F) should align with industry expectations. In the future, we will construct a manually labeled benchmark dataset consisting of commodity and county-profile combinations with expected container rankings by industry references. Then the rankings from the LLM can be evaluated using comparative metrics like Top-1 Accuracy, Spearman Rank Correlation, or Kendall's Tau against the industry rankings. Alternately, we can introduce a secondary independent agent to evaluate the same manufacturing and county-profile inputs to perform cross-agent validation. The rankings and contextual explanations generated by both agents can be compared to reduce single-agent bias and improve the confidence of the inferred container selection and behavior.

3.6 Seasonality and Time Scaling

During high demand periods, furniture shipments between Profile 7 counties intensify. Typically, in the spring and summer months the industry sees a surge in demand forming the peak season, with January to March acting as the warm-up period. During the warm-up period gradual increases in container types E and F are required to represent staged distribution. During reduced demand, slight increase in the share of container type D can be expected. For

priority adjustment, rushed furniture deliveries may increase the likelihood of container types E and D despite overall volume trends. Low-priority shipments for general restocking or inventory will reinforce full truckload shipments for cost-efficient transport. The model captures the influence or real-world levels for consolidation, seasonality, and delivery speed for bulky goods.

3.7 Results

Using the estimation framework designed for hyper-connected logistic networks promoting full truckload packing with modular containerization the following results are generated. Across OD paths within Profile 7, larger handling and transport container types, E and F, receive larger proportions, often accounting for about 45% to 60% of total allocation shown in Table 1 confirming the model's ability to capture bulk shipment behavior for furniture and high-volume flows. While container F (full truckload) frequently is given the largest proportion, the distribution remains balanced among type E and D indicating that the model preserves stochasticity and does not collapse to a single container type. OD paths with high priority show relatively higher shares of mid-sized containers E and D, compared to low-priority flows, reflecting lower levels of consolidation preference which reflects adjustments shippers make to meet urgent delivery windows. Despite the variability for profile 7 county paths, the model maintains structural consistency while balancing the influence of county profiles aligning with real-world freight behavior.

Table 2: Output for furniture freight between profile 7 counties during peak season

	Origin GeoID	Origin County	Destination GeoID	Destination County	Priority	A	B	C	D	E	F (Transport Containers)
1	22019	Calcasieu, LA	28047	Harrison, MS	Low	12.04%	11.38%	10.84%	16.40%	23.08%	26.25%
2	22033	East Baton Rouge, LA	12031	Duval, FL	Low	8.75%	9.69%	10.07%	15.19%	25.33%	30.97%
3	45051	Horry, SC	12099	Palm Beach, FL	High	11.25%	9.56%	12.02%	14.74%	20.34%	32.09%
4	12009	Brevard, FL	22033	East Baton Rouge, LA	Low	11.45%	11.74%	13.61%	15.38%	25.25%	22.57%
5	12057	Hillsborough, FL	22033	East Baton Rouge, LA	High	10.67%	12.98%	14.86%	17.08%	20.29%	24.11%
6	45091	York, SC	13051	Chatham, GA	Medium	10.81%	11.27%	13.03%	14.37%	22.81%	27.70%
7	13067	Cobb, GA	47187	Williamson, TN	Medium	11.79%	12.08%	11.90%	15.34%	23.78%	25.10%
8	45045	Greenville, SC	12117	Seminole, FL	Low	11.35%	8.83%	13.49%	13.20%	22.39%	30.75%
9	12071	Lee, FL	45083	Spartanburg, SC	Low	12.09%	13.33%	11.63%	15.55%	24.07%	23.33%
10	12086	Miami-Dade, FL	37051	Cumberland, NC	High	10.96%	14.51%	12.89%	16.27%	19.90%	25.48%
11	12086	Miami-Dade, FL	12009	Brevard, FL	High	13.63%	11.72%	12.88%	14.77%	19.33%	27.67%
12	45045	Greenville, SC	45045	Greenville, SC	Medium	10.52%	11.09%	12.58%	14.51%	25.37%	25.94%
13	12105	Polk, FL	45045	Greenville, SC	Medium	10.72%	8.90%	12.21%	14.15%	23.38%	30.64%
14	13135	Gwinnett, GA	22019	Calcasieu, LA	Medium	10.58%	13.61%	13.61%	16.81%	19.83%	25.56%
15	12101	Pasco, FL	12021	Collier, FL	Medium	11.54%	12.60%	12.39%	14.44%	23.52%	25.52%

4 Conclusion

In the case study, the container ranking given by the LLM-based agent, the model encodes structural logistics preferences. Translation between qualitative logistics knowledge and quantifiable containerization decisions is a fundamental requirement for implementing PI-motivated systems and operations (Montreuil, 2011). The ranking provided above captures consolidation behavior, where larger containers are favored for high-volume and bulky goods particularly for high-volume commodities and counties in Profile 7. Larger containers (F and E) reflect the physical internet system design goals of maximizing load factors to minimize empty space in a given truckload.

4.1 Contributions and Real-World Application

Based on the need for realistic, data-driven estimation of freight containerization for more full truckload transport, the model framework introduces an LLM-agent to balance freight volume and spatial, temporal, and socio-economic characteristics. The model uses FAF freight data to organize County-level structural profiles. The LLM-based agent translates these resources into

container likelihood ranking and contextualization. A Plackett-Luce model calculates probabilistic choices from the rankings enabling freight estimation to include sub-commodity influence, seasonality, time-of-day, and priority. The framework was developed to translate static freight estimation into context-aware probabilistic modeling.

From this estimation, design of networks and granular operations can be studied to enable realistic shipment generations at transport container level. While designing hyper-connected systems, the estimation serves as an input for network development, routing, and capacity planning. Since the model captures freight across different regions based on local characteristics, it can also be used for infrastructure investment decisions and freight analysis.

4.2 Limitations and Future Work

Currently, the LLM-based agent relies on inferred logistics patterns rather than trends derived from industry datasets, which makes it very difficult to qualify the rankings themselves besides with the contextualization output. This difference may explain the differences between short-term and long-term behavior given by the container estimation. Also, the aggregation due to Freight Analysis Framework flows may smooth out intro-profile variation across the counties.

To address the limitations, future development includes calculating the temporal weights based on shipment-level data which will also allow multi-peaks within a year. Comparing the simulated container distributions with observed logistics will enable a layer of validation for the model framework. For the LLM ranking, more specific training for each sub-commodity will enhance the contextualization itself. Moreover, the model will distribute the container type proportions over a time frame to mimic real-world operations. This refines the model from estimating static periods to continuous waves of shipments.

References

[1] American Trucking Association Department of Research (ATA), (2025). American Trucking Trends 2025, *American Trucking Association.*

[2] Asgari, A., Akbari, P., Demissie, M. G. (2025). Freight Trip Generation Modeling: A Narrative Review. *Transportation Research Record: Journal of the Transportation Research Board.* https://doi.org/10.1177/03611981251393157.

[3] Bahrami-Bidoni, Z., Montreuil, B. (2021). Predictive Demand Modeling for New Services in Hyperconnected Urban Parcel Logistics, *IPIC 2021.*

[4] CoStar Group. (2025). CoStar. https://www.costar.com/.

[5] Cristaldo, L., Petitdemange, E., Lauras, M., Montreuil, B. (2025). Utility-Driven Demand Estimation Framework for Regional Transport Systems, IFAC-PapersOnLine, Vol. 59, 94- 99, 10.1016/j.ifacol.2025.09.018.

[6] Flock Freight. (2025). New study finds 58% of truckloads moved over half empty in 2024, Press release from Flock Freight, https://www.flockfreight.com/company/pressreleases/new-study-finds-58-of-truckloads-moved-over-half-empty-in-2024.

[7] Fried, T., Tejada, C., Dennis-Bauer, S., Bolbaatar, O., Goodchild, A., Marshall, J. D., Olmedo, O., Carcia, L. (2025). Logistics of Zoning, Zoning for Logistics: Toward Healthy and Equitable Development for Urban Freight*, Journal of the American Planning Association*, https://doi.org/10.1080/01944363.2025.2515134.

[8] Garbers, A., Muthukrishnan, P., & Montreuil, B. (2025). Pragmatic Hyperconnected Freight Systems Design: A Data-Driven Case Study in the U.S. Southeast Megaregion. 11th International Physical Internet Conference.

https://repository.gatech.edu/server/api/core/bitstreams/0efed6e1-6487-4980-b202-d45365aa8120/.

[9] Hirata, E., Firdausiyah, N. & Kusumaningdyah, W. (2025). Handling modular containers in a physical internet environment. *Marit Econ Logist 27, 449–475 (2025).* https://doi.org/10.1057/s41278-024-00307.

[10] Kim, M., Karuppaiyan, N., Muthukrishnan, P., Montreuil, B. (2025). Hydrogen Refueling Facilities in Logistic Hubs across the Physical Internet: Adapting the Deployment to Vehicle Autonomy. *Proceedings of IPIC 2025 International Physical Internet Conference.* https://repository.gatech.edu/server/api/core/bitstreams/a9fd49a5-dedb-44c3-aad9-0b8c3cb565df/content.

[11] Liu, X., Muthukrishnan, P., Montreuil, B. (2025). Network Design and Capacity Management in Hyperconnected Urban Logistic Networks. *11th International Physical Internet Conference.* https://doi.org/https://doi.org/10.35090/gatech/13235.

[12] Middela, M., Ramadurai, G. (2024). Effect of the measurement period and spatial dependence on the accuracy of urban freight trip generation models, *Transportation Research Part A: Policy and Practice,* Volume 179, 2024, 103884, ISSN 0965-8564, https://doi.org/10.1016/j.tra.2023.103884.

[13] Mishra, S., Iseki, H., Moeckel, R. (2014). Multi entity perspective freight demand modeling technique: Varying objectives and outcomes, *Transport Policy,* Vol. 35, 176-185, 10.1016/j.tranpol.2014.05.002.

[14] Montreuil, B. (2011). Toward a Physical Internet: meeting the global logistics sustainability grand challenge, *Logistics Research*, Vol. 3, 2, 71-87, 10.1007/s12159-011-0045-x.

[15] Montreuil, B., Ballot, E.Tremblay, W. (2014). Modular Design of Physical Internet Transport, Handling and Packaging Containers, *Progress in Material Handling Research Vol. 13*, https://digitalcommons.georgiasouthern.edu/pmhr_2014/1/

[16] Quan, Y., & Liu, Z. (2024). Invagent: A large language model based multi-agent system for inventory management in supply chains. arXiv preprint arXiv:2407.11384.

[17] Quan, Y., Xu, Y., Chen, G., Bénaben, F., Montreuil, B. (2025). Leveraging Large Language Models for Risk Assessment in Hyperconnected Logistic Hub Network Deployment, *Proceedings of the IISE Annual Conference 2025*, 10.48550/arXiv.2503.21115.

[18] Quan, Y., Liu, Z., Benaben, F., & Montreuil, B. (2026). Leveraging large language models to enhance multi-agent risk assessment in supply chain networks. International Journal of Production Research, 1-29.

[19] Shaikh, S., Kim, N., Rababah, M., Montreuil, B., & Smith, J. (2023). Modular Containerization of Parcel Logistic networks: Simulation-Based Impact Assessment, *IPIC 2023.*

[20] US Census Bureau (2024). County Business Patterns, https://www.census.gov/programssurveys/cbp/data.html.

[21] US Department of Transportation, Bureau of Transportation Statistics, Federal Highway Administration (2017). Freight Analysis Framework Data Set (FAF*), US Department of Transportation*, 10.21949/1529116.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Physical Internet Center

Automatic Risk-Conscious Hyperconnected Logistics Location Selection via LLM-Based Multi-Agent Systems

Yinzhu Quan[1], Frederick Benaben[1,2], and Benoit Montreuil[1]

[1]Physical Internet Center, H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA, USA
[2]Centre Genie Industriel, IMT Mines Albi, Albi, France

Corresponding author: yquan9@gatech.edu

Abstract: *In volatile, uncertain, complex, and ambiguous logistics environments, supply chain planners are frequently required to select regions for purposes such as hub siting, network deployment, or infrastructure investment. Each of these situations requires integrating heterogeneous and cross-domain risk information under changing priorities. Yet existing approaches rarely adapt their analytical structure to the specific decision context, nor do they automatically, flexibly and selectively invoke appropriate analytical tools in response to varying decision needs. Instead, they typically generate fixed absolute scores, while decision-makers often require purpose-dependent and contextualized comparisons. To address these limitations, this paper proposes a LLM-based multi-agent architecture that organizes the assessment process around a user-defined purpose. Rather than relying on a fixed evaluation structure, the framework dynamically decomposes the decision task according to the specified objective, enabling a coordinating agent to invoke appropriate domain-specific agents and analytical tools as needed. By aligning analytical workflows with decision intent, the architecture supports adaptive, transparent, and explainable risk-conscious regional comparison across diverse supply chain planning contexts.*

Keywords: *Physical Internet; Supply Chain; Risk Assessment; Automatic Decision-Making; Multi-Agent Systems; Large Language Models; Generative AI*

Physical Internet (PI) Roadmap Fitness: ⊠ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ⊠ *PI Networks.*

Targeted Delivery Mode-s: ⊠ *Paper,* ⊠ *In-Person presentation*

1 Introduction

Global supply chains are increasingly exposed to structural uncertainty, climate variability, and systemic disruptions that challenge long-term infrastructure planning. As logistics networks expand across regions and become more interconnected, strategic decisions must account not only for economic efficiency but also for resilience and risk exposure. Within this context, risk-conscious logistics location planning has become increasingly critical in volatile, uncertain, complex, and ambiguous (VUCA) environments. Decisions such as hub siting, network deployment, and infrastructure investment shape long-term structural dependencies within supply networks and are exposed to heterogeneous risk factors including environmental hazards, extreme weather conditions, and transportation vulnerabilities. Because these risks may interact and propagate across interconnected regions, location selection cannot be treated as an isolated, single-domain evaluation problem but instead requires systematic and comparative cross-regional assessment.

As logistics networks become increasingly interconnected across regions, planning decisions are

embedded within complex structural dependencies that extend beyond individual facilities. The growing emphasis on shared infrastructure, modular coordination, and cross regional interoperability, which are consistent with the vision of the Physical Internet (Montreuil, 2011; Montreuil et al., 2012), further amplifies the need for systematic and risk-conscious planning mechanisms. In hyperconnected logistics systems, disruptions (Ivanov, 2022; Katsaliaki et al., 2022; Quan et al., 2023) or environmental pressures in one region may influence network performance elsewhere, making comparative and context-sensitive location selection essential. However, existing approaches to logistics location planning typically rely on fixed indicator sets or static scoring schemes, limiting their ability to automatically adapt evaluation structures to varying objectives and cross-domain risk considerations (Quan et al., 2026).

Recent work has begun exploring large language model (LLM)-based agents for inventory control, CRM automation, and risk-related analysis in supply chains (Quan and Liu, 2024b; Quan et al., 2025a,b, 2026). Building on these advances, we propose a LLM-driven multi-agent architecture for objective-oriented and automated regional comparison in logistics location planning. Rather than relying on predefined evaluation structures or fixed scoring schemes, the framework dynamically decomposes decision tasks according to user-defined objectives and coordinates specialized domain agents to retrieve, analyze, and synthesize heterogeneous evidence from live data sources, including environmental, weather-related, and transportation information. A central coordinating mechanism ensures consistency and comparability across regions while allowing domain-specific assessments to remain modular and independently structured. It further incorporates a self-criticism process to evaluate the completeness and consistency of outputs from Domain Agents, enabling the system to identify potential gaps, inconsistencies, and refine the final comparative assessment. Through structured comparison, selective analytical tool invocation, and transparent aggregation of risk factors, the proposed system enables context-sensitive, interpretable, and adaptive regional evaluation, thereby supporting robust and risk-conscious location selection within hyperconnected logistics networks.

Our main contributions in this paper are as follows:

1. **Objective-oriented logistics location selection:** We propose an objective-oriented and risk-conscious formulation for logistics location selection, where decision tasks are dynamically structured according to user-defined objectives. Instead of relying on predefined evaluation criteria, the framework adapts comparison dimensions to the specific decision context. This enables flexible and context-sensitive comparison between alternative regions in hub siting, infrastructure deployment, and strategic network planning.
2. **LLM-driven multi-agent coordination:** We develop a LLM-based multi-agent framework in which a Manager Agent coordinates multiple Domain Agents to perform domain-specific analysis and collect heterogeneous evidence. Domain Agents are systematically invoked to ensure comprehensive coverage, while analytical tools are selectively utilized based on requirements. The framework enables automatic orchestration of task execution, tool usage, and cross-domain reasoning, supporting structured aggregation of evidence.
3. **Explainable decision support:** The framework provides transparent and interpretable decision outputs by explicitly linking comparative conclusions to domain-specific evidence and intermediate analytical steps. It reveals how different risk factors contribute to the final assessment, making the reasoning process auditable and structurally explicit. This design enhances accountability and trust by grounding comparative decisions in consistent, evidence-based reasoning.

2 Related work

This section reviews prior work on supply chain risk management, multi-agent systems, and large language model (LLM)-based agents, highlighting the gap addressed in this study.

2.1 Supply chain risk management

In volatile, uncertain, complex, and ambiguous (VUCA) logistics environments, supply chain risk management (SCRM) has become a central concern for strategic and tactical decision-making (Ho et al., 2015; Katsaliaki et al., 2022). Location-related decisions, such as logistics hub placement, network expansion, and regional capacity allocation, are particularly exposed to disruption risks because they shape long-term structural dependencies within supply networks (Ivanov, 2017; Chu et al., 2020). Prior studies have shown that risks originating at specific regions or facilities may propagate through interconnected supply chain structures, leading to cascading and amplified impacts beyond local boundaries (Giannakis and Louis, 2011; Wu et al., 2012). As a result, effective location selection increasingly requires comparative assessment across regions, integrating heterogeneous risk factors related to infrastructure, operations, and network interactions (Giannakis and Louis, 2016).

2.2 Multi-agent systems for supply chain risk management

Multi-agent systems (MAS) have been widely adopted in SCRM as a modeling paradigm for capturing decentralized decision-making and the presence of heterogeneous, interacting stakeholders under uncertainty (Giannakis and Louis, 2016; Hassan et al., 2018). By representing key decision units within a supply chain as autonomous agents, prior research has shown that MAS-based formulations are well suited to characterize how localized disruptions can propagate through interconnected network structures and give rise to system-level risk outcomes (Bi et al., 2022, 2024). Compared with centralized or static modeling approaches, agent-based representations provide greater flexibility in describing interaction mechanisms, adaptive behaviors, and structural dependencies that are critical for analyzing complex and evolving risk scenarios (Ganesh and Kalpana, 2022).

2.3 LLM-based agents for risk-aware planning

Recent advances in LLMs have enabled a new class of agent-based systems that extend traditional multi-agent architectures with natural language reasoning, tool use, and autonomous coordination capabilities. In addition to domain-specific applications, recent studies have evaluated their reasoning, planning, and sequential decision-making abilities in economic and temporal settings, highlighting both their potential and limitations in structured problems (Quan and Liu, 2024a; Liu and Quan, 2024; Liu et al., 2025; Liu and Quan, 2025a). Related research on temporal event retrieval and tokenization further demonstrates how LLMs can represent and reason over event-driven processes linked to risk evolution and disruption dynamics (Liu and Quan, 2025b). Meanwhile, LLM-based MAS have been applied to automation and incident response scenarios, illustrating coordinated agent interaction and tool integration capabilities (Liu, 2025). Despite these advances, existing studies largely focus on task-level automation or isolated decision settings, leaving a gap in objective-oriented and systematically integrated planning frameworks that enable transparent, explainable, and risk-aware comparative evaluation across alternative regions. In particular, limited attention has been given to how agent interactions are structured as coordinated execution pipelines, including task decomposition, controlled agent invocation, and iterative evidence integration. This study addresses this gap by proposing an objective-oriented multi-agent framework with an explicit coordination pipeline that enables structured,

cross-domain comparative risk assessment for strategic location decisions.

3 Methodology

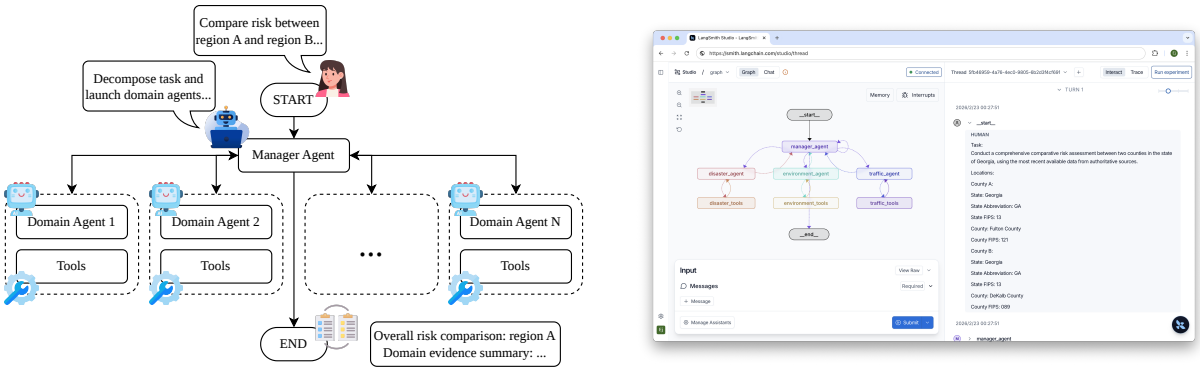
This section first introduces the problem setting and then presents the overall architecture of the proposed system, describing how large language model (LLM)-based agents are coordinated and supported by analytical tools for structured comparative evaluation.

3.1 Problem definition

In hyperconnected logistics systems, decision-makers are required to select or prioritize alternative regions for purposes such as logistics hub siting, network expansion, or infrastructure investment under heterogeneous and evolving risk conditions. Let a set of candidate regions be defined, each characterized by multi-domain risk attributes including environmental exposure, weather-related variability, and transportation infrastructure conditions. These attributes are heterogeneous in scale, source, and temporal resolution, and may exhibit interdependencies across regions. The problem is to comparatively evaluate candidate regions according to a user-defined planning objective, integrating cross-domain risk evidence while ensuring consistency, transparency, and contextual relevance. Unlike traditional location selection models that rely on fixed indicators or absolute scoring schemes, the decision context requires adaptive structuring of evaluation criteria and relative assessment aligned with the specified objective.

3.2 Framework

The proposed system in Figure 1a is implemented through a set of LLM-based agents organized into managerial and domain-specific roles. Manager Agent functions as the central coordinating component and organizes the overall process. It interprets the user-defined planning objective, decomposes the task into structured subtasks, and determines which Domain Agents should be activated at each stage of the pipeline for comparative assessment. In addition, Manager Agent integrates domain-level outputs into a consistent representation that supports cross-regional comparison. It monitors intermediate results and iteratively refines the analysis by re-invoking Domain Agents when necessary, forming a controlled execution loop that ensures coherent evidence integration and adaptive decision-making.



(a) Overall framework *(b) Agent pipeline demonstration*

Figure 1: Framework of LLM-driven multi-agent system for risk-conscious logistics location selection.

Domain Agents are responsible for conducting specialized risk analyses within specific domains. In this study, three primary categories of Domain Agents are defined: Disaster Agent, Environment Agent, and Traffic Agent. Disaster Agent evaluates hazard prevalence and disaster-related

exposure based on historical and climate-related events. Environment Agent assesses chronic environmental burdens and long-term exposure factors such as air quality and pollution indicators. Traffic Agent analyzes transportation infrastructure conditions, mobility patterns, and safety-related risks. Each Domain Agent leverages LLM reasoning capabilities to retrieve, interpret, and synthesize heterogeneous evidence from relevant data sources, producing structured outputs aligned with the comparative objective.

3.3 Analytical Tools

To support domain-level analysis, each agent is equipped with dedicated analytical tools that enable structured data retrieval and evidence integration, as summarized in Table 1. These tools include publicly accessible application programming interfaces (APIs), such as climate, disaster, and environmental data services, as well as structured databases containing transportation safety and infrastructure records. Domain Agents invoke these tools through structured queries, retrieve quantitative or categorical indicators, and transform raw data into interpretable domain-specific evidence. The invocation of analytical tools is governed by the coordination logic of Manager Agent and the reasoning process of Domain Agents, which determine when and how tools should be used based on intermediate results and analytical requirements. Overall, the architecture ensures modularity and extensibility, allowing tools to be selectively activated according to the planning objective while maintaining consistency across comparative evaluations.

Table 1: Analytical tools used in the LLM-driven multi-agent framework, with agent assignment, data source type, and functional descriptions.

Tool	Agent	Type	Description
Disaster Declaration Tool[1]	Disaster	API	Identify federal disaster declaration events
Multi-Hazard Risk Index Tool[2]	Disaster	API	Provide composite multi-hazard risk scores
Climate Tool[3]	Disaster	API	Query historical climate observations
Severe Weather Tool[4]	Disaster	API	Aggregate hazard event frequency and impact
Air Quality Tool[5]	Environment	API	Aggregate air quality exposure indicators
Environmental Facility Tool[6]	Environment	API	Identify regulated environmental facilities
Social Vulnerability Index Tool[7]	Environment	API	Assess relative social vulnerability levels
Traffic Accident Tool[8]	Traffic	Database	Aggregate traffic accident frequency and severity

3.4 Workflow

The proposed framework adopts a coordinated multi-agent architecture to support objective-oriented and comparative regional risk assessment. As illustrated in Figure 1a, a central Manager Agent receives the user-defined planning objective and structures the overall comparison task. Manager Agent decomposes the task into domain-specific analytical components and selectively invokes specialized Domain Agents responsible for environmental, weather-related, and transportation assessments. These Domain Agents operate and retrieve objective evidence through dedicated analytical tools. Manager Agent then aggregates domain-level outputs to ensure cross-regional consistency and structured comparison. This architecture enables mod-

[1] https://www.fema.gov/openfema-data-page/disaster-declarations-summaries-v2
[2] https://hazards.fema.gov/nri/data-resources
[3] https://www.ncdc.noaa.gov/cdo-web/webservices/v2
[4] https://www.ncdc.noaa.gov/stormevents/
[5] https://www.epa.gov/aqs
[6] https://www.epa.gov/frs
[7] https://www.atsdr.cdc.gov/place-health/php/svi/svi-data-documentation-download.html
[8] https://smoosavi.org/datasets/us_accidents

ular analysis while maintaining coordinated evaluation across heterogeneous risk dimensions, thereby supporting adaptive and risk-conscious logistics location selection.

4 Case studies

To demonstrate the practical applicability of the proposed multi-agent framework, we conduct a county-level comparative risk assessment within the state of Georgia, United States. The case study focuses on evaluating relative risk exposure between Fulton County and DeKalb County using the most recent authoritative data available across disaster, environmental, and transportation domains. This task reflects a realistic regional planning scenario in which decision-makers must compare alternative locations under heterogeneous and cross-domain risk conditions. The objective is to generate a structured, evidence-based comparison that supports transparent and explainable regional risk evaluation. The demonstration of the proposed agent pipeline can be observed in Figure 1b.

4.1 Manager Agent: Orchestration and Task Coordination

Manager Agent functions as the central orchestration component of the multi-agent framework, responsible for planning the assessment workflow, determining the sequence of domain consultations, and integrating heterogeneous domain-specific evidence into a unified comparative risk assessment. Unlike Domain Agents that operate within bounded analytical scopes, Manager Agent maintains global oversight, orchestrates task progression, integrates domain outputs, and produces a coherent, evidence-based cross-domain assessment for transparent decision support. The condensed version of the prompt used to guide Manager Agent is provided below.

Prompt for Manager Agent (Condensed Version)

You are the Manager Agent, responsible for orchestrating domain-specific agents to conduct a comparative risk assessment between Location A and Location B. Your role is to structure the overall evaluation process, determine which Domain Agents should be consulted, and integrate their outputs into a coherent cross-domain comparison. Domain Agents provide empirical evidence within their respective domains and do not influence coordination logic. At the beginning of the assessment, you formulate an analysis plan outlining the domains to be consulted and the strategy for integrating evidence. Domain Agents are invoked sequentially based on remaining evidence gaps. After collecting domain-level outputs, you synthesize the results into a structured comparative assessment that explicitly determines which location exhibits higher overall risk. The final output must be concise, evidence-based, and consistent across domains, supporting transparent and explainable decision-making.

Manager Agent begins by presenting an explicit initial analysis plan that identifies disaster, environment, and traffic as the core domains for comparative risk evaluation. The plan clarifies the analytical role of each domain and establishes a cross-domain integration framework based on metric normalization and transparent aggregation. Rather than invoking multiple agents simultaneously, Manager Agent adopts a sequential strategy, selecting one domain at a time according to remaining evidence gaps. In this instance, hazard exposure is prioritized as the foundational risk driver, and the workflow is routed to Disaster Agent for initial empirical assessment. This design constrains cross-domain integration through predefined rules, enhancing transparency and ensuring that comparative conclusions emerge from structured evidence rather than unconstrained reasoning.

4.2 Disaster Agent: Hazard and Disaster Risk Assessment

Disaster Agent operates as the hazard-specific analytical component within the multi-agent framework, responsible for evaluating exposure to natural and climate-related hazards across locations. It focuses exclusively on empirical disaster and hazard indicators, such as historical event frequency, severity characteristics, and associated impact metrics, without engaging in workflow coordination or cross-domain integration. By producing structured, data-driven comparisons of hazard exposure between locations, Disaster Agent provides domain-specific risk indicators that contribute to the overall comparative assessment. The condensed version of the prompt used to guide Disaster Agent is provided below.

Prompt for Disaster Agent (Condensed Version)

You are the Disaster Agent. Your sole responsibility is to execute the hazard and disaster risk analysis task assigned by the Manager Agent. You do not provide recommendations, workflow guidance, or planning suggestions. You must retrieve, analyze, and compare disaster- and hazard-related indicators for Location A and Location B, including historical disaster occurrences, hazard frequency and intensity, severity characteristics, and climate-related extreme event indicators where supported by data. Your objective is to characterize relative hazard exposure and disaster risk between the two locations using empirical evidence. Outputs must be structured, fact-based, and explicitly comparative to support risk assessment and decision-making contexts.

Table 2 presents the analytical tools invoked by Disaster Agent and the corresponding comparative hazard indicators for Fulton County and DeKalb County. The summarized results indicate that both counties exhibit high hazard exposure, albeit with distinct risk profiles. Fulton County records a slightly higher composite risk score and one additional disaster declaration since 2020, along with a marginally larger 2023 severe weather count and property damage. In contrast, DeKalb County presents substantially lower resilience despite a comparable risk score and similar disaster footprint. While hazard intensity and event patterns appear broadly similar across two counties, the resilience gap suggests potentially different recovery capacities under comparable hazard conditions. Overall, Fulton reflects slightly stronger exposure signals, whereas DeKalb's lower resilience indicates heightened structural vulnerability.

Table 2: Tools invoked by Disaster Agent and corresponding comparative hazard indicators for Fulton and DeKalb Counties.

Tool	Fulton County (Georgia)	DeKalb County (Georgia)
Disaster Declaration Tool (since 2020)	5 declarations; includes 2026 Winter Storm and 2024 Hurricane Helene events	4 declarations; similar Winter Storm and COVID-19 footprint
Multi-Hazard Risk Index Tool	Risk Score: 95.80 (relatively high); Resilience Score: 76.43 (higher resilience)	Risk Score: 94.27 (relatively moderate); Resilience Score: 33.40 (lower resilience)
Climate Tool (2023)	Total events: 379; Total property damage: $37.9M; dominant drought & heat	Total events: 334; Total property damage: $37.3M; includes 1 tornado
Severe Weather Tool (2023)	Drought & heat signals evident	Similar variability patterns

4.3 Environment Agent: Environmental Exposure Assessment

Environment Agent functions as the environmental exposure assessment component within the multi-agent framework, responsible for evaluating chronic pollution burdens and regulatory risk conditions across locations. It operates strictly within the environmental domain, focusing on empirical indicators such as air quality metrics, pollutant concentrations, toxic emissions, and

compliance records, without engaging in workflow coordination or cross-domain synthesis. By producing structured and data-driven comparisons of environmental exposure and long-term pollution risk, Environment Agent provides critical insight into sustained health-relevant and regulatory stressors that may amplify or interact with hazard-related impacts identified in other domains. The condensed version of the prompt used to guide Environment Agent is provided below.

Prompt for Environment Agent (Condensed Version)

You are the Environment Agent. Your sole responsibility is to execute the environmental exposure analysis task assigned by the Manager Agent. You do not participate in planning, prioritization, or decision-making. You must retrieve, analyze, and compare environmental condition indicators for Location A and Location B, including air quality metrics, pollutant concentration levels, toxic release and emissions indicators, regulatory compliance records where available, and measures of long-term environmental burden. Your objective is to characterize differences in environmental exposure and chronic pollution risk using validated and authoritative data sources. Outputs must be structured, evidence-based, and explicitly comparative to support environmental risk assessment and cross-location evaluation.

Table 3 presents the analytical tools invoked by Environment Agent and the corresponding comparative environmental indicators for Fulton County and DeKalb County. The summarized results indicate that Fulton County exhibits moderately higher ambient pollution signals across several indicators, including PM2.5 and CO concentrations, as well as consistently higher O_3 levels across most months. In contrast, NO_2 levels appear broadly similar between the two counties. With respect to emissions-related infrastructure, Fulton shows a substantially higher number of registered power-related facilities compared with DeKalb. Overall, while pollutant levels are comparable in magnitude, Fulton reflects stronger chronic exposure signals and a denser energy infrastructure footprint, whereas DeKalb exhibits relatively lower measured pollutant concentrations within the scope of the retrieved indicators.

Table 3: Tools invoked by Environment Agent and corresponding comparative environmental indicators for Fulton and DeKalb Counties.

Tool	Fulton County (Georgia)	DeKalb County (Georgia)
Air Quality Tool (2023)	PM2.5 $\approx$ 10.6 μg/m^3; CO $\approx$ 0.63 ppm; Higher O_3 across most months; NO_2 similar to DeKalb	PM2.5 $\approx$ 9.9 μg/m^3; CO $\approx$ 0.31 ppm; Lower O_3 AQI; NO_2 similar to Fulton
Environmental Facility Tool	Approx. 27 power-related facilities identified	Approx. 3 power-related facilities identified

4.4 Traffic Agent: Traffic Accident Risk Assessment

Traffic Agent functions as the transportation safety assessment component within the multi-agent framework, responsible for evaluating mobility-related risk and crash exposure across locations. It operates strictly within the transportation domain, focusing on empirical indicators such as accident frequency, severity distributions, fatality and injury counts, and other infrastructure- or mobility-related safety metrics, without participating in workflow coordination or cross-domain synthesis. By producing structured and data-driven comparisons of transportation risk, Traffic Agent contributes an essential exposure dimension that complements hazard and environmental analyses, enabling a more comprehensive and balanced cross-domain risk assessment. The condensed version of the prompt used to guide Traffic Agent is provided below.

Table 4 summarizes the traffic accident indicators retrieved by Traffic Agent for 2023. The results show that Fulton County experienced substantially more crashes than DeKalb County,

with approximately 2.2 times the total crash count. Although both counties exhibit a similar severity structure dominated by Severity 2 crashes, Fulton presents a slightly higher share of severe or fatal crashes and a marginally higher mean severity level. Overall, the comparative signal in this domain is primarily driven by differences in crash frequency rather than substantial differences in per-crash severity patterns.

Prompt for Traffic Agent (Condensed Version)

You are the Traffic Agent. Your sole responsibility is to execute the transportation safety analysis task assigned by the Manager Agent. You do not influence planning decisions or propose actions. You must retrieve, analyze, and compare transportation-related risk indicators for Location A and Location B, including traffic accident frequency, severity distributions (e.g., fatal, injury, property damage), and other mobility- and safety-related indicators supported by empirical data. Your objective is to characterize relative transportation and mobility risk between the two locations based on validated evidence. Outputs must be structured, fact-based, and explicitly comparative to support transportation risk assessment and cross-location evaluation.

Table 4: Tools invoked by Traffic Agent and corresponding comparative traffic accident indicators for Fulton and DeKalb Counties.

Tool	Fulton County (Georgia)	DeKalb County (Georgia)
Traffic Accident Tool (2023)	Total crashes: 2,393; Severity 2: 2,269; Severity 4: 124; Severe/Fatal share: 5.18%; Mean severity: 2.10; Approx. 2.2× higher crash count than DeKalb	Total crashes: 1,085; Severity 2: 1,041; Severity 4: 44; Severe/Fatal share: 4.06%; Mean severity: 2.08; Lower crash count relative to Fulton

4.5 Manager Agent: Cross-Domain Risk Synthesis

Manager Agent integrates disaster, environmental, and traffic domain outputs to produce a cross-domain risk judgment. Based on the synthesized evidence, Fulton County exhibits higher overall risk than DeKalb County. This is driven by stronger aggregate exposure in Fulton, including more disaster declarations, higher severe weather activity and losses in 2023, a larger crash count with a slightly higher severe and fatal share, and a higher density of emissions-related facilities. Although DeKalb shows lower resilience and greater recovery sensitivity, overall exposure across domains remains higher in Fulton. The assessment therefore identifies Fulton as the higher-risk location, with medium confidence given remaining environmental data gaps.

5 Conclusion

In this paper, we propose an objective-oriented and automatic LLM-driven multi-agent architecture for risk-conscious logistics location selection. By dynamically decomposing decision tasks according to user-defined objectives and coordinating specialized Domain Agents across disaster, environmental, and transportation domains, the framework enables structured, context-sensitive, and explainable cross-regional comparison. The proposed system supports adaptive tool invocation and transparent cross-domain integration through a coordinating Manager Agent. The Georgia county-level case study demonstrates how heterogeneous exposure, vulnerability, and mobility indicators can be synthesized into a coherent and traceable comparative judgment. The results highlight the importance of structured orchestration when evaluating multi-dimensional risk profiles in interconnected logistics systems.

References

Bi, M., Chen, G., Tilbury, D. M., Shen, S., and Barton, K. (2022). A model-based multi-agent framework to enable an agile response to supply chain disruptions. In *2022 IEEE 18th International Conference on Automation Science and Engineering (CASE)*, pages 235–241. IEEE.

Bi, M., Estrada-Garcia, J.-A., Tilbury, D. M., Shen, S., and Barton, K. (2024). Heterogeneous risk management using a multi-agent framework for supply chain disruption response. *IEEE Robotics and Automation Letters*, 9(6):5126–5133.

Chu, C.-Y., Park, K., and Kremer, G. E. (2020). A global supply chain risk management framework: An application of text-mining to identify region-specific supply chain risks. *Advanced Engineering Informatics*, 45:101053.

Ganesh, A. D. and Kalpana, P. (2022). Future of artificial intelligence and its influence on supply chain risk management–a systematic review. *Computers & Industrial Engineering*, 169:108206.

Giannakis, M. and Louis, M. (2011). A multi-agent based framework for supply chain risk management. *Journal of Purchasing and Supply Management*, 17(1):23–31.

Giannakis, M. and Louis, M. (2016). A multi-agent based system with big data processing for enhanced supply chain agility. *Journal of Enterprise Information Management*, 29(5):706–727.

Hassan, M. H., Mostafa, S. A., Mustapha, A., Abd Wahab, M. H., and Nor, D. M. (2018). A survey of multi-agent system approach in risk assessment. In *2018 International Symposium on Agent, Multi-Agent Systems and Robotics (ISAMSR)*, pages 1–6. IEEE.

Ho, W., Zheng, T., Yildiz, H., and Talluri, S. (2015). Supply chain risk management: a literature review. *International journal of production research*, 53(16):5031–5069.

Ivanov, D. (2017). Simulation-based ripple effect modelling in the supply chain. *International Journal of Production Research*, 55(7):2083–2101.

Ivanov, D. (2022). Viable supply chain model: integrating agility, resilience and sustainability perspectives—lessons from and thinking beyond the covid-19 pandemic. *Annals of operations research*, 319(1):1411–1431.

Katsaliaki, K., Galetsi, P., and Kumar, S. (2022). Supply chain disruptions and resilience: a major review and future research agenda. *Annals of operations research*, 319(1):965–1002.

Liu, Z. (2025). Autobnb: Multi-agent incident response with large language models. In *2025 13th International Symposium on Digital Forensics and Security (ISDFS)*, pages 1–6. IEEE.

Liu, Z., Nguyen, N. H., Quan, Y., and Zhang, S.-X. (2025). Temporal tokenization strategies for event sequence modeling with large language models. *arXiv preprint arXiv:2512.13618*.

Liu, Z. and Quan, Y. (2024). Tpp-llm: Modeling temporal point processes by efficiently fine-tuning large language models. *arXiv preprint arXiv:2410.02062*.

Liu, Z. and Quan, Y. (2025a). Econwebarena: Benchmarking autonomous agents on economic tasks in realistic web environments. *arXiv preprint arXiv:2506.08136*.

Liu, Z. and Quan, Y. (2025b). Retrieval of temporal event sequences from textual descriptions. In *Proceedings of the 4th International Workshop on Knowledge-Augmented Methods for Natural Language Processing*, pages 37–49.

Montreuil, B. (2011). Toward a physical internet: meeting the global logistics sustainability grand challenge. *Logistics Research*, 3(2):71–87.

Montreuil, B., Meller, R. D., and Ballot, E. (2012). Physical internet foundations. *IFAC Proceedings Volumes*, 45(6):26–30.

Quan, Y., Li, X., and Chen, Y. (2025a). Crmagent: A multi-agent llm system for e-commerce crm message template generation. *arXiv preprint arXiv:2507.08325*.

Quan, Y. and Liu, Z. (2024a). Econlogicqa: A question-answering benchmark for evaluating large language models in economic sequential reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 2273–2282.

Quan, Y. and Liu, Z. (2024b). Invagent: A large language model based multi-agent system for inventory management in supply chains. *arXiv preprint arXiv:2407.11384*.

Quan, Y., Liu, Z., Benaben, F., and Montreuil, B. (2026). Leveraging large language models to enhance multi-agent risk assessment in supply chain networks. *International Journal of Production Research*, pages 1–29.

Quan, Y., Pothen, A., and Montreuil, B. (2023). Predictive demand disruption signals for supply chain networks. In *IISE annual conference and expo. IISE. https://www. researchgate. net*.

Quan, Y., Xu, Y., Chen, G., Benaben, F., and Montreuil, B. (2025b). Leveraging large language models for risk assessment in hyperconnected logistic hub network deployment. *arXiv preprint arXiv:2503.21115*.

Wu, T., Huang, S., Blackhurst, J., Zhang, X., and Wang, S. (2012). Supply chain risk management: an agent-based simulation to study the impact of retail stockouts. *IEEE Transactions on Engineering Management*, 60(4):676–686.

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE
IMT Mines Albi
École Mines-Télécom

A Generic Meta-Model-Driven Simulation Engine for Configurable Supply Chain Structures

Camélia Bellepeau[1], Camille Durthaller-Renard[1], Benoit Morvan[1], Victor Romero[1], Guillaume Martin[1], Benoit Montreuil[2], Frederick Benaben[1,2]

1. Centre de Génie Industriel, IMT Mines Albi, Albi, France
2. H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, USA

Corresponding author: camelia.bellepeau@mines-albi.fr

Keywords: *Physical Internet, Hyperconnected Networks, Meta-model-driven Simulation, Discrete-time Simulation, Topology-agnostic Modelling, Configurable Supply Chain.*

Abstract: Supply chain simulation is widely used to analyze complex logistics systems, yet existing models remain largely case-specific and difficult to reuse across heterogeneous configurations. In parallel, the Physical Internet (PI) promotes open, hyperconnected logistics networks, but reusable end-to-end simulation frameworks aligned with these principles remain limited. This paper proposes a generic, meta-model–driven discrete-time simulation engine that decouples supply chain structure from execution logic. Entities are modeled as unified logistics nodes connected through configurable relations, enabling the simulation of production, inventory, and multimodal transport across diverse network topologies. Experiments on networks ranging from 8 to 60 entities show execution times below one second for medium-scale scenarios. The approach supports reusable, topology-agnostic modeling of supply chain networks.

Physical Internet (PI) Roadmap Fitness: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),☐ Transportation Equipment, ☒ PI Networks, ☒ System of Logistics Networks, ☐ Vertical Supply Consolidation, ☐ Horizontal Supply Chain Alignment, ☐ Logistics/Commercial Data Platform, ☐ Access and Adoption, ☐ Governance.

Targeted Delivery Mode-s: ☒ Paper, ☐ Poster, ☐ Flash Video, ☐ In-Person presentation

1 Introduction

Supply chain simulation has grown significantly as a research field, with a substantial surge in publications since 2019, reflecting the increasing complexity and dynamism of global logistics networks (Korder et al., 2024). Despite this progress, a persistent limitation lies in the limited reusability of simulation models across different supply chain configurations. Existing studies predominantly develop context-specific models tailored to particular cases, while general-purpose simulation tools (e.g., AnyLogic, SimPy) require substantial effort to design and adapt models for each new scenario. As a result, models tend to be large, time-consuming to build, and rely on non-standardized development processes, which complicates replication and hinders systematic comparison across heterogeneous supply chain structures (Cimino et al., 2010).

In parallel, the Physical Internet (PI) paradigm, introduced by (Montreuil, 2011), proposes a fundamental rethinking of global logistics toward open, hyperconnected, and modular networks. While a substantial share of PI research has concentrated on logistics network design, transportation, hubs, and urban freight systems, reusable end-to-end simulation frameworks

that jointly capture production, procurement, inventory, and multimodal transport across heterogeneous PI-enabled configurations remain limited. Existing tools either target single-firm, closed supply chains, or require case-specific recoding when the network topology, ownership structure, or transport configuration changes.

This paper addresses this gap by proposing a generic, meta-model-driven discrete-time simulation engine for configurable supply chain structures, separating structural definitions from the execution logic, that support multi-owner, open topology configurations where resources are shared across actors. The remainder of the paper is organized as follows: Section 2 reviews related work, Section 3 describes the simulation model, Section 4 presents computational performance experiments, and Section 5 concludes with future work directions.

2 Related Works

Supply chain simulation has been extensively studied through discrete-event simulation (DES) approaches. DES has established itself as a dominant paradigm in logistics and supply chain management research, with a growing number of studies published since 2019 demonstrating its applicability to a wide range of operational problems (Bottani & Casella, 2024). A variety of simulation environments have been adopted, including commercial platforms such as AnyLogic, Arena, FlexSim, and Simio, as well as programming-based frameworks such as SimPy (Cimino et al., 2010; Pegden, 2007; Pinho et al., 2016). However, despite their flexibility, these tools often require case-specific model development, where system structure and behavior are embedded within each implementation.

Agent-based simulation (ABS) is a modeling paradigm in which a system is represented as a collection of autonomous, interacting agents, each following behavioral rules, whose local interactions give rise to emergent system-level dynamics (Macal & North, 2010). ABS enables the modeling of dynamic supply chain configurations that adapt to market changes, where connectivity and competition between agents significantly impact performance indicators (Hernández & Pedroza, 2016). However, similarly, ABS models are typically developed for specific scenarios, which limits their reuse across heterogeneous supply chain configurations without major re-engineering.

In the context of the Physical Internet (PI), simulation has played a central role in assessing the potential of hyperconnected logistics networks. (Hakimi et al., 2012) developed a large-scale “mobility web” simulator representing interconnected logistics infrastructures, enabling simulations involving thousands of daily orders. (Sarraj et al., 2014) used simulation to evaluate the performance of interconnected logistics networks and protocols, reporting improvements in transport efficiency. More recently, (Kaboudvand & Montreuil, 2024) proposed a hybrid discrete-event and agent-based simulation platform for hyperconnected urban logistics, enabling the representation of multiple stakeholders and supporting scenario analysis across strategic, tactical, and operational decision levels.

Despite these advances, simulation models remain largely tailored to specific supply chain configurations, limiting their reuse across heterogeneous scenarios. This is partly due to the need to embed context-specific assumptions, structures, and data within each model implementation. These limitations highlight the need for more generic simulation frameworks that can support multiple supply chain structures within a unified modeling approach.

The proposed engine addresses this need by decoupling the structural definition of the supply chain from the simulation execution logic through a formal meta-model. The proposed framework targets end-to-end supply chain dynamics, capturing key dimensions of supply chain dynamics within a unified, topology-agnostic architecture. This enables comparative

analysis across heterogeneous supply chain configurations without case-specific recoding, supporting the analysis of interconnected logistics networks and the evaluation of alternative configurations.

3 The Simulation Model

The proposed simulation framework constitutes a discrete-time, event-driven modelling paradigm tailored to represent heterogeneous supply chain architectures within a unified computational framework.

3.1 General Architecture and State Evolution

The system evolves over a predefined temporal horizon by iteratively generating successive state representations of the supply chain. At each discrete time increment, the entire system configuration, composed of entities (actors), material flows, inventory levels, production statuses, transportation states, backlogs, and pending deliveries, is systematically updated via a **deterministic state-transition mechanism**.

To enhance computational efficiency, the simulation engine adopts a state-preservation strategy, where the subsequent time step is initialized through a deep-copy operation of the preceding state, followed by selective modifications of state variables. This approach ensures temporal continuity of the dynamic attributes.

Mathematically, the system's evolution can be formalized as a discrete dynamical system, where the state at time t, denoted as $X(t)$, transitions to $X(t+1)$ via the application of a fixed transition operator T, such that $X(t+1) = T(X(t))$. Concretely, T encodes the sequential execution of procurement, production, transportation, and inventory update rules applied uniformly across all entities in the configuration.

Moreover, the operator T is structurally invariable across diverse supply chain configurations, and thus its underlying logic remains independent from the configuration. Structural modifications, such as the integration of new distribution hubs or alterations to transport networks, are confined to adjustments in configuration parameters and topological definitions, thereby preserving the underlying simulation logic.

3.2 Actors as Logistics Nodes (π-Nodes)

Within this framework, supply chain entities are formalized as actors derived from a unified abstract entity class, where each actor embodies a logistics facility endowed with the functional capabilities of storage, transformation, and/or routing of goods. Conceptually, these actors align with the **π-node** abstraction, as they serve as intermediary facilities where goods undergo temporary storage, modal transfer, or consolidation prior to their subsequent progression through the distributed logistics network.

Table 1: Actors, roles and functions

Actor	PI concept (Montreuil, 2011)	Function
Entity (abstract)	π-node - standard interface	Standard interface: order, receive, schedule, produce, deliver, update forecast
Supplier	π-store + π-gateway	Receive orders, deliver products
Maker	π-composer + π-store	Receive demand and orders, schedule procurement and production, produce goods, deliver products, update forecast

Transfer Centre (DC or FC)	π-store + π-sorter	Receive demand and orders, schedule procurement, deliver products, update forecast
Hub	π-hub	Receive orders, deliver products
Customer	π-gateway	Generate demand, receive orders, update forecasts

Within the conceptual framework of each actor, a group of operational components, composed of warehouses, production lines, backlog registries, and transport allocations are integrated to cover different functional roles. Specifically, production lines serve as representations of transformation workflows occurring at nodal points, whereas warehouses embody storage and buffering functionalities. This cohesive modelling paradigm facilitates the abstraction of diverse facilities under a unified node-based structure, thereby enabling a standardized and homogeneous representation across heterogeneous operational environments.

3.3 Flows, Relations, and Transport Corridors

In this simulation framework, actor interactions are resolved through flow objects that contain the exchange of information and physical goods. Entities engage in interactions by emitting and receiving flows stored in shared data structures, which are dynamically managed by the simulation engine.

The system's interactions are structured around four primary flow categories: **demand flows**, **production flows**, **delivery flows**, and **forecast flows**.

- **Demand flows** are defined as product orders initiated by downstream nodes, which, upon receipt by upstream nodes, are transformed into backlog entries, composed of product specifications, quantities, customer identifiers, and delivery constraints. These backlogs, representing pending orders awaiting fulfilment, serve as the principal driver for subsequent delivery decisions.
- **Production flows** represent the internal transformation of materials within manufacturing nodes. Input stocks are consumed according to bill-of-material relationships and production line capacities, yielding finished goods that are subsequently packaged into π-containers and made available for outbound dispatch. Production flows thus bridge inbound delivery flows and outbound shipment preparation, and are generated exclusively by transformation nodes such as Makers.
- **Delivery flows** take care of the physical transit of goods between actors, through which the initiated shipments are instantiated as delivery objects and scheduled based on the travel time associated with predefined transport itineraries. Deliveries are ordered by their expected arrival step, ensuring accurate synchronization with the simulation timeline. Each delivery object - encapsulating product specifications, quantities, routing constraints, and transport mode - is conceptually analogous to a π-container (Montreuil et al., 2010).
- **Forecast flows** propagate anticipated future demand upstream through the network, enabling actors to pre-emptively generate procurement and production plans prior to actual order placement. The propagation of forecasts relies on bill-of-material (BOM) relationships and supplier connectivity, translating downstream demand projections into upstream material requirements.

The **commercial connectivity** between actors is formalized through actor relations, which delineate supplier-customer compatibility, operational parameters (selling prices, customer priority rankings, and permissible delivery delays ...), and thus define the logical topology of the supply chain network.

Transport connectivity is modelled via transport relations and itineraries, delineating the physical routes available for shipments between actors. An itinerary comprises one or more transport segment, each characterized by travel time and transport mode, with the cumulative travel time determining the delay between shipment departure and arrival. These itineraries model the π-network layer, representing the physical and contractual corridors through which π-containers are relayed across successive π-nodes.

Transport resources are represented by transport models, which encapsulate vehicle capacities, operational costs, carbon emissions, and performance attributes. Upon scheduling a delivery, the simulation engine selects an appropriate transport model based on a configurable strategy, which may prioritize capacity optimization, cost minimization, or emissions reduction.

3.4 Step Execution Sequence and Strategy Integration

As previously mentioned, the system's evolution adheres to a deterministic, time-stepped operational framework, with each simulation iteration deriving from the previous system state through discrete updates.

The first step is the resolution of in-transit deliveries: shipments scheduled for arrival at the current simulation step are dequeued based on their arrival time and allocated to the respective receiving entities. After deliveries, backlog records undergo consolidation, where expired orders identified via predefined expiration thresholds are flagged and purged, and redundant entries are agglomerated together.

For each time step (T), the engine then executes two sequential passes over all entities. The first pass propagates from end-customer to root supplier: each actor evaluates incoming demand, synthesizes production and procurement schedules, and generates procurement orders. These orders are injected into the demand pipeline and relayed to upstream suppliers, thereby cascading demand signals through the network. The dominant cost per entity (E) in this pass is *O(AR * P)*, driven by the need to match incoming demands against actor relations (AR) and product catalogs (P). The second pass propagates in reverse, from root to end: each actor receives delivered materials, executes production routines, adjusts inventory balances, and initiates outbound shipments along the predefined transport corridors. The delivery function must select transport routes across actor and transport relations (TR) for each product, giving a dominant cost of *O(AR * TR * P)* per entity. Combining these, the theoretical **worst-case complexity** of the engine would be ***O(T * E * AR * TR * P).***

This modular and topology-agnostic design of the execution logic enables the simulation framework to represent different logistics architectures, ranging from centralized hub-based configurations to decentralized, hyperconnected systems.

4 Case study

4.1 Case Study description

The study is conducted on four supply chain configurations, comprised of respectively 8, 15, 30 and 60 actors, over varying temporal horizons (10, 20, 100, 500 or 1000 steps). The scenarios share most of the operational parameters. The simulated supply chain produces a single product (skateboards) assembled from two components (wheels and wooden boards) using a predefined bill of materials and production line. Economic parameters such as purchase prices, selling prices, and transportation costs remain constant across scenarios. Four transportation modes, truck, rail, sea, and air, are utilized for inter-node logistics, while storage policies, warehouse capacities, and demand forecasts for recurring customers across recurring timesteps are the

same across scenarios. Consequently, the configurations differ only in their structural characteristics, namely the number of nodes and the density of actor and transport relations. As network size increases, the number of possible interactions and routing paths grows, resulting in increasingly complex logistics structures. Table 2 summarizes the main characteristics of each configuration.

As the number of nodes increases, both actor relations and transport relations grow significantly, reflecting a progressive increase in network interconnectivity and routing possibilities. The multiplication of operational relations leads to a substantial increase in interaction complexity across scenarios.

Table 2: Network Characteristics Across Scenarios

Scenario	Nodes	Actor Relations	Transport Relations
A	8	7	7
B	15	10	21
C	30	28	32
D	60	56	108

4.2 Computational Performance

The computational performance of the simulation engine was evaluated by measuring the runtime required to execute scenarios of increasing network size or time steps. All experiments were executed on a workstation equipped with an Intel Core i9-12900H processor (2.50 GHz) and 32 GB of RAM, running a 64-bit operating system.

The computational complexity of the simulation is theoretically influenced by three main factors: the number of simulation steps, the number of entities, and the number of interactions between entities.

4.2.1 Complexity with Respect to the Simulation Horizon

To evaluate the impact of the simulation horizon on computational performance, experiments were conducted with increasing numbers of simulation steps. Scenarios were executed with 10, 100, 500, and 1000 simulation steps, while keeping the network configuration constant (30 entities case-study). Each configuration was run 30 times. Outliers were removed using Tukey's fence method (Q1 − 1.5×IQR, Q3 + 1.5×IQR). The average, median, Interquartile Range (IQR), min, and max (expressed in seconds) of the remaining runs are reported.

Table 3: Statistic analysis by number of steps

Nb of steps	10	100	500	1000
Average	0,1160	1,3889	7,6525	14,3060
Median	0,1155	1,3920	7,2950	14,3510
IQR	0,0173	0,0987	1,3460	0,2698
Standard Deviation	0,0100	0,0533	0,6424	0,2196
Min	0,1040	1,3130	6,9850	13,8210
Max	0,1340	1,4820	8,8820	14,6500

The results show that runtime increases with the number of simulation steps, as anticipated.

To better understand the scaling behaviour, runtime was normalized by the number of steps. The per-step computational cost remains approximately stable across the tested temporal

horizons, with runtime/T varying within a narrow range of [0.012, 0.016] seconds per step, consistent with the theoretical O(T) scaling assumption.

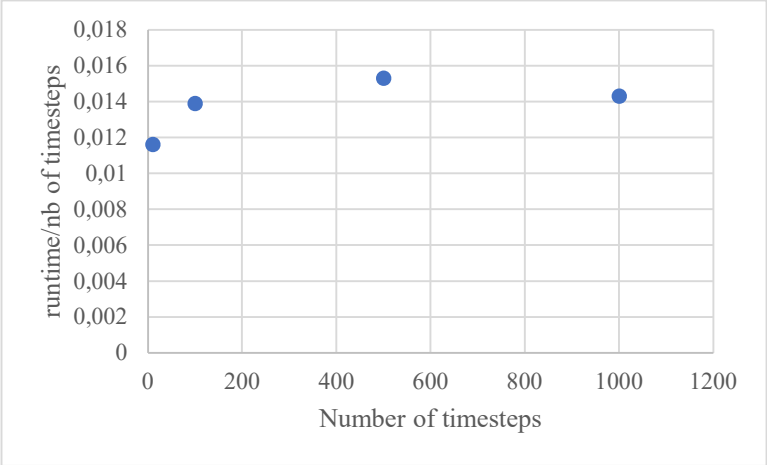


Figure 1: Runtimes relative to number of timesteps

4.2.2 Complexity based on the number of entities, actor relations and transport relations

Experiments were conducted on scenarios A, B, C and D over 20 simulation steps. Each configuration was run 30 times. Outliers were removed using Tukey's fence method. The average, median, IQR, min, and max (expressed in seconds) of the remaining runs are reported.

Table 2: Statistic analysis by scenarios

Scenario	A	B	C	D
Average	0,1210	0,1414	0,2449	0,5727
Median	0,1270	0,1370	0,2410	0,5725
IQR	0,0285	0,0260	0,0250	0,0123
Standard Deviation	0,0216	0,0165	0,0156	0,0087
Min	0,0790	0,1180	0,2190	0,5540
Max	0,1510	0,1730	0,2760	0,5860

The results show that the total runtime increases with the number of entities and relations in the network. For instance, the scenario A executes in approximately 0.12 seconds, while the scenario D requires approximately 0.57 seconds. This increase is expected since larger networks require processing a greater number of actors, flows, and inventory updates at each simulation step.

Relative to complexity, results reveal two systematic departures from the hypothesized worst-case bound.

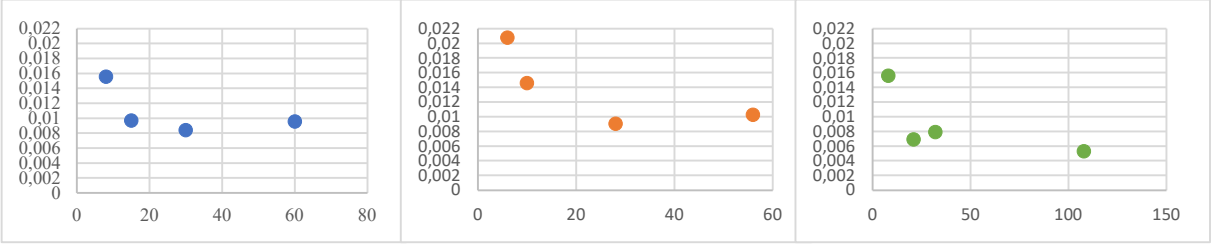


Figure 2: Runtimes relative to E, AR and TR (left to right)

Regarding the number of entities, normalized runtime per entity decreases from approximately 0.016 s/entity at E=8 to 0.009 s/entity at E=30 and E=60, indicating sub-linear growth rather than the linear scaling suggested by the theoretical bound. This reflects the fact that larger networks distribute fixed per-step overhead across more entities, reducing the apparent per-entity cost. In practice, E remains a moderate scaling factor with no evidence of quadratic behavior.

Regarding actor relations, normalized runtime per actor relation decreases from approximately 0.021 s/AR at AR=6 to 0.009 s/AR at AR=28, with a marginal uptick to 0.010 at AR=56. This overall decreasing trend departs from the theoretical expectation of a full multiplicative role for AR. The explanation lies in the engine's architectural design: demand reception, production scheduling, and procurement planning all perform entity-filtered lookups over actor relations rather than full scans of the global relation set. Each entity only processes the subset of actor relations in which it is directly involved, meaning additional relations not activated by a given entity incur negligible cost. AR therefore behaves as a sparse indexed parameter rather than a true multiplicative factor.

Regarding transport relations, normalized runtime per transport relation similarly decreases from approximately 0.016 s/TR at TR=8 to 0.005 s/TR at TR=108. This sub-linear behavior is architecturally guaranteed: inspection of the engine's delivery logic reveals that itinerary resolution is performed via a provider and customer-filtered lookup, meaning each entity only queries the transport relations directly relevant to its outbound shipments. Relations not activated by a given entity at a given step are never traversed, removing TR from the dominant inner loop by the same mechanism as AR.

Based on all four empirical findings, with **T** confirmed linear, AR and TR dropped out of the equation, the practical complexity becomes O(T*E*P). P (number of distinct products) remains as a theoretical inner-loop factor but was not independently varied in your experiments, so it is retained by construction rather than confirmed empirically. However, this complexity estimate is based on code analysis and a limited set of data points and a more thorough empirical validation would be enriched by additional configurations.

Overall, the results indicate that the proposed simulation engine executes supply chain scenarios **in well under 1 second** for networks up to 60 entities, 56 actor relations and 108 transport relations for 20 timesteps. For middle sized networks like scenario C, its performance is under 15 seconds for 1000 timesteps. These runtimes demonstrate that the simulator can efficiently handle moderately sized supply chain networks while maintaining short execution times, enabling **rapid scenario exploration and interactive decision-support analyses**.

4.3 Benchmark

Table 4 presents a qualitative comparison of our simulation model as well as others based on 5 features: configurability, PI alignment, E2E Supply Chain dynamics, reusability, and Open logistics network support.

Table 4: Qualitative comparison of supply chain simulation approaches

Feature	Proposed Engine	AnyLogic	SimPy	PI-specific simulators
Configurable meta-model	Native	Manual coding	Manual coding	Case-specific
PI alignment	Explicit	Possible but not native	Possible but not native	Yes
End-to-end SC dynamics	Yes	Yes	Partial	Transport-focused

Reusable across configurations without structural recoding	Yes	No	No	No
Open logistics network support	Yes[1]	Possible	Possible	Yes

5 Conclusion and future works

This paper presented a generic, meta-model-driven discrete-time simulation engine for configurable supply chain structures, decoupling structural definitions from execution logic to enable reusable, topology-agnostic modeling. All actor types are unified under a single abstract interface mapped to current PI node functions (π-store, π-composer, π-sorter, π-transit, and π-gateway) with delivery objects modeled as π-containers and transport itineraries as π-networks, enabling structural alignment with the Physical Internet paradigm. The engine natively supports multi-owner, open-topology configurations where transport resources are shared across actors and network structure is fully configurable. Experiments on networks of 8 to 60 entities confirmed execution times under 1.2 seconds over 20 steps, with runtime scaling linearly with the simulation horizon. Complexity analysis suggests an O(T*E*P) behavior driven by the forecast propagation step.

The proposed model abstracts several operational aspects (e.g., detailed production constraints and dynamic routing policies), which may limit its applicability to highly specialized industrial contexts.

Several directions are identified for future work. A rigorous empirical complexity analysis across additional network configurations would consolidate the estimates presented here. The transport selection mechanism could be extended to support dynamic open-access routing, consistent with full PI protocols for hyperconnected logistics. Integration with optimization solvers would further enable prescriptive decision-support.

Finally, a dedicated modelling sub-layer is envisioned, comprising a generic network editor enabling practitioners to interactively configure supply chain topologies, actor roles, and relational structures without requiring direct manipulation of the underlying configuration schemas. Building upon this, a performance projection tool leveraging the simulator's time-stepped state outputs would translate simulation results into actionable performance indicators, supporting predictive analysis and decision-making at both operational and strategic levels. Together, these two components would constitute a complete end-to-end environment, from network design to performance exploitation, significantly broadening the engine's applicability to practitioners and researchers alike.

[1] open network topologies can be configured but dynamic open-access routing is not natively enforced

References

- Bottani, E., & Casella, G. (2024). Discrete-event simulation in logistics and supply chain management: A scientometric perspective. *Production & Manufacturing Research*, *12*(1), 2415038. https://doi.org/10.1080/21693277.2024.2415038
- Cimino, A., Longo, F., & Mirabelli, G. (2010). *A General Simulation Framework for Supply Chain Modeling: State of the Art and Case Study*. *7*(2).
- Hakimi, D., Montreuil, B., Sarraj, R., Ballot, E., & Pan, S. (2012). *Simulating a physical internet enabled mobility web: The case of mass distribution in France*.
- Hernández, J. M., & Pedroza, C. (2016). *The influence of the network topology on the agility of a supply chain* (arXiv:1611.10094). arXiv. https://doi.org/10.48550/arXiv.1611.10094
- Kaboudvand, S., & Montreuil, B. (2024). Simulation-Based Assessment of Hyperconnected Megacity Parcel Logistics. *Logistics*, *8*, 66. https://doi.org/10.3390/logistics8030066
- Korder, B., Maheut, J., & Konle, M. (2024). Simulation Methods and Digital Strategies for Supply Chains Facing Disruptions: Insights from a Systematic Literature Review. *Sustainability*, *16*(14), 5957. https://doi.org/10.3390/su16145957
- Macal, C. M., & North, M. J. (2010). Tutorial on agent-based modelling and simulation. *Journal of Simulation*, *4*(3), 151–162. https://doi.org/10.1057/jos.2010.3
- Montreuil, B. (2011). Toward a Physical Internet: Meeting the global logistics sustainability grand challenge. *Logistics Research*, *3*, 71–87. https://doi.org/10.1007/s12159-011-0045-x
- Montreuil, B., Meller, R., & Ballot, E. (2010). Towards a Physical Internet: The Impact on Logistics Facilities and Material Handling Systems Design and Innovation. *Progress in Material Handling Research: 2010*. https://digitalcommons.georgiasouthern.edu/pmhr_2010/40
- Pegden, C. (2007). *SIMIO: A new simulation system based on intelligent objects.* (p. 2300). https://doi.org/10.1109/WSC.2007.4419867
- Pinho, T. M., Coelho, J. P., & Boaventura-Cunha, J. (2016). Forest-based supply chain modelling using the SimPy simulation framework. *IFAC-PapersOnLine, 7th IFAC Conference on Management and Control of Production and Logistics MCPL 2016*, *49*(2), 90–95. https://doi.org/10.1016/j.ifacol.2016.03.016
- Sarraj, R., Ballot, E., Pan, S., Hakimi, D., & Montreuil, B. (2014). Interconnected logistic networks and protocols: Simulation-based efficiency assessment. *International Journal of Production Research*, *52*(11), 3185–3208. https://doi.org/10.1080/00207543.2013.865853

IPIC 2026
12th International
Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE

A Conceptual Digital Twin Architecture for Implementing Physical Internet Logistics Networks

Masoud Kahalimoghadam[1*], Russell Thompson[1], Lele Zhang[2], Michael Kirley[3]

1. University of Melbourne, Department of Infrastructure Engineering, Melbourne, Australia
2. School of Mathematics and Statistics, The University of Melbourne, Australia
3. School of Computing and Information Systems, The University of Melbourne, Melbourne, Australia

Corresponding author: masoud.kahalimoghadam@unimelb.edu.au

Abstract

The Physical Internet (PI) has emerged as a transformative paradigm for freight systems, seeking to replace traditional logistics systems with shared, standardized, and interoperable networks. However, despite growing interest, progress toward PI implementation remains limited, as many studies and projects address only isolated aspects of the PI. This paper identifies the key requirements for implementing a mature PI-enabled logistics network by reviewing and synthesizing the relevant literature. Standardization, interconnectivity, real-time visibility, and collaborative logistics are identified as the main features that must be jointly established to enable integrated PI implementation. The analysis further reveals that these requirements can be met through digital twin technology, which provides a data-driven virtual representation of the physical system, supporting real-time monitoring, modeling, and informed decision-making. Accordingly, this paper not only examines the relationship between key PI features and digital twin capabilities, but also proposes a digital twin framework for PI implementation. Hence, this paper provides an implementation-oriented foundation that clarifies the role of digital twins in enabling PI logistics networks.

Keywords: Physical Internet; Digital twin; Freight systems; Collaborative logistics.

1. Introduction

Logistics networks face various challenges, such as rapid urbanization, the boom in e-commerce, and disruptive events, which threaten network sustainability and efficiency (Kahalimoghadam et al., 2024). The Physical Internet (PI) has been introduced as a transformative paradigm to address these challenges by applying the principles of the digital internet to logistics networks (Montreuil et al., 2012). Accordingly, the PI aims to define interfaces and protocols that enable a shared global logistics network, facilitating interoperable freight movement across multiple stakeholders and logistics assets. Additionally, to ensure seamless freight movement, the PI introduces encapsulated freight flows using standardized PI-containers. However, implementing a PI-enabled network requires interoperable information exchange across PI components and stakeholders to support coordinated and responsive operations.

The PI has been studied for more than a decade worldwide, and PI-based logistics networks have evolved during this period. Over time, this evolution has shifted the PI from conceptual foundations, modular standardization, and the design of core elements such as PI-containers, PI-nodes, and PI-movers (Crainic & Montreuil, 2016; Montreuil et al., 2010) to implementation-oriented topics,

including digital twins, real-time coordination, and synchronized PI hubs (Leung et al., 2022; Wu et al., 2025). This shift suggests that technological adoption is increasingly central to PI implementation. However, existing studies still lack an integrated analysis of the PI implementation features that must be addressed and the technological capabilities that may support them.

A digital twin is a virtual representation of a real-world entity or process, synchronized at a specified frequency and fidelity with its physical counterpart, enabling real-time monitoring, simulation, predictive analysis, and operational support (Thompson et al., 2025). In logistics networks, a digital twin can support system-wide visibility, real-time monitoring, integrated modeling, and decision support, thereby helping improve operational responsiveness and coordination. Building on this synergy, digital twin technology can potentially support PI implementation. For example, synchronization is a key issue in a PI network, especially in intermodal and multimodal systems, because delays or mismatches between transport modes can disrupt transfers, increase waiting times, and reduce overall system efficiency. A digital twin can facilitate synchronization between physical PI-containers and their virtual counterparts by continuously integrating real-time location, status, and schedule data, thereby enabling predictive analytics and timely monitoring to reduce delay risks. Hence, the question is not whether a digital twin supports the PI in general, but which digital twin capabilities are most relevant to the specific features required for PI implementation.

Despite this strong connection between the PI and digital twin technology, recent research reveals that only 2.82% of developed digital twin models are related to transportation and logistics (Tao et al., 2023). This limited uptake highlights the need for a clearer implementation-oriented understanding of which PI features must be addressed to develop a mature PI network and how digital twin capabilities may support them. To address this gap, this paper first identifies the key requirements for PI implementation and then examines which digital twin capabilities are most relevant to these requirements. Based on this analysis, this study proposes a conceptual digital twin architecture for PI network implementation.

The rest of this paper is organized as follows. Section 2 identifies the key requirements for PI implementation. Section 3 examines the digital twin capabilities relevant to these requirements. Section 4 presents the proposed digital twin framework for PI network implementation. Finally, Section 5 concludes the paper and outlines directions for future research.

2. Key Requirements for Physical Internet Implementation

This section examines the key PI features that must be translated into operational practices to support implementation in logistics networks.

2.1 Standardization and Modularity

Standardization and modularity are fundamental to PI implementation because freight must move through the network in standardized, compatible, and easily handled units (Prakash & Vats, 2025). This is mainly achieved through PI-containers designed with standardized dimensions, modular structures, and interlocking capabilities (Tordecilla et al., 2025). Such designs include hierarchical container systems and modular PI-containers, including T, H, and P containers (Leung et al., 2022). They also include reconfigurable container compositions that allow containers to be assembled and disassembled as needed (Wu et al., 2025). These features enhance compatibility across different cargo

types and logistics assets while improving storage, handling, and transport efficiency. They also reduce wasted space, facilitate consolidation, and reduce empty vehicle movements. Therefore, standardization and modularity provide the physical and operational basis for seamless freight movement and more efficient network-wide operations in PI-enabled logistics systems.

2.2 Interoperability and Interconnectivity

Interoperability and interconnectivity are essential to PI implementation because freight and logistics assets must move across a shared logistics network (Nguyen et al., 2022; Prakash & Vats, 2025). At the physical level, this is enabled by interconnected modular containers and standardized logistics entities, including PI-nodes, PI-movers, and PI-hubs, which support coordinated handling and movement across the network (Leung et al., 2022; Tordecilla et al., 2025). At the network level, PI-hubs, which are open and shared among different logistics firms, facilitate cross-docking, consolidation, storage, transfer, and rerouting within an open logistics web (Tordecilla et al., 2025; Wu et al., 2025). In addition, PI-movers connect PI-nodes, such as PI-hubs, ports, and factories, thereby supporting multimodal and intermodal freight movement across the logistics network (Prakash & Vats, 2025; Wu et al., 2025). Physical interoperability and interconnectivity therefore provide the structural foundation for more resilient and efficient freight movement in PI-enabled logistics networks. However, without information sharing about assets and freight among stakeholders, this physical connectivity is difficult to achieve.

2.3 Real-Time Visibility and Information Sharing

Real-time visibility and data exchange provide the informational basis for coordinated freight movement across an interconnected PI network. At the operational level, PI-containers must carry and communicate key shipment information so that logistics actors can coordinate handling, routing, and delivery activities across the network. This requires technologies that enable each container to be uniquely identified, tracked, monitored, and integrated within the PI network through real-time data collection and communication (Leung et al., 2022; Sallez et al., 2016). In addition, real-time data collection, storage, and exchange across organizations support synchronized information flows, continuous operational monitoring, and coordinated decision-making in areas such as scheduling and transport management (Nguyen et al., 2022; Wu et al., 2025). Finally, secure and transparent data exchange, supported by standardized service protocols and smart contracts, further enhances transparency, security, and operational efficiency across platforms and stakeholders (Prakash & Vats, 2025).

2.4 Resource Sharing and Collaborative Logistics

Resource sharing and collaborative logistics are central to PI implementation. They enable firms to move beyond isolated operations and coordinate assets, services, and decision-making processes across an open and interconnected logistics network. In this context, the PI promotes the integration of logistics resources across stakeholders through shared freight, crowd shipping, consolidation, and collaborative PI-container routing through common hubs, which together improve flexibility, reduce transport costs, and enhance overall network performance (Tordecilla et al., 2025; Wu et al., 2025). This collaborative logic also extends to planning activities, where horizontal collaboration for resource sharing can be combined with vertical collaboration for demand forecasting, inventory positioning, and production decisions to improve responsiveness and strengthen supply chain competitiveness (Nguyen et al., 2022; Prakash & Vats, 2025). Stakeholder-based modeling in

construction logistics further shows that the decisions of governments, firms, and service providers are interdependent, and that policy and financial mechanisms can shape the adoption of new practices (Kahalimoghadam & Thompson, 2026). However, such collaboration cannot be sustained only through physical connectivity and information exchange. It also requires planning and allocation mechanisms that support fair and efficient participation, including resource planning and allocation, inventory management, auctions, and cost-sharing mechanisms among actors (Kahalimoghadam, 2024). By enabling coordinated use of assets, services, and planning mechanisms, resource sharing and collaborative logistics help reduce logistics costs and inefficiencies, while improving flexibility, responsiveness, and overall performance in PI-enabled logistics networks.

Table 1 summarizes the key features required for implementing PI-enabled freight networks, along with their roles and impacts within these networks.

Table 1: Key requirements for PI implementation and their operational contributions.

PI implementation requirement	Role in PI implementation	Key operational elements	Expected contribution
Standardization and modularity	Provides the basis for seamless freight movement.	Standardized PI-containers, modular structures, interlocking units, and reconfigurable container compositions.	Improves compatibility, handling, storage, consolidation, and transport efficiency.
Interoperability and interconnectivity	Enables freight and assets to move across a shared logistics network.	PI-nodes, PI-movers, PI-hubs, cross-docking, transfer, rerouting, and multimodal movement.	Supports coordinated, resilient, and efficient freight movement.
Real-time visibility and information sharing	Provides the informational basis for network coordination.	Container identification, tracking, monitoring, real-time data exchange, service protocols, and smart contracts.	Supports synchronization, monitoring, scheduling, transparency, and decision-making.
Resource sharing and collaborative logistics	Enables firms to coordinate assets, services, and planning.	Shared freight, consolidation, collaborative routing, common hubs, auctions, and cost-sharing mechanisms.	Reduces costs and inefficiencies while improving flexibility, responsiveness, and performance.

3. Digital Twin Capabilities for PI Implementation

A digital twin is an emerging technology that uses data-driven modeling to provide an accurate representation of a physical entity (Tao et al., 2023). It is typically supported by simulation, optimization, and prediction, and, through connections to sensor systems and, more broadly, IoT, it can access live data for modeling and enable continuous monitoring (Thompson et al., 2025). The main applications of digital twin have been in aerospace, manufacturing, city management, healthcare, maritime, and logistics (Semeraro et al., 2021). This section discusses the main digital twin features relevant to the PI. It then introduces a digital twin framework for PI network implementation.

3.1 Data Integration and Synchronization

Data integration and synchronization provide the informational basis for a digital twin by ensuring that the virtual model remains aligned with the current state of the physical system. A digital twin draws on heterogeneous data sources, such as design data, operational data, and sensor data, to build a coherent representation of the physical entity and its conditions over time. In this context, data integration is important because the twin must combine information from multiple domains, rather than rely on a single source, to represent structure, status, and behavior in a consistent manner (Dihan et al., 2024). Synchronization then ensures that these data are updated in near real time, so the virtual model can reflect changes in the physical system as they occur (Pan et al., 2021). This continuous

alignment improves the reliability of simulation, state monitoring, and performance evaluation under changing operating conditions. By maintaining data consistency across the physical and virtual spaces, integration and synchronization support an accurate and up-to-date digital representation throughout the system life cycle.

3.2 Modelling: Simulation, Prediction, and Optimization

Modeling is a foundational element of digital twin systems, as it enables the creation of a high-fidelity virtual representation that mirrors the physical system and supports decision-making. A digital twin is a knowledge model that integrates physical models, data, and analytical methods to interpret reality and guide actions (Thelen et al., 2022). On this basis, modeling provides the structure through which system behavior, properties, and interactions are represented, allowing the twin to simulate, evaluate, and predict system performance. Moreover, digital twins rely on continuously updated, multi-domain, and dynamic models that capture the physical form, system behavior, and operational logic of the asset (Wu et al., 2025). This modeling capability enables both simulations, for analysis and design, and emulation, for real-time synchronization with the physical system, extending across design, production, and service phases. Through this integration, modeling supports continuous performance evaluation and optimization by linking real-time data with system behavior, enabling informed and adaptive decisions in complex and uncertain environments.

3.3 Bidirectional Interaction

Bidirectional interaction characterizes the operational logic of digital twin systems by enabling a closed-loop relationship between the physical and virtual spaces. Rather than functioning as a passive replica, a digital twin both receives real-time data from the physical system and actively returns insights, predictions, and control actions. This two-way exchange allows the virtual model to continuously update its state while influencing physical operations through model-driven recommendations (Tao et al., 2023). As a result, the system can evaluate alternative scenarios, anticipate responses to changes, and adjust operational parameters in a timely manner. Such interaction supports coordinated monitoring and control by informing decisions through both current system conditions and simulated outcomes, thereby improving responsiveness and overall system performance across the life cycle.

3.4 Cross-System Integration

Cross-system integration extends the scope of a digital twin from an isolated asset to a broader system composed of interconnected components, data sources, and functions. A digital twin should support a comprehensive representation of the system by integrating information about structure, behavior, functionality, and communication interfaces, rather than focusing only on a single physical element (Tao et al., 2023; Wu et al., 2025). This broader integration is important because physical systems operate within larger operational environments where multiple subsystems, processes, and actors interact. Accordingly, effective digital twins combine data and models across different domains and life cycle stages to maintain a coherent view of the overall system. Such integration improves the ability to analyze system performance, understand interdependencies, and evaluate the effects of changes across connected elements (Thompson et al., 2025). By supporting a unified representation of the wider operational setting, cross-system integration strengthens the capacity of the digital twin to provide reliable analysis and coordinated decision support.

3.5 Service and Operational Intelligence

In digital twin systems, service and operational intelligence are mainly associated with use and maintenance, where the twin supports monitoring, diagnostics, reconfiguration, and performance improvement after deployment. Based on this, digital twins enable remote commissioning, diagnostics, and performance optimization, while also supporting the virtual monitoring of interconnected systems to validate operating conditions such as energy use and faults (Semeraro et al., 2021). They can further enhance prognostics and health management by improving failure prevention, predicting reliability and remaining useful life, and increasing operational accuracy and efficiency. Beyond maintenance, digital twins can also support product reconfiguration during use and enable personalized functional upgrades based on user data, extending their role from operational supervision to adaptive service delivery.

The digital twin features that can support PI network implementation and their expected impacts are summarized in Table 2.

Table 2: Key digital twin capabilities for PI implementation.

Digital twin capability	Main role in PI implementation	Key operational elements	Expected contribution
Data integration and synchronization	Keeps the virtual model aligned with the physical PI system.	Heterogeneous data sources, data integration, real-time updates, and physical-virtual synchronization.	Supports accurate representation, reliable monitoring, simulation, and performance evaluation.
Modeling, simulation, prediction, and optimization	Represents system behavior and supports decision-making.	Physical models, operational logic, dynamic models, simulation, prediction, and optimization.	Enables performance evaluation, future-state estimation, and adaptive decisions.
Bidirectional interaction	Creates a closed-loop link between physical and virtual spaces.	Real-time data input, insights, predictions, recommendations, and control actions.	Supports timely adjustment, coordinated monitoring, and responsive operations.
Cross-system integration	Extends the twin from isolated assets to connected systems.	Interconnected components, data sources, functions, interfaces, subsystems, and actors.	Improves analysis of interdependencies, system performance, and coordinated decision support.

Based on the findings in Sections 2 and 3, the main features that must be addressed to enable PI network implementation are identified. In addition, the digital twin capabilities that can address these PI requirements are discussed. Figure 1 summarizes the findings from Sections 2 and 3.

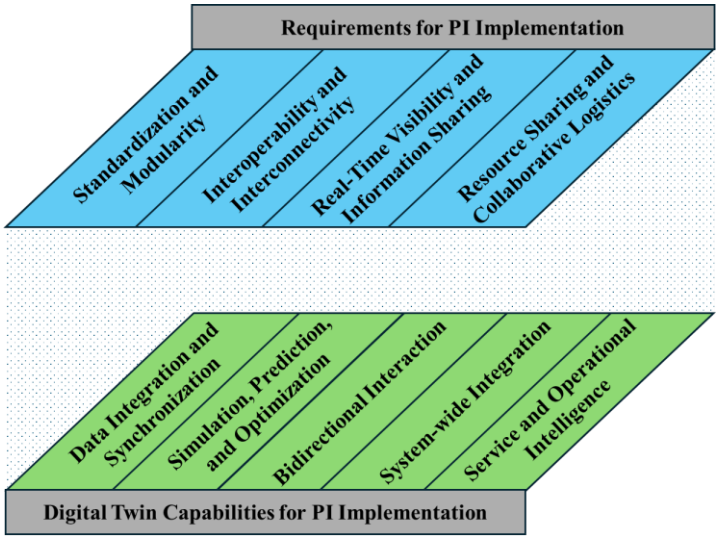


Figure 1: Key PI implementation requirements and the digital twin capabilities that support them.

4. Conceptual Digital Twin Architecture for PI Network Implementation

Building on the PI implementation requirements identified in Section 2 and the digital twin capabilities discussed in Section 3, this section proposes a conceptual digital twin architecture for PI network implementation. Figure 2 illustrates this framework and shows the connection between the physical and digital layers. Each layer consists of several components, which are discussed below.

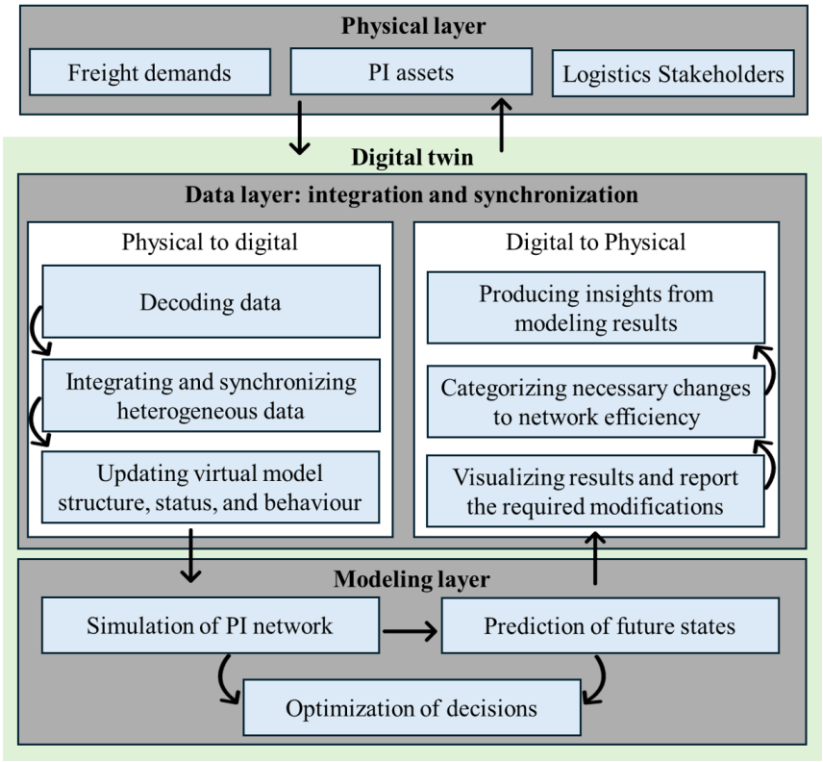


Figure 2: Proposed digital twin architecture for PI network implementation.

The physical layer represents the PI logistics network. One of its main components is freight demand, which is driven by market activity. Current demand directly affects operational decisions, such as PI-mover routing, whereas historical demand data can be used by the prediction component of the modeling module to support strategic decisions, such as PI-hub location planning. Freight demand also plays an important role in stakeholder interactions, particularly financial interactions among logistics actors.

Another major component of the physical layer is PI assets, namely PI-containers, PI-movers, and PI-hubs, which influence decisions from operational to strategic levels. For PI-containers, the main data to be collected include capacity, location, goods type, and origin-destination information. The firms involved in each shipment must also be identified. These may include receivers, carriers, freight forwarders, shippers, and producers. In addition, delivery schedules for PI-containers must be transferred from the physical layer to the digital twin.

These freight flows are moved by PI-movers, which may include any transport mode capable of carrying PI-containers, such as trucks, trains, and ships. Key data for PI-movers include delivery routes, container characteristics, and origin-destination information. PI-movers may also be owned and operated by different stakeholders, such as carriers, shippers, or PI-managers. Further details are provided in Kahalimoghadam et al. (2025). Logistics stakeholders may include state and local

governments, carriers, shippers, freight forwarders, PI-managers, receivers, and transport network managers across road, rail, and shipping systems.

The digital twin consists of two main layers, the data integration layer and the modeling layer. These layers are designed not only to replicate the components of the physical layer, but also to generate actionable insights. Data integration and synchronization operate bidirectionally. On the physical-to-digital side, raw data received from the physical layer must first be decoded. This involves filtering noise, structuring heterogeneous data, and converting the data into a unified format for consistent representation. Because the data originate from multiple sources, integration and standardization are essential. In addition, digital twins can support the development of protocol frameworks, standardized interfaces, and service platforms that define how logistics elements connect and operate within an interoperable PI network. The data must also be synchronized so the virtual model reflects time-aligned system states and can support reliable monitoring and decision-making. In addition, some of these data describe the current status of PI assets and therefore must be updated in the final physical-to-digital component before being passed to the modeling layer.

The second component of the data integration and synchronization layer is the digital-to-physical flow, through which data moves from the modeling layer back to the physical layer. Synchronization across transport modes and stakeholders is essential, as misalignment can lead to delays, inefficiencies, and underutilized resources. At this point, outputs from the modeling layer are visualized, and any changes from the previous system state are identified. These changes are then categorized and passed to the final component to generate insights for the physical layer. The purpose is to guide physical-layer decisions in ways that improve system efficiency. Such decisions may involve optimized routing or, in a collaborative network, determining which firm should fulfill a given demand.

The last layer in the proposed digital twin framework is the modeling layer, which includes simulation, prediction, and optimization. Although some digital twins may rely on only one of these components, the proposed framework integrates all three to support more comprehensive PI network implementation.

In a digital twin, simulation is the component that reproduces the behavior of the physical layer under specific conditions using integrated and time-synchronized data. It allows the modeling layer to represent system operations not only in the current state, but also under alternative scenarios. This supports performance evaluation without disrupting real operations and enables different decisions, configurations, and disruptions to be tested in a risk-free environment. Simulation also provides the basis for prediction by projecting future system states under different inputs, and it supports optimization by helping identify strategies that improve efficiency, reduce costs, and make better use of resources. Through continuous data updates, simulation also helps maintain alignment between the physical and virtual systems so that the resulting insights remain reliable and actionable.

Prediction of future states and performance in a digital twin emerges from its ability to use continuously updated data to infer how the system will evolve over time. While prediction is primarily data-driven, it may require input from simulation because simulated system responses can enrich the data space, especially under conditions that have not yet been observed in reality, thereby improving the robustness and coverage of predictive models. In this context, prediction focuses on estimating the most likely future states based on current trends and learned patterns, enabling early

detection of risks such as failures or performance degradation. This distinguishes it from scenario analysis in simulation, which evaluates multiple hypothetical conditions and alternative decisions to explore a range of possible outcomes rather than identifying the most probable one. As a result, prediction supports proactive and targeted decision-making, while scenario analysis supports the exploratory and comparative assessment of different system configurations.

Optimization in a digital twin aims to identify the best decisions or system configurations to improve performance, e.g., minimizing costs, reducing emissions, or enhancing resource utilization. While optimization determines what should be done, it often requires input from both simulation and prediction, because simulation evaluates how different decisions affect system behavior and prediction provides estimates of future conditions under which those decisions will operate. This integration ensures that optimization is not only based on current system states but also accounts for expected changes and uncertainties. Moreover, optimization differs from prediction by being prescriptive rather than descriptive, as it actively searches for the best course of action rather than estimating what is likely to happen. It also differs from scenario analysis in simulation, as it does not simply compare predefined alternatives, but systematically explores the decision space to identify optimal or near-optimal solutions under given objectives and constraints.

5. Conclusion

PI logistics network implementation cannot be achieved effectively by addressing only selected features. To accelerate implementation, the critical characteristics that enable a PI network to operate as an integrated system must first be identified. These include standardization through modular container design, interconnectivity across logistics assets, real-time visibility through live data sharing, and collaborative logistics through shared resource use. The review further shows that digital twin technology can support these requirements by enabling integrated data exchange, synchronized system representation, continuous monitoring, and decision support across assets and stakeholders. Based on these findings, the proposed framework in Figure 2 illustrates how digital twin can provide the architectural link between the physical layer of the PI network and the digital environment. The framework also identifies the minimum requirements of each digital twin layer and shows how data and information flow across the system. The main contribution of this research is not only to identify the core implementation requirements of the PI, but also to explain how digital twin capabilities can be aligned with them in a coherent, implementation-oriented framework. This provides a conceptual foundation for future research on the design and practical deployment of PI-enabled logistics networks. In addition, future research should extend beyond a sole focus on digital twin technology, as a more rigorous literature review is needed to identify the full range of technologies that may support PI implementation and to clarify how their roles compare with or complement that of digital twins.

References

Crainic, T. G., & Montreuil, B. (2016). Physical Internet Enabled Hyperconnected City Logistics. *Transportation Research Procedia*, *12*, 383-398. https://doi.org/https://doi.org/10.1016/j.trpro.2016.02.074

Dihan, M. S., Akash, A. I., Tasneem, Z., Das, P., Das, S. K., Islam, M. R., Islam, M. M., Badal, F. R., Ali, M. F., Ahamed, M. H., Abhi, S. H., Sarker, S. K., & Hasan, M. M. (2024). Digital twin: Data exploration, architecture, implementation and future. *Heliyon*, *10*(5), e26503. https://doi.org/10.1016/j.heliyon.2024.e26503

Kahalimoghadam, M. (2024). *Developing and Implementing Collaborative Urban Distribution Networks to Improve the Sustainability of Last-mile Logistics* The University of Melbourne]. https://hdl.handle.net/11343/343143

Kahalimoghadam, M., & Thompson, R. G. (2026). Assessing the viability of construction and demolition waste recycling through agent-based modeling. *International Journal of Construction Management*, 1-26. https://doi.org/10.1080/15623599.2026.2653085

Kahalimoghadam, M., Thompson, R. G., & Rajabifard, A. (2024). Assessing unsustainable trends in city logistics. *Transportation Research Procedia*, *79*, 401-408. https://doi.org/https://doi.org/10.1016/j.trpro.2024.03.053

Kahalimoghadam, M., Thompson, R. G., & Rajabifard, A. (2025). An intelligent multi-agent system for last-mile logistics. *Transportation Research Part E: Logistics and Transportation Review*, *200*, 104191. https://doi.org/https://doi.org/10.1016/j.tre.2025.104191

Leung, E. K. H., Lee, C. K. H., & Ouyang, Z. (2022). From traditional warehouses to Physical Internet hubs: A digital twin-based inbound synchronization framework for PI-order management. *International Journal of Production Economics*, *244*, 108353. https://doi.org/10.1016/j.ijpe.2021.108353

Montreuil, B., Meller, R. D., & Ballot, E. (2010). *Towards a Physical Internet: the impact on logistics facilities and material handling systems design and innovation*. Material Handling Industry of America.

Montreuil, B., Meller, R. D., & Ballot, E. (2012). Physical Internet Foundations. *IFAC Proceedings Volumes*, *45*(6), 26-30. https://doi.org/https://doi.org/10.3182/20120523-3-RO-2023.00444

Nguyen, T., Duong, Q. H., Nguyen, T. V., Zhu, Y., & Zhou, L. (2022). Knowledge mapping of digital twin and physical internet in Supply Chain Management: A systematic literature review. *International Journal of Production Economics*, *244*, 108381. https://doi.org/10.1016/j.ijpe.2021.108381

Pan, Y. H., Wu, N. Q., Qu, T., Li, P. Z., Zhang, K., & Guo, H. F. (2021). Digital-twin-driven production logistics synchronization system for vehicle routing problems with pick-up and delivery in industrial park. *International Journal of Computer Integrated Manufacturing*, *34*(7-8), 814-828. https://doi.org/10.1080/0951192x.2020.1829059

Prakash, S., & Vats, P. (2025). Decision-making models for implementing physical internet to build resilience against supply chain and logistics risks. *Logistics Research*, 1-24. https://doi.org/10.1108/lore-10-2024-0003

Sallez, Y., Pan, S., Montreuil, B., Berger, T., & Ballot, E. (2016). On the activeness of intelligent Physical Internet containers. *Computers in Industry*, *81*, 96-104. https://doi.org/10.1016/j.compind.2015.12.006

Semeraro, C., Lezoche, M., Panetto, H., & Dassisti, M. (2021). Digital twin paradigm: A systematic literature review. *Computers in Industry*, *130*, 103469. https://doi.org/10.1016/j.compind.2021.103469

Tao, F., Xiao, B., Qi, Q., Cheng, J., & Ji, P. (2023). Digital twin modeling. *Journal of Manufacturing Systems*, *64*, 372-389, Article 1. https://doi.org/10.1016/j.jmsy.2022.06.015

Thelen, A., Zhang, X., Fink, O., Lu, Y., Ghosh, S., Youn, B. D., Todd, M. D., Mahadevan, S., Hu, C., & Hu, Z. (2022). A comprehensive review of digital twin—part 2: roles of uncertainty quantification and optimization, a battery digital twin, and perspectives. *Structural and Multidisciplinary Optimization*, *66*(1). https://doi.org/10.1007/s00158-022-03410-x

Thompson, R. G., Kahalimoghadam, M., Zhang, L., & Sabri, S. (2025). Digital Twins for Hyperconnected City Logistics. In G. Rabadi & B. Soykan (Eds.), *Optimizing Supply Chains Through Digital Twins: Navigating the Modern Landscape* (pp. 221-240). Springer Nature Switzerland. https://doi.org/10.1007/978-3-032-08147-6_12

Tordecilla, R. D., Montoya-Torres, J. R., & Guerrero, W. J. (2025). Resilient design of hyperconnected multiactor Physical Internet supply chain networks. *International Transactions in Operational Research*, *32*(6), 3528-3564. https://doi.org/10.1111/itor.13615

Wu, H., Li, M., Yu, C., Ouyang, Z., Lai, K.-h., Zhao, Z., Pan, S., Wang, S., Zhong, R. Y., Kuo, Y.-H., Zhang, F., Huang, W., Shen, Z.-J. M., Ballot, E., & Huang, G. Q. (2025). Towards cyber-physical internet: A systematic review, fundamental model and future perspectives. *Transportation Research Part E: Logistics and Transportation Review*, *197*, Article 104051. https://doi.org/10.1016/j.tre.2025.104051

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Interoperable π-Containers: Physical-Internet Network Design and Multi-Firm Adoption under Uncertainty

Jinghan Huang[1], Mingxi Chen[2], Yumang Ye[3], Wenbin Wang[4], Timothy Tian[5], Libin Zhang[6], Zhaohui Lyu[7]

1, 2, 4: Shanghai University of Finance and Economics, Shanghai, China

3, 6: COSCO SHIPPING Logistics Supply Chain, Shanghai, China

5: WaterMirror Technology, Shanghai, China

7: Wenzhou University of Technology, Wenzhou, China

Corresponding author: jinghanhuang@stu.sufe.edu.cn

Keywords: *Physical Internet; π-containers; multi-firm adoption; digital twin; simulation-based optimization*

Conference Domain Fitness: *This contribution advances Physical Internet design and operations for cross-border and multimodal logistics through interoperable, standardized modular π-containers and open, interoperable handling and data interfaces. Grounded in the COSCO SHIPPING context, it quantifies when this PI design creates value beyond generic "shared capacity" by jointly modeling open hub–relay network design, π-container pooling, and a verifiable digital-twin chain-of-custody protocol. Using simulation-based optimization under demand volatility, turnaround uncertainty, lane imbalance, and disruption/recall shocks, it derives implementable access/pricing rules and governance levers that sustain participation, preserve accountability, and guide multi-firm adoption on regulated global lanes.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☒ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☒ *Transportation Equipment,* ☒ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☒ *Horizontal Supply Chain Alignment,* ☒ *Logistics/Commercial Data Platform,* ☒ *Access and Adoption,* ☒ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☐ *In-Person presentation*

Research Contribution Abstract

Cross-border ocean logistics for high-value, regulation-sensitive cargo (e.g., lithium-battery modules and other tightly controlled industrial components) is still hindered by single-use transport packaging, empty repositioning, and costly traceability breakdowns (disputed custody, slow recalls). Most countermeasures are implemented within a single firm's closed loop, which limits reuse turns and leaves imbalanced lanes largely unaddressed. Grounded in the COSCO SHIPPING context, we study when Physical Internet (PI) design—standardized modular containers plus open, interoperable handling and data interfaces—creates incremental value in a multi-firm shipping ecosystem, beyond generic "shared capacity." We model a carrier-operator that commits to three coupled PI choices: (i) an open hub–relay structure that supports cross-firm consolidation and transshipment through standardized handling interfaces, (ii) a π-container pool of reusable,

stackable/collapsible modules that can circulate across firms and lanes, and (iii) a verifiable digital-twin protocol that records chain-of-custody and condition at interoperable nodes. Heterogeneous shippers (volumes, lane portfolios, directional imbalances, compliance sensitivity) endogenously choose whether to adopt π-containers and the protocol and how to route shipments through the network. Total cost captures (a) utilization and handling/dwell time improvements enabled by modular standardization and interface compatibility, (b) packaging lifecycle cost driven by pool turns and turnaround times, (c) empty-return costs driven by cross-firm backhaul matching, and (d) disruption/recall losses reduced by verifiable custody and responsibility assignment. We solve the joint design–adoption problem using simulation-based optimization: the optimization layer selects network nodes, pool sizing, and access/pricing rules; the simulation layer evaluates performance under demand volatility, turnaround uncertainty, lane imbalance, and disruption/recall shocks. The analysis delivers three findings. First, interoperability generates super additive savings only when lane complementarity/imbalance is sufficiently strong; in that regime, open-node transshipment and fungible π-containers jointly reduce both handling frictions and empty repositioning, whereas closed loops plateau quickly. Second, openness introduces free-riding and data-boundary concerns; an implementable scheme that couples access pricing with turn-based circulation credits linked to verified protocol compliance sustains participation and preserves accountability. Third, uncertainty shifts optimal design toward fewer but more strategically placed open relays and larger, more flexible pool capacity to buffer imbalances and shocks. Overall, the paper provides a tractable, PI-specific account of how standardization plus openness changes pooling economics, and yields concrete design and governance levers for carriers and shippers operating regulated global lanes.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Place your logo here

Multi-Agent Reinforcement Learning Framework with Resource-Affinity Modeling for Production-Logistics Collaborative Scheduling in MiC Factories

Yujie Han[1] and Zhiheng Zhao[1,2], and George Q. Huang[1,2]
1.Department of Industrial and Systems Engineering, The Hong Kong Polytechnic University, Hong Kong, China
2.Research Institute for Advanced Manufacturing, The Hong Kong Polytechnic University, Hong Kong, China
Corresponding author: zhiheng.zhao@polyu.edu.hk

Abstract: *Driven by the Internet of Things (IoT) and Physical Internet (PI), MiC factory production is evolving toward a more interconnected and data-driven manufacturing mode, where machines and transport units operate as tightly coupled cyber-physical resources. Efficient module production therefore depends on real-time collaboration between production and logistics resources. However, the resulting production-logistics collaborative scheduling problem is highly challenging due to heterogeneous resources, dynamic resource states, and strong interdependence between machine assignment and transport-unit dispatch. Although recent deep reinforcement learning methods have shown promise for dynamic scheduling, most existing approaches still rely on coarse state representations and cannot explicitly capture the evolving relationships among heterogeneous resources. To address this gap, this paper proposes a two-stage resource-affinity multi-agent reinforcement learning framework for IoT-enabled MiC factory scheduling. A production-state embedding network first captures the affinity between operations and machines, and a logistics-state embedding network then models the dynamic competition relationships among transport units under the updated production context. The learned embeddings are integrated into an actor-critic architecture and optimized by proximal policy optimization. Experimental results show that the proposed method consistently outperforms heuristic rules, an ablation variant without two-stage state encoding, and the advanced DANIEL baseline. The results verify that explicit modeling of dynamic heterogeneous resource relationships can effectively improve collaborative scheduling performance in interconnected MiC production systems.*

Keywords: *Internet of Things; Physical Internet; modular integrated construction; production-logistics collaborative scheduling; multi-agent reinforcement learning.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan: ☐ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories), ☒ Transportation Equipment, ☐ PI Networks, ☒ System of Logistics Networks, ☐ Vertical Supply Consolidation, ☐ Horizontal Supply Chain Alignment, ☒ Logistics/Commercial Data Platform, ☐ Access and Adoption, ☐ Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

1 Introduction

In modular integrated construction (MiC), factory production plays a decisive role in the efficiency of downstream logistics and on-site installation. Unlike conventional construction, MiC shifts major construction activities into factory environments, where each module undergoes multiple processing and transportation stages before delivery. During this process,

operations are assigned to eligible machines, while semi-finished modules are transferred by logistics resources. As a result, factory production in MiC inherently depends on the close coordination of production and logistics resources, which has been increasingly recognised in recent MiC and DfMA studies (Bao et al., 2022; Pan et al., 2023).

This problem is commonly formulated as the Production-Logistics Collaborative Scheduling Problem (PLCSP), which aims to jointly optimize production and transportation decisions under process and logistics constraints. However, PLCSP in MiC factories is highly challenging. First, the system involves complex heterogeneous resources, including operations, machines, and transport units, whose functions, states, and interactions are fundamentally different. Second, the relationships among these resources are highly dynamic, varying with processing progress, machine occupancy, and transport availability. Third, production and logistics decisions are strongly coupled: production assignments determine subsequent transport demands, while logistics delays directly affect machine start times and workflow continuity. Similar coordination difficulties have also been noted in MiC transportation and synchronized execution studies (Wang et al., 2024; Jiang et al., 2024). These characteristics make it difficult to represent the scheduling state effectively and to learn high-quality coordination policies.

To address such complexity, deep reinforcement learning (DRL), especially multi-agent reinforcement learning (MARL), has emerged as a promising approach for collaborative scheduling in dynamic environments. By decomposing decisions into multiple agents, MARL improves scalability and reduces the difficulty of large action spaces. Recent progress in flexible job shop scheduling has shown the potential of graph-based and attention-based DRL for complex scheduling state modeling (Zhang et al., 2023; Wan et al., 2024). However, existing MARL-based methods usually rely on manually designed state features or simple feature concatenation, and thus cannot explicitly capture the dynamic relationships among heterogeneous resources. As a result, the learned policies often lack sufficient awareness of cross-resource dependencies, limiting coordination performance in highly coupled production-logistics systems.

Motivated by this gap, this paper proposes a two-stage resource-affinity modeling framework for MARL-based production-logistics collaborative scheduling in MiC factories. In the first stage, a production-state embedding network models the affinity between operations and machines to capture their interactive dependencies. In the second stage, a logistics-state embedding network further characterizes the dynamic relationships among transport resources under the updated production context. By explicitly modeling resource affinity in a staged manner, the proposed method provides relation-aware state representations that improve collaborative policy learning for heterogeneous scheduling environments. The main contributions of this paper are shown below.

- We develop a two-stage resource-affinity modeling method to explicitly encode the dynamic relationships among heterogeneous resources.
- We integrate the proposed relation-aware state representation into a MARL framework, enabling effective collaborative scheduling in dynamic MiC factory environments.

2 Literature review

2.1 MiC Production in the Factory

Research on MiC production initially focused on the broader industrialisation benefits and implementation conditions of modular construction. Bao et al. (2022) examined design for manufacture and assembly (DfMA) enablers for offsite interior design and construction,

showing that early design standardisation, process simplification, and supply-chain coordination are critical to improving offsite production. Pan et al. (2023) further discussed MiC implementation in high-rise and high-density cities and highlighted that logistics constraints, limited capable suppliers, and regulatory restrictions remain major barriers in Hong Kong. These studies established the managerial and organisational foundations of MiC, but they mainly viewed factory production from a strategic or adoption-oriented perspective rather than from the viewpoint of real-time workshop scheduling. Recent studies have gradually moved toward factory-level operational issues. Yang and Lu (2023) investigated facility layout design for modular construction manufacturing and showed that factory configuration has a significant impact on throughput and cost efficiency. Peiris et al. (2023) then demonstrated that digitalised manufacturing execution and control systems can improve information visibility and offsite manufacturing productivity. Along the logistics dimension, Wang et al. (2024) formulated a mathematical programming model for MiC transportation scheduling, while Jiang et al. (2024) proposed a data-driven out-of-order model for synchronized planning, scheduling, and execution in modular construction fit-out management. Along this timeline, the literature has progressed from general MiC adoption and process design toward more operational and data-driven decision support.

2.2 Production Scheduling Solutions

In the scheduling literature, early studies on flexible job shop scheduling mainly relied on mathematical programming and metaheuristics. Ozguven et al. (2010) developed mathematical models for job-shop scheduling with routing and process-plan flexibility, and Demir and Isleyen (2013) systematically evaluated representative mathematical formulations for the flexible job shop scheduling problem (FJSP). To address dynamic environments, Shen and Yao (2015) proposed a multi-objective evolutionary approach that jointly considered efficiency and stability under disturbances. As summarised by Dauzere-Perez et al. (2024), these model-based and heuristic approaches laid the methodological foundation of FJSP research, but they often face scalability limitations in large and highly dynamic scheduling environments. With the rise of data-driven scheduling, recent studies have increasingly adopted deep reinforcement learning (DRL). Gui et al. (2023) applied DRL to dynamic flexible job shop scheduling, showing the advantage of adaptive rule selection over static dispatching policies. When transportation constraints were considered explicitly, Zhang et al. (2023) proposed a graph-neural-network-based DRL framework for dynamic flexible job shop scheduling with insufficient transportation resources. Wan et al. (2024) further introduced a meta-path-based heterogeneous graph neural network to improve FJSP state representation, and Dong et al. (2025) developed a heuristic-assisted DRL algorithm for flexible job shop scheduling with transport constraints. Overall, the literature has evolved from exact and heuristic optimisation to graph-based and DRL-based scheduling.

2.3 Research Gaps and Motivations

Two gaps can be identified from the above review. First, MiC studies have increasingly recognised the importance of factory operations, digital control, and transportation planning (Bao et al., 2022; Pan et al., 2023; Peiris et al., 2023; Wang et al., 2024; Jiang et al., 2024), but they have not yet sufficiently addressed the integrated scheduling of processing and intra-factory logistics resources at the workshop level. Second, although recent scheduling studies have introduced DRL and graph-based learning to model dynamic shop-floor states (Gui et al., 2023; Zhang et al., 2023; Wan et al., 2024; Dong et al., 2025), most of them are designed for generic FJSP settings and do not explicitly model the staged and heterogeneous dependency between operations, machines, and transport units. In MiC factories, however, these resource

relationships are dynamic and strongly interdependent. This motivates the present study to develop a two-stage resource-affinity modeling framework that explicitly captures such dynamic heterogeneous interactions and supports more effective multi-agent collaborative scheduling.

3 Problem Definition

The considered problem is a production-logistics collaborative scheduling problem in a MiC factory. Let $J = \{J_1, \dots, J_n\}$ denote the set of modules, where each module J_i consists of an ordered set of operations $O_i = \{O_{i1}, \dots, O_{in_i}\}$. Each operation O_{ij} can be processed by one eligible machine from the candidate set M_{ij}, where M_{ij} is a subset of the machine set $M = \{M_1, \dots, M_m\}$. The inter-machine transportation of semi-finished modules is executed by one logistics unit from $V = \{V_1, \dots, V_q\}$. Therefore, each scheduling decision jointly determines machine assignment and logistics-unit dispatch under process, resource, and transportation constraints. Let $x_{ijk} \in \{0,1\}$ denote whether operation O_{ij} is assigned to machine M_k, and let $y_{ijr} \in \{0,1\}$ denote whether the transportation required before processing O_{ij} is served by logistics unit V_r. Let S_{ij} and C_{ij} denote the start and completion times of O_{ij}, respectively, and let p_{ijk} be the processing time of O_{ij} on M_k. The transportation time associated with assigning V_r is denoted by τ_{ijr}, which depends on the current position of V_r and the destination machine selected for O_{ij}. The objective is to minimize the final makespan $C_{\max}$ of all modules:

$$\min \quad C_{\max}$$

$$\text{s.t.} C_{\max} \geq C_{ij}, \text{for all } i, j, \tag{1}$$

$$\sum_{M_k \in M_{ij}} x_{ijk} = 1, \text{for all } i, j, \tag{2}$$

$$\sum_{V_r \in V} y_{ijr} = 1, \text{for all } i, j, \tag{3}$$

$$C_{ij} = S_{ij} + \sum_{M_k \in M_{ij}} p_{ijk}\, x_{ijk}, \tag{4}$$

$$S_{ij} \geq C_{i,j-1} + \sum_{V_r \in V} \tau_{ijr}\, y_{ijr}, \text{ for } j > 1, \tag{5}$$

Equation (1) defines the makespan. Equation (2) and (3) ensure unique machine assignment and logistics dispatch. Equation (4) calculates the completion time after machine assignment. Equation (5) links the start of an operation to the completion of its predecessor and the required transportation process. In addition, no machine can process more than one operation at a time, and no logistics unit can execute overlapping transportation tasks.

4 Methodology

This paper proposes a two-stage resource-affinity modeling framework embedded in a multi-agent reinforcement learning architecture, as shown in Figure 1. The framework contains a production agent and a logistics agent. At each decision step, the production agent first selects a feasible operation-machine pair based on production-state embedding. The logistics state is then updated according to the selected production action, after which the logistics agent

dispatches a transport unit. In this way, the framework captures the sequential coupling between production allocation and logistics response.

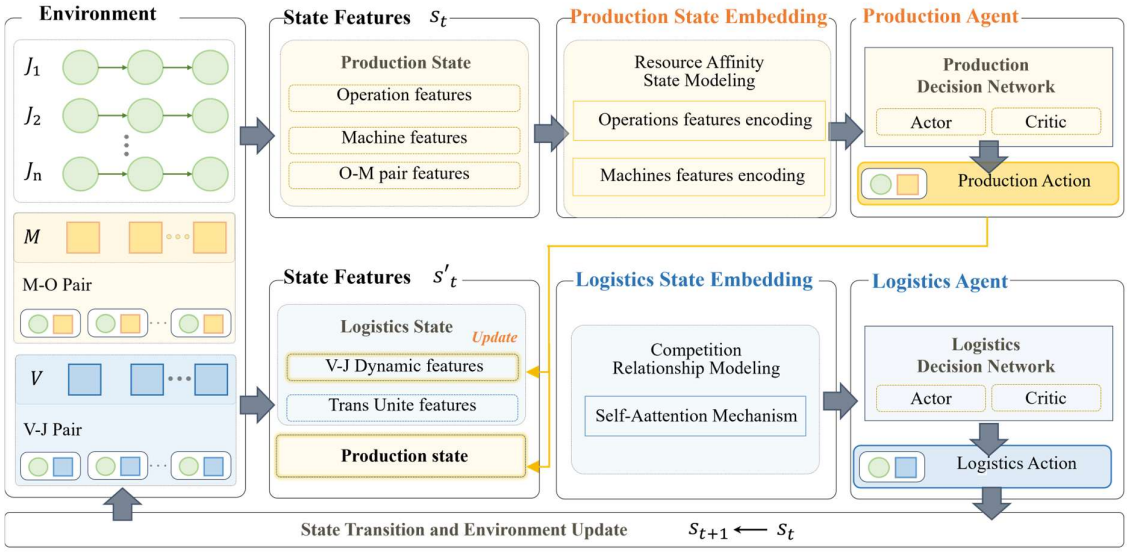


Figure 1: Overall Architecture of the Proposed Two-Stage Resource-Affinity MARL Framework

4.1 MDP Formulation

The collaborative scheduling process is modeled as a sequential Markov decision process $M = \langle S, A, P, R, \pi \rangle$. The global state at decision step t is represented as $s_t = \{s_t^p, s_t^l\}$, where s_t^p contains production information such as operation features, machine features, and operation-machine pair features, and s_t^l contains logistics-unit features, vehicle-job dynamic features, and the updated production context. The action is decomposed into two parts, $a_t = \{a_t^p, a_t^l\}$.
The production action a_t^p selects a feasible operation-machine pair, and the logistics action a_t^l selects the transport unit for the corresponding transportation task. After the two actions are executed, the environment updates operation status, machine occupancy, transport-unit location, and the set of feasible candidates, yielding the next state s_{t+1} according to the transition function $P(s_{t+1} \mid s_t, a_t)$. To optimize the scheduling objective, a shared reward is defined by the reduction in estimated makespan between two consecutive decisions: $r_t = C_{\max}(s_t) - C_{\max}(s_{t+1})$. This reward encourages both agents to choose actions that shorten the final completion time. The collaborative policy is factorized as $\pi(a_t \mid s_t) = \pi_p(a_t^p \mid h_t^p)\, \pi_l(a_t^l \mid h_t^l, a_t^p)$, where h_t^p and h_t^l denote the production and logistics embeddings generated by the two-stage state modeling network.

4.2 Production Resource State Modeling

As shown in Figure 2, the first stage models the affinity between operations and machines. Let $H_o \in \mathbb{R}^{N_t \times d_o}$ and $H_m \in \mathbb{R}^{m \times d_m}$ denote the raw operation and machine feature matrices at step t, where N_t is the number of unscheduled operations. To capture cross-resource dependency, a bidirectional cross-attention mechanism is used to mutually enhance operation and machine representations, as shown below.

$$Q_o = H_o W_o^Q, K_m = H_m W_m^K, V_m = H_m W_m^V, \tag{6}$$

$$H_o^+ = \text{Softmax}\left(\frac{Q_o K_m^T}{\sqrt{d}}\right) V_m + H_o, \tag{7}$$

$$Q_m = H_m W_m^Q, K_o = H_o W_o^K, V_o = H_o W_o^V, \tag{8}$$

$$H_m^+ = \text{Softmax}\left(\frac{Q_m K_o^T}{\sqrt{d}}\right) V_o + H_m. \tag{9}$$

Through this design, each operation representation is aware of machine availability and processing capability, while each machine representation is aware of the workload and priority of candidate operations. Then, the feasible operation-machine pair features are concatenated with the enhanced node features to form the production embedding h_t^p. This stage provides the production agent with a relation-aware representation instead of isolated handcrafted features.

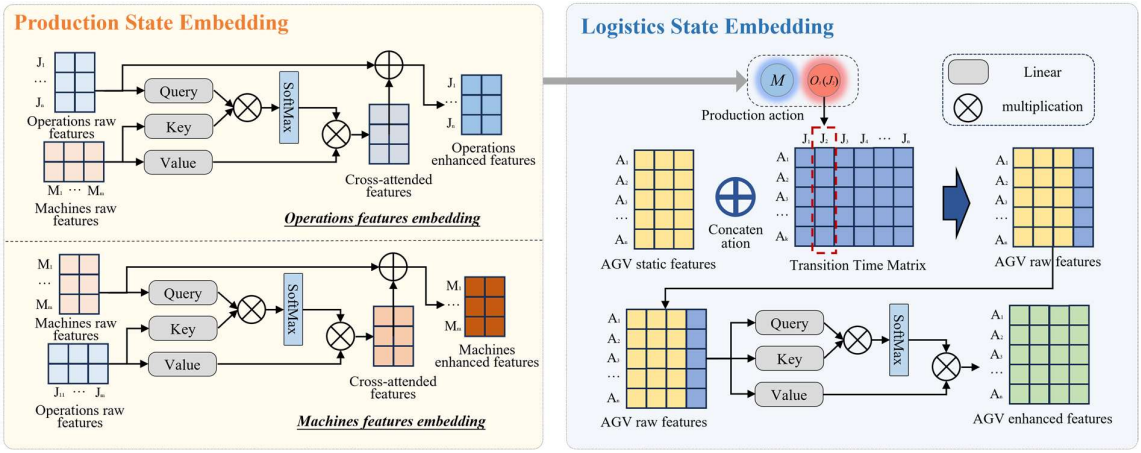


Figure 2: Production-State Embedding with Resource Affinity Modeling

4.3 Logistics Resource State Modeling

As shown in Figure 2, the second stage captures the dynamic relationships among logistics resources under the updated production context. After the production agent selects an operation-machine pair, the transportation demand and destination machine become explicit. Based on this information, a transition-time matrix T_t is constructed to describe the travel cost from each transport unit to the newly generated service target. The logistics raw feature matrix is formed as $H_v^0 = [H_v \parallel \Phi(T_t, a_t^p)]$, where H_v denotes the static and dynamic attributes of transport units, and $\Phi(\cdot)$ maps the production action and transition-time information into vehicle-job pair features. To model resource competition, self-attention is applied over all transport units:

$$Q_v = H_v^0 W_v^Q, K_v = H_v^0 W_v^K, V_v = H_v^0 W_v^V, \tag{10}$$

$$H_v^+ = \text{Softmax}\left(\frac{Q_v K_v^T}{\sqrt{d}}\right) V_v + H_v^0. \tag{11}$$

The resulting logistics embedding h_t^l encodes not only the individual state of each transport unit, but also the competitive relationships among them. This is important because dispatching one unit immediately changes the future availability and relative advantage of others.

4.4 Policy Optimization

The production and logistics embeddings are respectively fed into two actor-critic decision networks. The production actor outputs a probability distribution over feasible operation-machine pairs, and the logistics actor outputs a probability distribution over feasible transport units. The two critics estimate the state values for policy evaluation. The whole framework is

trained by PPO to ensure stable policy updates in the multi-agent scheduling process. The final objective combines the clipped policy loss, value loss, and entropy regularization. Through iterative interaction with the scheduling environment, the model learns to jointly optimize machine assignment and logistics dispatch.

5 Experiments

5.1 Experimental Setup

Training and Validation Dataset. The proposed model is trained and evaluated on simulated production–logistics scheduling instances for modular construction factories. To reflect heterogeneous and dynamic MiC production environments, three problem scales are considered: 10×5×5, 20×10×10, and 30×20×10, where the numbers denote modules, machines, and transport units, respectively. These settings represent small-, medium-, and relatively large-scale collaborative scheduling scenarios. Following common practice in flexible job shop scheduling, all instances are generated using the same stochastic rules. The number of operations per module is sampled from [0.8m, 1.2m], the number of eligible machines for each operation is sampled from [1,m], and processing times are drawn from [21, 59]. All transport units are assumed feasible, with transportation times sampled from [10, 20] based on current transport location and destination machine. Validation instances are generated independently under the same rules and parameter ranges to ensure consistency while testing robustness across different factory states.

Baseline Methods. To evaluate the proposed method, five baselines are considered, including three rule-based methods and two learning-based methods. FIFO+SPT selects the earliest available operation and assigns it to the machine with the shortest processing time. MOPNR+SPT prioritizes the module with the most operations remaining, while MWKR+SPT prioritizes the module with the maximum workload remaining; both use SPT for machine assignment. Ours w/o two-stage state encoding is an ablation variant that removes the staged resource-affinity state modeling while keeping the rest of the MARL framework unchanged. DANIEL (Wang et al., 2024) is a state-of-the-art DRL method with dual-attention state encoding for flexible job shop scheduling and serves as a strong learning-based benchmark.

5.2 Comparative Results and Analysis

5.2.1 Overall Comparison of Makespan

The comparison results in Table 1 show that the proposed method consistently achieves the best makespan across all three problem scales. The advantage is already visible on the small-scale instances and becomes more pronounced as the scale increases from 10×5×5 to 30×20×10. This result indicates that the proposed two-stage resource-affinity modeling is effective not only in limited scheduling environments but also in more congested production-logistics systems with a larger combinatorial search space. Among the three dispatching rules, MWKR+SPT generally performs better than FIFO+SPT and MOPNR+SPT, but all rule-based methods remain clearly inferior to the learning-based methods. This indicates that fixed priority rules are insufficient for capturing the dynamic interaction between production progress, machine status, and transport availability, especially in larger scheduling instances. Among the DRL methods, DANIEL achieves competitive performance and consistently outperforms the heuristic baselines, confirming the benefit of attention-based state representation. However, the proposed method still obtains the best makespan across all scales. This improvement mainly comes from the two-stage state encoding design, which better reflects the sequential dependency between production decisions and subsequent logistics decisions. The ablation

study further supports this finding. Although Ours w/o two-stage state encoding performs better than the heuristic rules, it remains consistently worse than the full model. This shows that the performance gain comes not only from reinforcement learning itself, but also from the explicit modeling of heterogeneous resource interactions across decision stages.

Table 1: Comparison of Makespan and Computation Time Across Different Problem Scales

Size		FIFO+SPT	MOPNR+SPT	MWKR+SPT	Ours w/o state encoding	Ours	DANIEL
10×5×5	makespan	669.88	787.45	780.31	578.66	523.25	542.23
	Time(s)	0.09	0.09	0.09	0.36	0.39	0.40
20×10×10	makespan	1440.09	1690.69	1698.69	1130.42	1037.26	1098.85
	Time(s)	0.42	0.42	0.42	1.64	1.68	1.71
30×20×10	makespan	2713.5	3406.25	3407.84	1933.58	1774.10	1851.02
	Time(s)	1.76	1.75	1.75	5.57	5.64	5.70

5.2.2 Training Behavior of DRL Methods

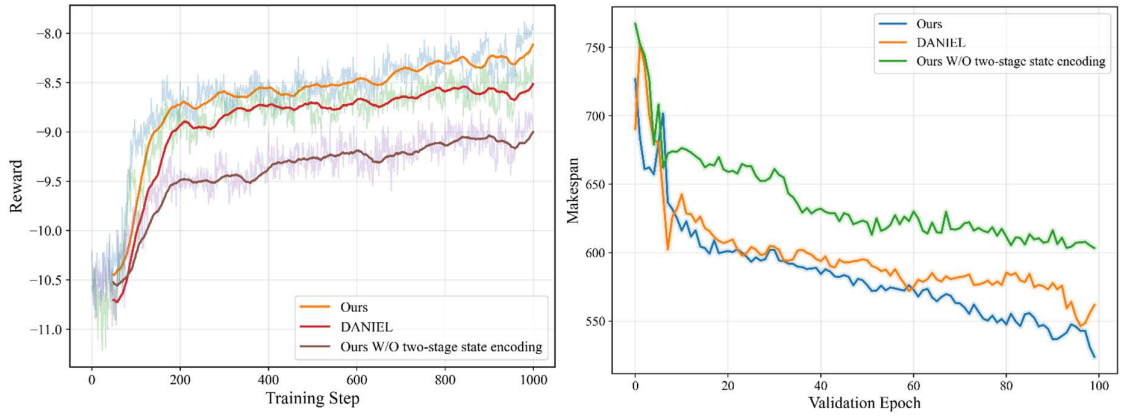


Figure 3: Training Curves of DRL-based Methods. (Left: Reward, Right: Makespan)

The training reward curves shown in Figure 3 further support the superiority of the proposed method. Compared with DANIEL and the ablation variant, the proposed model reaches higher reward values and converges more stably, indicating that the two-stage state encoding provides more informative features for policy learning. By contrast, the ablation model shows slower improvement and larger fluctuations, suggesting weaker decision guidance when resource interactions are not explicitly modeled. A similar trend is observed in the makespan curves. The proposed method reduces makespan faster during early training and converges to a lower final level. Although DANIEL also demonstrates effective learning, its final performance remains inferior. These results indicate that the proposed resource-affinity modeling offers a more suitable inductive bias for coupled production–logistics scheduling.

5.2.3 Qualitative Gantt-Chart Analysis

A qualitative comparison based on representative Gantt charts shown in Figure 4 further highlights the advantage of the proposed method over DANIEL. The proposed method produces more compact schedules with fewer machine idle periods and better-aligned transportation responses, resulting in smoother production–logistics coordination. In contrast,

the schedules generated by DANIEL show more fragmented waiting time, especially when transport timing is less well synchronized with machine execution. This difference is consistent with the model design. By explicitly capturing operation–machine affinity in the first stage and updating transport-unit competition in the second stage, the proposed method better aligns downstream transportation decisions with upstream production needs. As a result, it achieves shorter makespan and more continuous workshop flow.

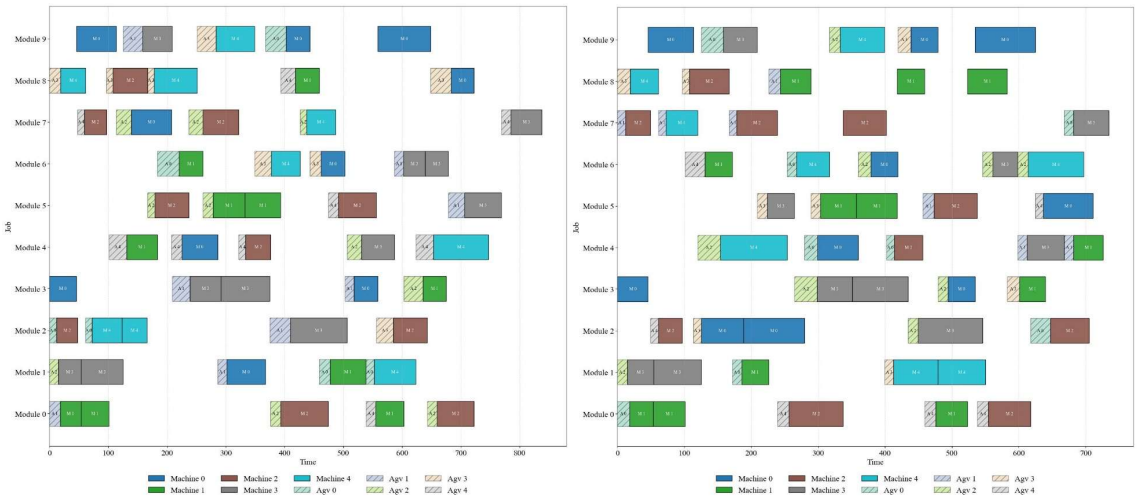


Figure 4: Representative Gantt charts of DANIEL and the proposed method. (Left: DANIEL, Right: Ours)

5.3 Discussion

The experimental results demonstrate three main findings. First, explicit modeling of heterogeneous resource relationships is critical for MiC production-logistics collaborative scheduling, especially when the system scale increases and resource interaction becomes more complicated. Second, the proposed two-stage state encoding contributes substantially to performance improvement, as verified by the consistent gap between the full model and the ablation variant. Third, compared with a strong attention-based DRL baseline, the proposed method better matches the sequential coupling structure of the problem and therefore learns more effective multi-agent scheduling policies. These findings confirm the importance of relation-aware state representation in dynamic heterogeneous factory scheduling environments.

6 Conclusion

This paper investigates the production-logistics collaborative scheduling problem in MiC factories from the perspective of interconnected intelligent manufacturing. A two-stage resource-affinity multi-agent reinforcement learning framework is proposed to capture the dynamic interaction between production and logistics resources. The method first models operation-machine affinity and then characterizes competition among transport units under updated production states, generating relation-aware representations for collaborative decision-making. Experimental results show that the proposed method outperforms classical dispatching rules, the ablation variant without two-stage state encoding, and an advanced DRL baseline in makespan performance. These results highlight the importance of explicitly modeling heterogeneous resource relationships for intelligent scheduling in IoT-enabled MiC factories. Future work will incorporate real-time sensing, uncertainty, and multi-objective decision-making to enhance practical applicability.

References

- Bao Z., V. Laovisutthichai, T. Tan, Q. Wang, W. Lu (2022): Design for manufacture and assembly (DfMA) enablers for offsite interior design and construction. Building Research & Information, v50, no3, 325-338.
- Dauzere-Perez S., J. Ding, L. Shen, K. Tamssaouet (2024): The flexible job shop scheduling problem: A review. European Journal of Operational Research, v314, no2, 409-432.
- Demir Y., S. K. Isleyen (2013): Evaluation of mathematical models for flexible job-shop scheduling problems. Applied Mathematical Modelling, v37, no3, 977-988.
- Dong X., G. Wan, P. Zeng (2025): A heuristic-assisted deep reinforcement learning algorithm for flexible job shop scheduling with transport constraints. Complex & Intelligent Systems, v11, no5, 210.
- Gui Y., D. Tang, H. Zhu, Y. Zhang et al. (2023): Dynamic scheduling for flexible job shop using a deep reinforcement learning approach. Computers & Industrial Engineering, v180, 109255.
- Jiang Y., M. Li, B. J. Ma, R. Y. Zhong, G. Q. Huang (2024): Data-driven out-of-order model for synchronized planning, scheduling, and execution in modular construction fit-out management. Computer-Aided Civil and Infrastructure Engineering, v39, no16, 2457-2480.
- Ozguven C., L. Ozbakir, Y. Yavuz (2010): Mathematical models for job-shop scheduling problems with routing and process plan flexibility. Applied Mathematical Modelling, v34, no6, 1539-1548.
- Pan W., Y. Yang, M. Pan (2023): Implementing modular integrated construction in high-rise high-density cities: perspectives in Hong Kong. Building Research & Information, v51, no3, 354-368.
- Peiris A., F. Hui, C. F. Duffield, J. Wang, M. Gil Garcia, Y. Chen, T. Ngo (2023): Digitalising modular construction: Enhancement of off-site manufacturing productivity via a manufacturing execution & control (MEC) system. Computers & Industrial Engineering, v178, 109117.
- Shen L., X. Yao (2015): Mathematical modeling and multi-objective evolutionary algorithms applied to dynamic flexible job shop scheduling problems. Information Sciences, v298, 198-224.
- Wan L., L. Fu, C. Li, K. Li (2024): Flexible job shop scheduling via deep reinforcement learning with meta-path-based heterogeneous graph neural network. Knowledge-Based Systems, v296, 111940.
- Wang H., L. Liao, W. Yi, L. Zhen (2024): Transportation scheduling for modules used in modular integrated construction. International Journal of Production Research, v62, no11, 3918-3931.
- Wang R., G. Wang, J. Sun, F. Deng, J. Chen (2024): Flexible Job Shop Scheduling via Dual Attention Network-Based Reinforcement Learning. IEEE Transactions on Neural Networks and Learning Systems, v35, no3, 3091-3102.
- Yang Z., W. Lu (2023): Facility layout design for modular construction manufacturing: A comparison based on simulation and optimization. Automation in Construction, v147, 104713.
- Zhang M., L. Wang, F. Qiu, X. Liu (2023): Dynamic scheduling for flexible job shop with insufficient transportation resources via graph neural network and deep reinforcement learning. Computers & Industrial Engineering, v186, 109718.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Place your logo here

LLM-Driven Multi-Agent Group Chat Framework for Decentralized Garment Supply Chain Coordination

Kexin Sun[1] and Zhiheng Zhao[1,2], and George Q. Huang[1,2]
1. Department of Industrial and Systems Engineering, The Hong Kong Polytechnic University, Hong Kong, China
2. Research Institute for Advanced Manufacturing, The Hong Kong Polytechnic University, Hong Kong, China
Corresponding author: zhiheng.zhao@polyu.edu.hk & gq.huang@polyu.edu.hk

Abstract: *The growing demand for mass customization and rapid fulfillment in the garment industry has intensified coordination challenges among distributed supply chain actors. While the Physical Internet paradigm promotes decentralized and intelligent logistics systems, practical mechanisms for real-time multi-party collaboration remain insufficiently developed. Existing studies primarily emphasize optimization models or digital tools at individual stages, leaving a gap in dynamic cross-actor coordination and collaborative decision-making under uncertainty. This limitation motivates the need for an intelligent framework that enables structured communication and decentralized planning across heterogeneous stakeholders. This paper proposes a Large Language Model (LLM)-driven multi-agent group chat framework for distributed garment supply chain management. The framework represents five key stakeholders, including Designer, Manufacturer, Warehouse, Logistics, and Retailer Agents, as role-based decision entities interacting through a structured conversational protocol. Within the group chat, agents exchange order details, capacity constraints, inventory levels, replenishment conditions, and delivery requirements, while a role-guided prompting strategy combined with shared state memory supports negotiation and consensus formation toward a feasible operational plan. A proof-of-concept case study involving a customized garment order with multi-actor interactions, capacity limitations, material shortages, and delivery-time pressure demonstrates the feasibility of the proposed approach in a moderately complex coordination scenario. The results indicate that LLM-mediated structured communication can enhance collaborative planning efficiency and decision transparency in decentralized garment supply chains.*

Keywords: *Physical Internet; Garment Supply Chain Coordination; Large Language Model; Multi-Agent Group Chat.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan: ☐ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories), ☐ Transportation Equipment, ☐ PI Networks, ☒ System of Logistics Networks, ☐ Vertical Supply Consolidation, ☒ Horizontal Supply Chain Alignment, ☐ Logistics/Commercial Data Platform, ☐ Access and Adoption, ☐ Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

1 Introduction

The garment industry is undergoing a significant transformation driven by increasing demand for mass customization and rapid order fulfillment (Bruce et al., 2004; Mikkola and Skjott-

Larsen, 2004; Xu et al., 2021). Customers now expect personalized products with shorter delivery times, which places considerable pressure on supply chain coordination. Unlike traditional mass production systems, customized garment production involves multiple distributed stakeholders—including designers, manufacturers, warehouses, logistics providers, and retailers—whose decisions must be dynamically aligned to meet specific order requirements.

In response to these challenges, the Physical Internet (PI) paradigm has been proposed to promote decentralized, open, and intelligent logistics systems (Montreuil, 2011). By enabling modular and interconnected operations, PI aims to enhance flexibility and efficiency across supply chain networks. However, despite its conceptual advantages, practical mechanisms for real-time coordination and collaborative decision-making among heterogeneous actors remain underdeveloped. In particular, decentralized environments require structured communication and negotiation processes to ensure that local decisions can collectively lead to globally feasible solutions (Swaminathan et al., 1998; Chan and Chan, 2004).

Existing research on garment supply chain management primarily focuses on optimization models or digital tools designed for individual stages, such as production planning, inventory control, or transportation scheduling (Wen et al., 2019). While these approaches provide valuable insights, they often assume centralized control or rely on static data, limiting their ability to support dynamic, cross-actor coordination under uncertainty. Meanwhile, recent advances in multi-agent systems and large language models (LLMs) have introduced new opportunities for distributed intelligence and autonomous decision-making (Guo et al., 2024; Wang et al., 2024). LLM-driven agents, in particular, demonstrate strong capabilities in natural language understanding, reasoning, and interactive communication. Nevertheless, current applications of LLM-based multi-agent systems largely lack domain-specific coordination mechanisms for decentralized garment supply chain management.

To address these gaps, this paper proposes a large language model (LLM)-driven multi-agent group chat framework for decentralized garment supply chain coordination. In the proposed framework, key stakeholders are modeled as role-based intelligent agents that interact through a structured group chat environment. Through iterative communication, agents exchange critical operational information—such as order details, capacity constraints, inventory levels, and delivery requirements—and collaboratively negotiate feasible plans. A role-guided prompting strategy combined with a shared state memory mechanism supports consistent information sharing and coordinated decision-making across agents. The main contributions of this study are threefold.

- First, we develop a novel role-based multi-agent framework that enables structured communication and decentralized coordination in garment supply chains.
- Second, we design an integrated mechanism that combines group chat interaction, shared memory, and role-guided decision-making to facilitate collaborative planning.
- Third, we demonstrate the feasibility of the proposed approach through a proof-of-concept case study, highlighting its potential to improve coordination efficiency and decision transparency in distributed supply chain environments.

2 Related Work

2.1 Coordination under Demand Uncertainty in Garment Supply Chains

Early studies mainly framed garment and fashion supply chains as high-volatility systems characterized by short life cycles, low predictability, and strong time pressure. Bruce et al.

(2004) argued that textile and clothing supply chains require agile or leagile structures to cope with unstable demand and fast market changes. Building on this logic, Choi (2006) showed that quick response supported by Bayesian information updating can improve ordering decisions for seasonal fashion products under uncertainty. Later, Xiao and Jin (2011) explicitly modeled lead-time-dependent demand uncertainty in a fashion apparel channel and demonstrated that markdown money and revenue-sharing contracts can improve coordination. As Wen et al. (2019) summarized, fashion retail supply chains remain strongly consumer-demand-driven and continue to face intertwined demand and supply uncertainties. More recent research moves from simple dyadic settings toward richer and more realistic coordination problems. Adhikari et al. (2020) extended the analysis to a five-level textile supply chain with simultaneous demand and supply uncertainty and showed that a wholesale price contract alone is insufficient for coordination. Recent results (Li et al., 2024; Alom et al., 2024) further indicate that stochastic demand, production uncertainty, shipment policy, and risk-sharing design jointly shape coordination performance in textile and fashion supply chains. Overall, the literature has evolved from quick-response thinking toward contract-based coordination under multiple sources of uncertainty.

2.2 Multi-Agent Systems for Supply Chain Decision-Making

Research on multi-agent systems (MAS) provides another important stream for decentralized supply chain decision-making. Swaminathan et al. (1998) offered one of the earliest foundational studies by modeling supply chain dynamics through interacting agents, control elements, and communication protocols. Chan and Chan (2004) then emphasized that distributed manufacturing supply chains need higher-level coordination mechanisms beyond isolated local planning systems. These classic studies established the basic idea that distributed actors can make local decisions while still contributing to system-level coordination. The recent literature shows a shift from simulation-oriented MAS toward language-enabled autonomous coordination. Srivastava et al. (2024) reviewed emerging applications of large language models in supply chain management, highlighting opportunities in optimization, risk management, and knowledge support. Guo et al. (2024) further noted that LLM-based multi-agent systems rely heavily on role profiling, communication design, and memory mechanisms for complex collaboration. This transition is especially relevant to supply chains, where many coordination tasks are inherently conversational and require iterative negotiation rather than one-shot optimization**.**

2.3 Research Gaps and Motivation

Despite the above progress, two limitations remain. First, garment and fashion supply chain studies still focus predominantly on optimization or contract models, usually centered on manufacturer-retailer interactions, even though recent work has begun to expand the chain structure (Adhikari et al., 2020; Alom et al., 2024). Such models are useful for analytical insight, but they provide limited support for real-time cross-actor communication and collaborative adjustment when order conditions change. Second, although LLM-enabled multi-agent research is developing quickly, its application to domain-specific and decentralized garment supply chain coordination remains limited (Guo et al., 2024; Srivastava et al., 2024). In particular, existing studies rarely combine stakeholder roles, structured group communication, and shared operational memory in one coordination framework. This gap motivates the present study, which develops an LLM-driven multi-agent group chat framework to support structured information exchange, negotiation, and consensus formation among distributed garment supply chain actors.

3 Methodology

The proposed framework considers five stakeholder agents in the case study: Designer, Manufacturer, Warehouse, Logistics, and Retailer. Each agent has its own role, capability setting, and local state, and is implemented using LLaMA-2 7B as the base large language model. For a customized order, the objective is to reach a coordinated operational plan through multi-agent interaction under capacity, inventory, delivery, and business-rule constraints.

3.1 Overall Framework

This study proposes an LLM-driven multi-agent group chat framework for decentralized garment supply chain coordination. As illustrated in Figure 1, the system consists of four tightly coupled components: a customized prompt strategy, a group chat mechanism, a shared memory mechanism, and a decision-making mechanism. Unlike centralized optimization approaches, the proposed framework models key supply chain stakeholders as role-based agents with explicit responsibilities and decision boundaries and enables them to collaboratively construct a feasible plan through structured conversational interaction. Operationally, the framework starts from an incoming customized order. The stakeholder agents then interact in a shared group chat environment to exchange and interpret order specifications, capacity constraints, inventory levels, delivery requirements, costs, and risks. During each round of interaction, each agent reasons from both its own local perspective and the shared global state summarized in the memory. Through repeated proposal generation, feedback, negotiation, and revision, the system gradually converges to a coordinated plan including task assignments, resource allocation, scheduling decisions, and delivery arrangements.

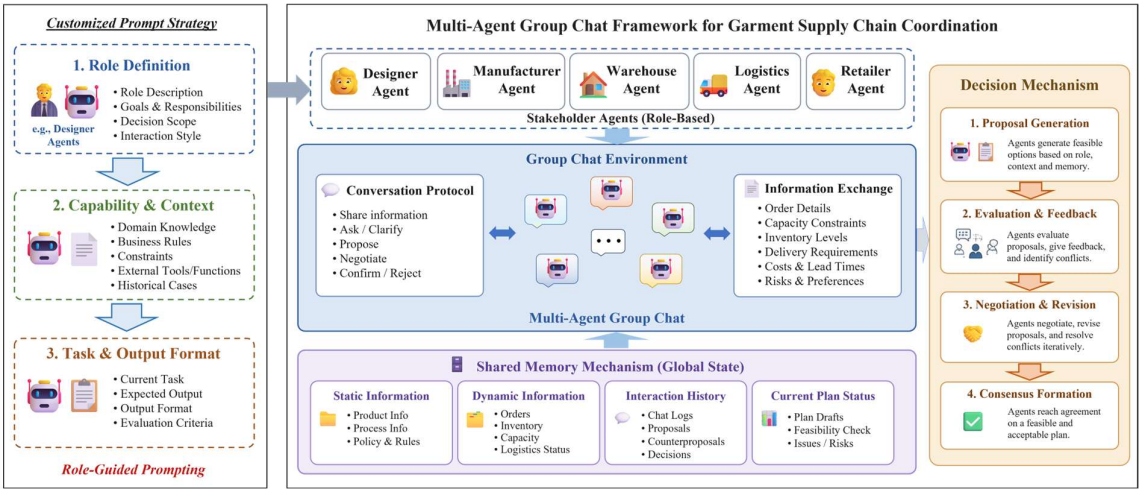


Figure 1: Overall framework of the proposed LLM-driven multi-agent group chat framework for garment supply chain Coordination

3.2 Customized Prompt Strategy

To ensure that heterogeneous agents behave consistently while preserving role specialization, a customized prompt strategy is designed for each stakeholder agent. The strategy contains three layers: role definition, capability and context, and task and output format. The role definition specifies the responsibilities, decision scope, and interaction style of the agent. The capability and context layer injects domain knowledge, business rules, operational constraints, available tools, and relevant background information. The task and output format layer specifies the current task, required output fields, and evaluation criteria. In practice, each prompt combines a global system instruction with role-specific information and a structured summary

of the latest shared memory, so that the agent can respond based on both its specialization and the current coordination state. This design plays three important roles. First, it enforces role-consistent reasoning. For example, the Manufacturer agent prioritizes production feasibility and capacity utilization, whereas the Warehouse agent focuses on stock availability and replenishment implications. Second, it improves output regularity by constraining the response structure, which facilitates later parsing, memory updating, and proposal integration. Third, the prompt remains dynamically adaptive to the evolving coordination state, reducing context drift across long interactions.

3.3 Group Chat Mechanism

The group chat mechanism serves as the interaction layer through which decentralized agents communicate and coordinate. All stakeholder agents are placed in a common conversational environment governed by a structured dialogue protocol. This protocol supports five major interaction types: information sharing, questioning/clarification, proposal generation, negotiation/revision, and confirmation/rejection. Compared with isolated pairwise communication, the group chat structure enables all agents to observe critical state changes and proposal updates in a synchronized manner, thereby reducing information fragmentation and repeated coordination effort. Let Γ_t denote the dialogue context accumulated up to round t. Then the message and local proposal generated by agent a_i at round t can be written as

$$(m_i^t, y_i^t) = f_i(\pi_i^t, \Gamma_t), \tag{1}$$

where m_i^t is the natural-language message produced by agent a_i, and y_i^t is the structured local proposal extracted from or explicitly provided in that message. If the agent only asks for clarification or shares information without advancing a plan, then $y_i^t = \emptyset$. After all agents complete their responses in round t, the dialogue context is updated as

$$\Gamma_{t+1} = \Gamma_t \cup \{m_i^t \mid a_i \in \mathcal{A}\}. \tag{2}$$

In the proposed framework, the group chat mechanism fulfills two core functions. First, it acts as a structured information exchange channel through which stakeholders communicate order details, capacity limitations, inventory conditions, delivery requirements, costs, and risks. Second, it acts as a negotiation arena in which locally feasible but potentially conflicting decisions can be compared, discussed, and revised. This is particularly important in garment supply chains, where many coordination constraints are tacit, context-dependent, or difficult to encode fully in advance.

3.4 Shared Memory Mechanism

To preserve state continuity across multiple rounds of interaction, the framework introduces a shared memory mechanism that stores and updates the evolving global state. Unlike approaches that rely only on raw chat history, the shared memory organizes coordination-relevant information in a structured form, providing all agents with a consistent reference point and enabling the system to track not only what has been said, but also what has been decided and what remains unresolved. The shared memory at round t is defined as

$$M_t = (S, D_t, H_t, P_t), \tag{3}$$

where S denotes static information, including product specifications, process information, policies, and rules; D_t denotes dynamic information, including order status, inventory levels, capacity availability, and logistics status; H_t denotes interaction history, including chat logs,

proposals, counterproposals, and intermediate decisions; and P_t denotes the current plan status, including draft plans, feasibility-check results, and identified issues or risks. After each interaction round, the shared memory is updated according to

$$M_{t+1} = U(M_t, \{m_i^t\}_{a_i \in \mathcal{A}}, \{y_i^t\}_{a_i \in \mathcal{A}}, e_t), \tag{4}$$

where $U(\cdot)$ is the memory update operator and $e_t$denotes optional execution feedback from the physical or simulated supply chain environment. Through this mechanism, the memory captures both the latest operational state and the rationale underlying the coordination process. An additional function of the shared memory is to generate a compact state summary for prompting. Instead of passing the full dialogue history to the language model, the framework derives a condensed context representation:

$$\sigma(M_t) = \rho(S, D_t, H_t, P_t), \tag{5}$$

where $\rho(\cdot)$ is a summarization function. The resulting summary may include updated inventory levels, unresolved conflicts, the latest plan draft, and major risks. This design reduces context length while improving consistency, traceability, and decision relevance in subsequent rounds.

3.5 Decision-Making Mechanism

Building on the group chat and shared memory mechanisms, the framework adopts a four-stage decision-making process: proposal generation, evaluation and feedback, negotiation and revision, and consensus formation. The goal is to transform distributed stakeholder inputs into a coordinated operational plan. In the proposal generation stage, each agent produces a local proposal based on its role, local state, dialogue context, and shared memory. These proposals may involve production scheduling, inventory allocation, transportation choices, or delivery commitments. The system then aggregates them into a candidate global plan:

$$\hat{y}^t = G(\{y_i^t\}_{a_i \in \mathcal{A}}, M_t), \tag{6}$$

where $G(\cdot)$ denotes the proposal aggregation operator. In the evaluation stage, the candidate plan is checked against global constraints such as capacity, inventory, and delivery requirements. Its infeasibility is measured by

$$\Phi(\hat{y}^t, M_t) = \sum_{k=1}^{K} \lambda_k \max(0, g_k(\hat{y}^t, M_t)), \tag{7}$$

where λ_k is the penalty weight for the k-th constraint. A smaller value of $\Phi(\hat{y}^t, M_t)$ indicates a more feasible plan. At the same time, each agent evaluates whether the plan is acceptable from its own operational perspective. If the plan is infeasible or unacceptable to one or more agents, the process enters the negotiation and revision stage. Agents discuss the source of conflicts and revise relevant parts of the plan, such as delivery schedules, batch allocation, or resource commitments. This stage is important for incorporating tacit local knowledge that is difficult to fully formalize in advance. Finally, in the consensus formation stage, the system checks whether sufficient agreement has been reached across agents and whether the revised plan satisfies the main constraints. Once these conditions are met, the candidate plan is accepted as the final coordinated plan $y^\star$.

4 Case Study

4.1 Case Setting

Case Source and Scenario Design. This study adopted a simulation-based proof-of-concept case design to examine whether the proposed multi-agent group chat mechanism can support structured communication, conflict identification, iterative revision, and coordinated decision formation in a realistic supply chain context. The simulation cases were constructed based on typical coordination conflicts and operational rules reported in the garment and textile supply chain literature, including order urgency, distributed inventory information, short-term capacity limitations, replenishment delay, and delivery-time pressure. Based on these literature-informed conflict patterns, ChatGPT was used to simulate ten coordination cases under different combinations of order quantity, inventory availability, production constraints, and delivery requirements. The experimental design focused on comparing the initial decision and the final coordinated decision in each simulated case. The initial decision refers to the preliminary fulfillment idea formed before full cross-actor constraint disclosure, whereas the final coordinated decision refers to the plan reached after multi-agent group discussion and revision. The objective was to evaluate whether the group chat framework could improve the quality and traceability of decentralized planning.

Evaluation Metrics. Four evaluation indicators were retained in this study. Feasibility, due-date satisfaction, and inventory consistency were manually rated on a 1–5 scale according to predefined case evaluation rules. Feasibility measures whether the plan is operationally executable under disclosed constraints. Due-date satisfaction reflects the likelihood that the promised delivery schedule can be met. Inventory consistency evaluates whether the proposed plan is consistent with the disclosed material and inventory conditions. In contrast, decision transparency was rated by ChatGPT on a 1–5 scale based on the clarity and explicability of the decision process as reflected in the group chat records and shared memory updates.

4.2 Illustrative Case Demonstration

To qualitatively demonstrate the proposed framework, one representative simulated case is shown in Figure 2 and 3. The case involves a customized order of 300 hoodies with a 12-day delivery deadline, coordinated by five stakeholder agents: Retailer, Designer, Manufacturer, Warehouse, and Logistics.

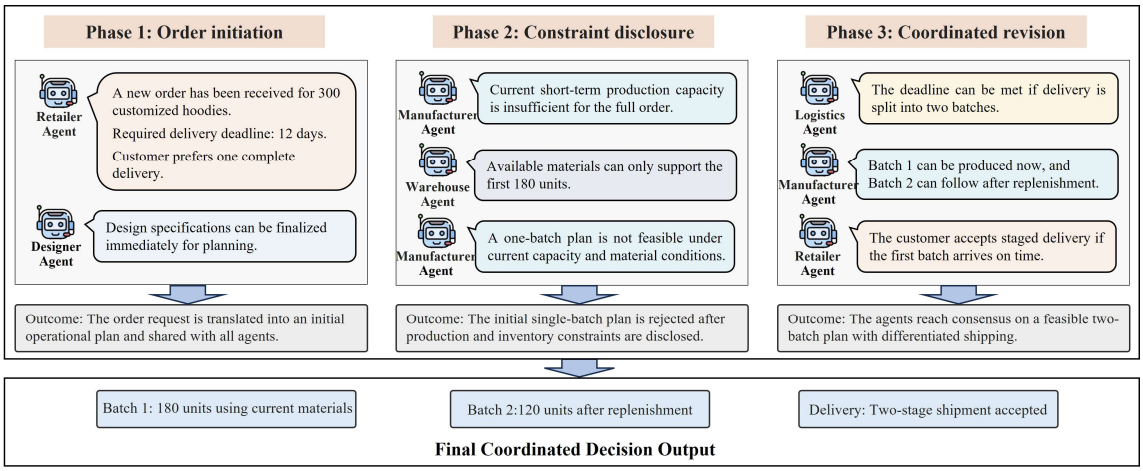


Figure 2: Iterative Coordination Process in Multi-Agent Group Chat

Figure 2 illustrates the coordination process in the multi-agent group chat. In Phase 1, the Retailer posts the order request and delivery expectation, and the Designer confirms that the specifications can be finalized for planning. In Phase 2, the Manufacturer and Warehouse disclose key constraints: short-term production capacity is insufficient, and available materials

can support only 180 units. Therefore, the initial single-batch plan is judged infeasible. In Phase 3, the agents revise the plan through discussion. The Logistics agent suggests staged delivery, the Manufacturer proposes batch-wise production after replenishment, and the Retailer confirms customer acceptance of split delivery. The agents finally agree on a feasible plan consisting of 180 units in Batch 1, 120 units in Batch 2, and a two-stage shipment arrangement.

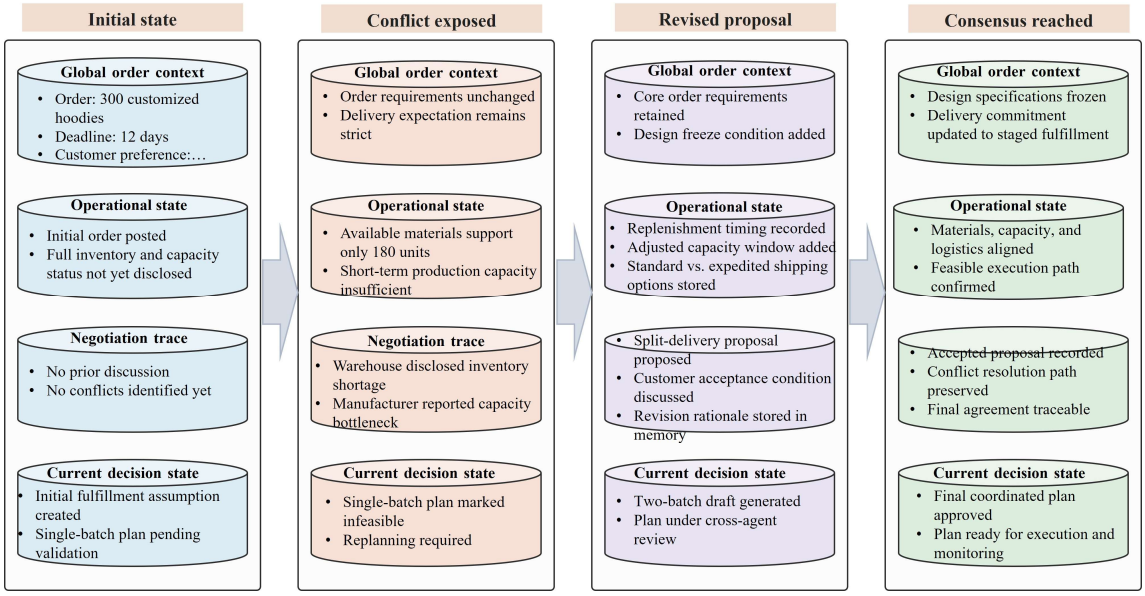


Figure 3: Shared Memory Evolution During Coordination

Figure 3 explains the same process from the perspective of shared memory. Initially, the memory stores the global order context, but full operational information is not yet available. After constraint disclosure, newly revealed information such as material shortage and capacity bottleneck is written into memory. During revision, replenishment timing, shipment options, and the split-delivery proposal are added. After consensus is reached, the accepted plan and its rationale are stored as the final decision state. Together, these two figures show that the framework can progressively expose hidden constraints, preserve coordination knowledge across rounds, and transform an initially infeasible plan into a feasible coordinated decision.

4.3 Multi-Case Evaluation Results

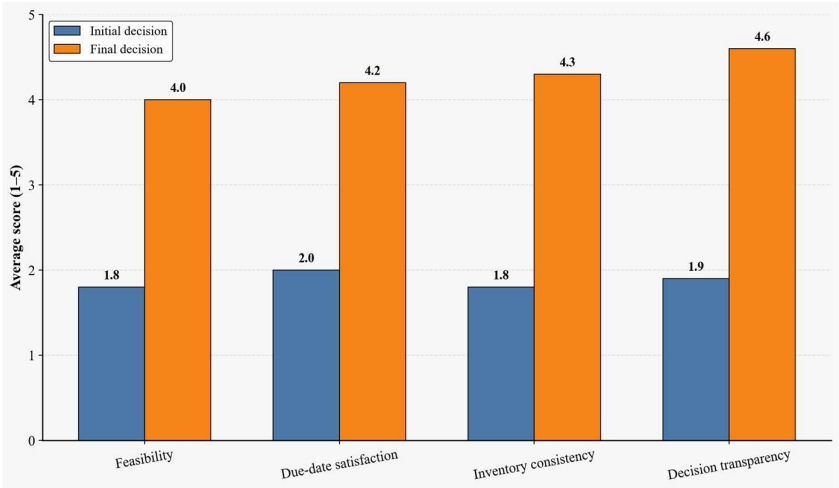


Figure 4: Performance Comparison Between Initial and Final Decisions

To complement the qualitative case, a small-scale evaluation was conducted on ten simulation-based cases. Figure 4 compares the average scores of the initial decisions and final coordinated decisions across four indicators. The results show consistent improvements after multi-agent coordination. Feasibility increased from 1.8 to 4.0, due-date satisfaction from 2.0 to 4.2, inventory consistency from 1.8 to 4.3, and decision transparency from 1.9 to 4.6. These results indicate that the final plans were more executable, better aligned with delivery and inventory conditions, and more traceable than the initial decisions.

Overall, the results suggest that the proposed framework improves decentralized planning both operationally and procedurally. It helps stakeholders generate more feasible and deadline-aware plans while also making the decision process more transparent through explicit chat interaction and shared memory updates. Since the evaluation is simulation-based, the findings should be interpreted as proof-of-concept evidence rather than full empirical validation. Even so, they provide initial support for the practical viability of the proposed multi-agent group chat framework.

4.4 Discussion

The case study and multi-case evaluation provide initial evidence of the practical viability of the proposed framework. The results suggest that multi-agent group chat can improve decentralized coordination by helping stakeholders reveal local constraints, revise infeasible plans, and reach mutually acceptable decisions. The observed improvements in feasibility, due-date satisfaction, inventory consistency, and decision transparency indicate that structured interaction combined with shared memory can support more operationally sound and traceable planning. These findings imply that the framework can serve as a lightweight alternative to purely centralized planning in garment supply chains with fragmented information, while also showing that LLM-based multi-agent systems can support collaborative planning under uncertainty when role design, memory, and decision rules are properly structured. Nevertheless, this study remains a proof-of-concept. The evaluation is based on simulated cases rather than real industrial deployment, and some indicators rely on qualitative or rule-based scoring. Therefore, the results should be interpreted as initial feasibility evidence rather than definitive validation. Future work should test the framework on larger case sets, more diverse disruption scenarios, and real-world operational data.

5 Conclusion

This paper proposed an LLM-driven multi-agent group chat framework for decentralized garment supply chain coordination. The framework integrates role-guided prompting, structured group chat, shared memory, and iterative decision-making to support cross-actor communication and collaborative planning in interconnected supply chain environments. A simulation-based case study and a small-scale multi-case evaluation were conducted to examine whether the framework can function effectively in literature-informed coordination scenarios. The results show that the proposed framework can progressively expose hidden constraints, preserve coordination-relevant information across interaction rounds, and generate more feasible, deadline-aware, and transparent plans than initial uncoordinated decisions. These findings suggest that multi-agent group chat is a promising mechanism for supporting decentralized coordination in customized garment fulfillment and can complement IoT-enabled visibility by turning distributed operational data into collaborative decisions. Overall, this study offers an initial step toward combining LLM-based multi-agent systems with intelligent and connected supply chain operations. While further empirical validation is still needed, the proposed framework demonstrates clear potential as a practical and explainable coordination approach for distributed garment supply chains.

References

- Adhikari A., A. Bisi, B. Avittathur (2020): Coordination mechanism, risk sharing, and risk aversion in a five-level textile supply chain under demand and supply uncertainty. European Journal of Operational Research, v282, no1, 93-107, https://doi.org/10.1016/j.ejor.2019.08.051
- Alom S., S. Basu, P. Basu, R. Joshi (2024): Shipment policy and its impact on coordination of a fashion supply chain under production uncertainty. Transportation Research Part E: Logistics and Transportation Review, v192, 103778, https://doi.org/10.1016/j.tre.2024.103778
- Bruce M., L. Daly, N. Towers (2004): Lean or agile: A solution for supply chain management in the textiles and clothing industry? International Journal of Operations & Production Management, v24, no2, 151-170, https://doi.org/10.1108/01443570410514867
- Chan F.T.S., H.K. Chan (2004): A new model for manufacturing supply chain networks: A multiagent approach. Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture, v218, no4, 443-454, https://doi.org/10.1243/095440504323055551
- Choi T.M. (2006): Quick response in fashion supply chains with dual information updating. Journal of Industrial and Management Optimization, v2, no3, 255-268, https://doi.org/10.3934/jimo.2006.2.255
- Guo T., X. Chen, Y. Wang, R. Chang, S. Pei, N.V. Chawla, O. Wiest, X. Zhang (2024): Large Language Model Based Multi-agents: A Survey of Progress and Challenges. Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, 8048-8057, https://doi.org/10.24963/ijcai.2024/890
- Li L., X. Liu, M. Hu (2024): Textile and apparel supply chain coordination under ESG related cost-sharing contract based on stochastic demand. Journal of Cleaner Production, v437, 140491, https://doi.org/10.1016/j.jclepro.2023.140491
- Mikkola J.H., T. Skjott-Larsen (2004): Supply-chain integration: implications for mass customization, modularization and postponement strategies. Production Planning & Control, v15, no4, 352-361, https://doi.org/10.1080/0953728042000238845
- Montreuil B. (2011): Towards a Physical Internet: Meeting the Global Logistics Sustainability Grand Challenge. Logistics Research, v3, no2, 71-87, https://doi.org/10.1007/s12159-011-0045-x
- Srivastava S.K., S. Routray, S. Bag, S. Gupta, J.Z. Zhang (2024): Exploring the Potential of Large Language Models in Supply Chain Management. Journal of Global Information Management, v32, no1, 1-29, https://doi.org/10.4018/JGIM.335125
- Swaminathan J.M., S.F. Smith, N. Sadeh (1998): Modeling supply chain dynamics: A multiagent approach. Decision Sciences, v29, no3, 607-632, https://doi.org/10.1111/j.1540-5915.1998.tb01356.x
- Wang L., C. Ma, X. Feng, Z. Zhang, H. Yang, J. Zhao, Z. Cheng, Y. Tang, X. Gao, Y. Chen, L. Zhang (2024): A survey on large language model based autonomous agents. Frontiers of Computer Science, v18, no6, 186345, https://doi.org/10.1007/s11704-024-40231-1
- Wen X., T.M. Choi, S.-H. Chung (2019): Fashion retail supply chain management: A review of operational models. International Journal of Production Economics, v207, 34-55, https://doi.org/10.1016/j.ijpe.2018.10.012
- Xiao T., J. Jin (2011): Coordination of a fashion apparel supply chain under lead-time-dependent demand uncertainty. Production Planning & Control, v22, no3, 257-268, https://doi.org/10.1080/09537287.2010.498601
- Xu Y., S. Thomassey, X. Zeng (2021): Garment mass customization methods for the cutting-related processes. Textile Research Journal, v91, no7-8, 802-819, https://doi.org/10.1177/0040517520957399

Inventory management of Traditional and Physical Internet-Enabled Logistics Network: A Simulation Study

Niyas K M[1], Sajan T John[2]

[1]Research Scholar, National Institute of Technology Calicut, Calicut, India

[2]Assistant Professor, National Institute of Technology Calicut, Calicut, India

Corresponding author: sajan@nitc.ac.in

Keywords: *Physical Internet, Inventory Control, Safety Stock, Ordering Cost, Logistics, Supply Chain*

Conference Domain Fitness:

This work addresses the transformation of the traditional hierarchical supply chain system into the Physical Internet (PI) enabled supply chain and investigates its impact on inventory control performance. The proposed Physical Internet-based network discussed in the work has the following PI features: inventory repositioning between hubs, multi-sourcing, and dynamic changes in inventory locations. The work helps to understand how PI principles influence operational efficiency, cost trade-offs, and resilience in distributed inventory systems.

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:☐ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),☐ Transportation Equipment, ☒ PI Networks, ☐ System of Logistics Networks, ☐ Vertical Supply Consolidation, ☐ Horizontal Supply Chain Alignment, ☐ Logistics/Commercial Data Platform, ☐ Access and Adoption, ☐ Governance*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

Research Paper Abstract

The current methods of managing logistics activities, particularly in transportation and inventory management, exhibit several limitations and challenges. The existing approaches are unsustainable from an environmental, social, and economic standpoint. Logistics services are customised to meet each customer's needs. The use of supply chain facilities often results in underutilisation of many of them. For example, the truck's capacity may not be fully utilised, as it may not carry a full truckload daily. Additionally, many vehicles make empty journeys because they need to be repositioned. Continuous long-haul trucking is another major problem. These extended hours of journey have an impact on the family life and health of the drivers. Due to the very dynamic nature of demand, dedicated hierarchical supply chains are somewhat inefficient.

The Physical Internet (PI) introduces an open, interconnected logistics network in which inventories are dynamically shared among PI hubs rather than a rigid hierarchical system. The existing studies related to inventory performance of Physical internet-based systems primarily focus on transportation and inventory-holding costs, without explicitly modelling ordering costs or the safety stock implications.

This research compares two supply chain scenarios: a traditional hierarchical multi-echelon network with fixed source-destination pairs, and a PI-enabled network that allows inventory repositioning between hubs, multi-sourcing, and dynamic changes in inventory locations in response to market demand variations. This study considers a single product scenario. The model operates under a continuous-review (Q, R) policy. The total logistics cost includes transportation, inventory holding, and ordering costs. The behaviour of the systems under safety stock is also examined.

An agent-based simulation framework evaluates both the supply chain scenarios. The supply chain scenarios are modelled and simulated using AnyLogic software. The experiments are conducted using hypothetical data. To compare the results, the same network of facilities is used across all scenarios. The demand and cost data are also the same for both scenarios. A comparative analysis focuses on total logistics costs and average inventory levels. Results suggest that the Implementation of Physical internet helps reduce average inventory levels and, thereby, total logistics costs. Safety stock analysis suggests that multi-sourcing flexibility may reduce buffer requirements in the presence of demand variability. This study extends PI-based inventory control models by integrating ordering cost and examining safety stock behaviour in interconnected logistics networks.

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE
The Hong Kong Polytechnic University 香港理工大學

A CPI-Enabled Resilient Flexible Job Shop Scheduling Framework for Shipyard Manufacturing in Industry 5.0

Yanying Wang[1] , Zhiheng Zhao[1] and Guoquan Huang[1]
1. Department of Industrial and Systems Engineering, The Hong Kong Polytechnic University, Hong Kong, China
Corresponding author: zhiheng.zhao@polyu.edu.hk

Abstract: *Shipyard manufacturing is a typical Engineer-to-Order environment characterized by high complexity, long process chains, and frequent uncertainty. To improve both efficiency and resilience in the Industry 5.0 context, this study proposes a Cyber-Physical Internet (CPI)-based scheduling framework for robotic welding lines. The scheduling problem is formulated as a Flexible Job Shop Scheduling Problem with sequence-dependent setup times (FJSP-SDST), where setup activities such as fixture changes and welding parameter adjustments are explicitly considered. To solve this problem, a setup-aware Hierarchical Attention Network-based Reinforcement Learning method is developed. In addition, CPI-enabled schedule caching and reuse support faster recovery and more resilient production under disruptions.*

Keywords: *Cyber-Physical Internet (CPI), Deep reinforcement learning (DRL), Flexible job shop scheduling.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan: ☐ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories), ☐ Transportation Equipment, ☒ PI Networks, ☐ System of Logistics Networks, ☐ Vertical Supply Consolidation, ☐ Horizontal Supply Chain Alignment, ☐ Logistics/Commercial Data Platform, ☐ Access and Adoption, ☐ Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

1 Introduction

Shipyard manufacturing is a typical Engineer-to-Order (ETO) process characterized by high complexity, strong discreteness, long process chains, and significant uncertainty. A complete shipbuilding process usually involves multiple interdependent stages, including cutting and forming, block assembly, pre-outfitting and painting, pre-erection, and erection. Meanwhile, Industry 5.0 emphasizes not only digitalization and automation, but also resilience and the ability to respond to uncertainty. In this context, shipyards must optimize production under normal conditions while also improving their ability to adjust and recover when equipment failures, new task arrivals, or on-site disruptions occur. Because a failure at one critical machine can easily propagate to downstream operations, maintaining production continuity under uncertainty has become a key challenge in shipyard manufacturing.

Against this background, this study focuses on production planning and shop-floor scheduling in shipyard manufacturing. Before final on-site assembly, robotic welding lines perform critical upstream tasks such as plate cutting, component processing, segment welding, and sub-assembly. In practice, these systems involve multiple robotic lines, multiple workstations, and alternative task-resource assignments, while also facing real-time disturbances such as new task arrivals, resource changes, and equipment failures. Therefore, the problem can be reasonably

modeled as a Flexible Job Shop Scheduling Problem (FJSP) (Song *et al.*, 2023; Wang *et al.*, 2025). When preparation and adjustment activities between tasks are further considered, it becomes an FJSP with sequence-dependent setup times (FJSP-SDST).

In robotic welding production, setup refers to the preparation and adjustment operations required when switching from one workpiece or segment to another. In this study, setup mainly includes fixture changes, welding program switching, torch or material parameter adjustment, workpiece repositioning, and robot trajectory calibration. Although these operations do not directly create production value, they strongly affect machine availability, waiting time, and the overall production rhythm. If setup cost is ignored, the resulting schedule may appear feasible but may still cause frequent switching, longer idle times, and increased completion time in practice. Therefore, explicitly considering setup cost is essential for improving both scheduling quality and execution efficiency in shipyard environments.

To cope with this complexity and uncertainty, this study further introduces the concept of the Cyber-Physical Internet (CPI). Originally explored in logistics and routing scenarios, CPI uses router-like nodes to collect state information, propagate updates, and support real-time decision-making in a manner similar to the Internet. Inspired by this idea, this study designs a CPI-based scheduling framework in which CPI routers are deployed to collect real-time production states, cache solved schedules, and rapidly retrieve reusable results when disruptions occur. After faults are identified and repaired, node states are updated so that scheduling inputs can be refreshed for subsequent recovery.

Based on this framework, this study makes two contributions. First, it reuses historical scheduling results through CPI, allowing solved schedules under repeated or similar states to be cached and reused, thereby reducing redundant rescheduling and improving real-time performance and resilience. Second, it proposes a setup-aware hierarchical attention network-based reinforcement learning method to capture complex relations among jobs, machines, setup costs, and competition, enabling more effective solutions to the FJSP-SDST than existing DRL baselines.

2 Related Works

2.1 Cyber -Physical Internet (CPI) and CPI Router

The idea of the CPI originated in the logistics domain (Zhao *et al.*, 2024). In that context, router-like nodes are responsible for sending and receiving freight "datagrams," enabling the dynamic routing of goods across physical transportation networks and thereby improving both efficiency and intelligence in logistics operations (You and Li, 2025). When transferred to manufacturing, CPI pursues a similar packet-switched mode of coordination: manufacturing tasks are "sent" and "received" across a factory network in much the same way that instant messages are transmitted over the Internet (Qu *et al.*, 2024). Although this vision resembles Cyber-Physical Systems (CPS), which coordinate physical manufacturing units through integrated information and control (Cecil *et al.*, 2019), CPI places greater emphasis on a concrete routing framework inspired by computer networks, where information is integrated to orchestrate collaboration among physical production units.

Accordingly, recent CPI research in logistics has mainly focused on router-level mechanisms, such as dynamic pricing for CPI routes (He *et al.*, 2024) and bandwidth allocation for routers (You and Li, 2025). In communication networks, routers connect subnetworks and forward datagrams hop by hop until they reach their destinations. By analogy, CPI routers installed at transportation hubs direct goods along the most appropriate physical routes (He *et al.*, 2024).

However, a dedicated CPI architecture for manufacturing has yet to be fully established. Work units on the shop floor and their associated information flows have not been tightly coupled with CPI, and the workflows and operational value that CPI could unlock in production systems remain underexplored.

2.2 Flexible Job-Shop Scheduling Problem

The Flexible Job-Shop Scheduling Problem (FJSP) is a common task-allocation problem in manufacturing. Given a set of jobs and machines, the objective is to assign each operation to one of several eligible machines and determine the processing sequence while satisfying technological constraints(Song *et al.*, 2023). In practice, the goal is usually to optimize performance measures such as makespan. As a classical NP-hard problem, FJSP has been widely studied using exact, heuristic, and metaheuristic methods.

Traditional FJSP studies often ignore setup times between operations. Since these costs depend on the processing sequence, they are referred to as sequence-dependent setup times (SDST) (Shen, Dauzère-Pérès and Neufeld, 2018). Research on FJSP with SDST has developed from heuristic and metaheuristic algorithms to reinforcement learning-based approaches. Various metaheuristics, such as variable neighborhood search, genetic algorithms and tabu search have been applied to SDST-related scheduling problems (Bagheri and Zandieh, 2011; Vallada and Ruiz, 2011). However, due to the strong coupling between jobs and machines, these methods often struggle to represent scheduling states accurately, while SDST further increases solution difficulty and slows convergence. As a result, they usually require case-specific redesign and extensive parameter tuning.

More recently, reinforcement learning has attracted increasing attention. Hybrid methods combining reinforcement learning with metaheuristics have been applied to SDST-related variants, but they still rely heavily on hand-crafted rules and numerous hyperparameters(Zhao, Jiang and Wang, 2023; Zhao, Zhou and Wang, 2023). With the emergence of deep reinforcement learning (DRL), end-to-end frameworks that automatically learn priority dispatching rules have become more promising. In particular, graph neural networks (GNNs), based on the disjunctive graph representation of FJSP, have been widely used as state encoders because they effectively capture complex job-machine interactions (Song *et al.*, 2023; Wang *et al.*, 2024, 2025). Nevertheless, most existing end-to-end DRL studies focus on JSP/FJSP rather than SDST-aware settings. Their graph-based representations generally hard to explicitly model

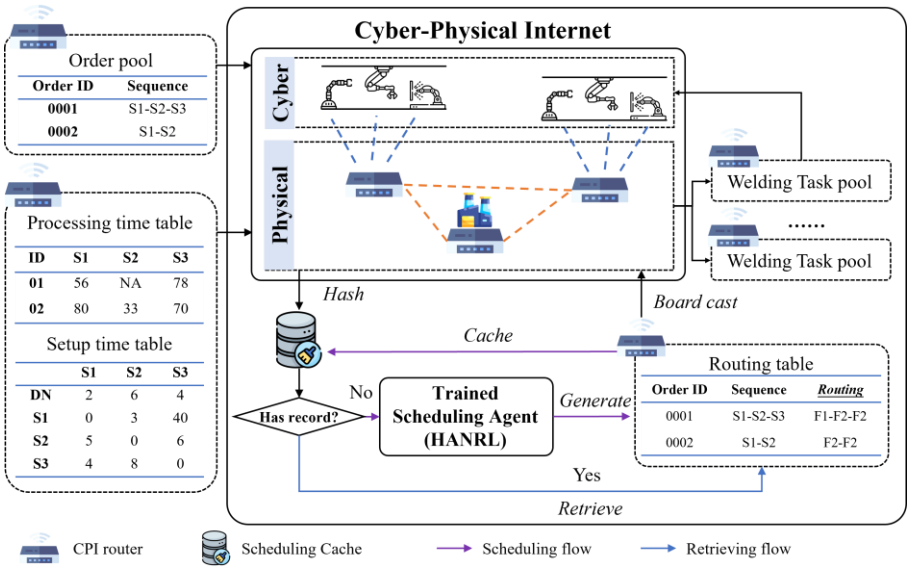


Figure 1: The CPI network framework

the substantial impact of setup times, indicating the need for a more accurate SDST-aware state representation framework.

3 Proposed Method

3.1 The CPI network framework

Shipyard manufacturing is a typical engineer-to-order (ETO) environment, characterized by high product variety, long and strongly coupled process chains, and substantial uncertainty. A complete shipbuilding process usually involves multiple interdependent stages. Within this production context, robotic welding lines are critical resources because they directly affect throughput, quality, and downstream assembly progress. Under the vision of Industry 5.0, shipyards are expected not only to improve productivity under normal conditions, but also to enhance resilience, and the ability to respond quickly to unexpected events. In practice, disturbances such as robot failures, urgent task arrivals and material delays often propagate across subsequent processes. As a result, maintaining continuity through timely production-line scheduling and recovery has become a central challenge in shipyard welding workshops.

To address this issue, this study designs a CPI-based scheduling framework deployed within a welding factory and across its robotic welding production lines.

3.1.1 The CPI network construction

The CPI network consists of one welding factory, multiple production lines with different processing capabilities, factory-level and line-level CPI routers, and the communication links among them. On the physical layer, the nodes are the welding factory and its robotic welding lines, while the edges represent their cooperative relationships for task assignment and execution. On the cyber layer, the nodes are the factory CPI router and the line CPI routers, each mapped one-to-one to its physical counterpart, and the edges indicate their ability to exchange digital information.

At initialization, CPI routers are deployed for the welding factory and each production line, and PIP addresses are automatically assigned. Time units and time zones are normalized for the job pool, processing-time table, and setup-time table, while job and operation codes are unified across the system. Each CPI router functions as an edge-computing microcomputer for collecting production data and updating routing tables through the scheduling agent, and communicates with other routers by sharing routing information. Earlier work suggests Raspberry Pi 3 B+ as a practical CPI router implementation.

The factory and line routers play different roles and store different data. The factory CPI router releases and archives jobs in a job pool, where each row corresponds to a welding job derived from ship blocks, panels, or sub-assemblies, and the Sequence field records the required operation order, e.g., “O1 → O2,” meaning Operation 1 must precede Operation 2. This is static information. By contrast, each line CPI router maintains local production data, including a processing-time table and a setup-time table. The processing-time table specifies, for every job and operation, whether the line can execute that operation and, if so, the required processing time; infeasible entries are marked “N/A.” These values are jointly determined by job specifications and line capabilities and thus become fixed once a batch of welding jobs is defined. The setup-time table records the time required to switch from one operation to another on the same production line. Setup time between identical operations is set to 0, and each operation has an initialization time from the dummy node (DN). Setup time is assumed to depend only on the previously executed and next operation, and the setup-time table is shared across all production lines once the job batch is determined.

3.1.2 Workflow of CPI network

Based on this construction, the CPI workflow supports both routine scheduling and disturbance-driven recovery. During data collection, production requirements are extracted from the platform database and enter the CPI system through Kafka. Based on the factory's BOM and process plans, these requirements are decomposed into materials and operations, forming job data recognizable by the scheduling system. For lines with MES, line CPI routers connect to local systems via APIs to obtain real-time shop-floor data and dispatch tasks through MES. For lines without MES, IoT devices can capture signals such as robot on/off status, machine utilization, workpiece counts, and operation progress.

During data aggregation, each line CPI router sends its processing-time and setup-time tables to the factory CPI router using a unified schema. On the factory side, these streams are normalized, deduplicated, and aligned through Kafka Streams to form a complete scheduling instance. The normalized instance is then hashed and stored with its scheduling result in a database so that identical or highly similar production states can later be recognized and reused.

Once the welding factory releases new jobs, the system decomposes them into schedulable job instances, which are collected in the factory job pool. Based on this information, the system retrieves from the $m$production-line CPI routers the corresponding processing-time and setup-time tables. In this way, an $n \times m$ scheduling instance is defined by the n jobs and the production data of the $m$robotic welding lines.

The hash values of the relevant scheduling tables are then checked in the Scheduling Cache. If no matching record exists, the instance is sent to the trained Scheduling Agent, which performs off-line scheduling and generates a production-line routing table. In the shipyard welding context, this routing table is essentially a line-level production schedule specifying, for each job, which production line should execute each operation and how these operations should be sequenced under current shop-floor conditions. The unseen instance and its resulting schedule are then stored in the cache. If the cache already contains the same hash, the corresponding schedule is retrieved directly, avoiding repeated computation in recurring scheduling environments. This cache-and-reuse mechanism is particularly valuable in shipyard welding workshops, where many blocks, panels, and repetitive welding tasks recur under similar processing conditions. By reusing previously solved scheduling instances, the CPI framework reduces scheduling calls and improves response efficiency. Finally, the retrieved or newly generated routing table is broadcast by the factory CPI router to all line routers, and each production line uses it to populate and update its task queue. When disturbances such as robot faults, urgent task insertions, or temporary line interruptions occur, the affected line router updates its node status in real time. Once the fault is identified and repaired, or a new task is confirmed, the updated status is sent back to the factory router so that the scheduling input can be refreshed for the next round of recovery or rescheduling. Thus, the CPI-based framework enables faster response, better reuse of historical schedules, and stronger resilience for shipyard robotic welding production under uncertainty. Communication among CPI routers can be implemented through Message Queuing Telemetry Transport (MQTT), while recently proposed dedicated intra-CPI communication protocols are also applicable.

3.2 Scheduling agent

This section details the construction of the scheduling agent within the CPI network. During training, the system draws on historical or simulated data cached in the CPI routers. We develop a Hierarchical Attention Network-based Reinforcement Learning (HANRL) framework to tackle the job dispatching problem with SDST. First, the production state is encoded by jointly describing job operations and machine conditions. These raw features are then fed into a

Hierarchical Attention Network, which combines GAT-style self-attention and classical cross-attention to capture both intra- and inter-entity interactions while explicitly incorporating setup times. Finally, the extracted representations are passed to a DRL-based decision-making module that determines which job operation should be assigned to which machine.

3.2.1 Problem Formulation

In a shipyard robotic welding line, the scheduling problem can be formulated as FJSP. Here, jobs correspond to welding tasks for different ship components, panels, or sub-assemblies, operations denote sequential welding processes, and machines represent alternative welding robots or robotic workstations. The task is to determine, for each job operation, the compatible machine, processing order, and start time, subject to technological precedence constraints and setup-related factors.

Formally, in an $n \times m$ instance, there are n jobs $J_i \in J$ and m machines $M_k \in M$. Each job consists of j operations processed in a fixed order, denoted by $O_i = \{O_{i1}, O_{i2}, \dots, O_{ij}\}$. An operation O_{ij} can be processed, without pre-emption, on any compatible machine $M_k \in M$ with processing time $p_{ij}^k > 0$. Switching from operation a to operation b incurs a setup time $u_{(a+1),b}$, operation b can start only after operation a is completed and the corresponding setup is finished. Thus, u forms an operation-to-operation setup-time matrix. Setup time is assumed to depend only on the previously processed operation and the next operation, and to be identical across machines. A feasible schedule assigns every operation O_{ij} to a compatible machine and determines its start time. The objective is to minimize the maximum completion time of all operations, i.e., the makespan. Therefore, the problem is an FJSP with SDST (FJSP-SDST).

For decision-making, the planning horizon is divided into discrete time slots $T = \{1,2, \dots, t\}$, and the complete operation set is $O = \bigcup_i O_i$. At time t, the state is defined as $s_t = (O^t, M^t, A^t)$, where O^t is the set of unfinished operations, M^t is the set of available machines, and A^t is the feasible action set. An action $a_t \in A^t$ selects a feasible operation-machine pair (O_{ij}, M_k), where feasibility requires all predecessors of O_{ij} to be completed and M_k to be idle, i.e., to have finished both processing and setup for its previous assignment.

The reward is designed to reduce makespan. Let $\underline{C}(O_{ij}, s_t)$ denote the lower bound of the completion time of operation O_{ij} at state s_t, with $\underline{C}(O_{i0}, s_t) = 0$. For unscheduled operations,

$$\underline{C}(O_{ij}, s_t) = \underline{C}(O_{i(j-1)}, s_t) + \min_{k \in M_{ij}} (p_{ij}^k + \bar{u}_i) \quad (1)$$

Accordingly, the makespan estimate at time t is $\max_{O_{ij} \in O} \underline{C}(O_{ij}, s_t)$. The immediate reward is

$$r_t = \max_{O_{ij} \in O} \underline{C}(O_{ij}, s_t) - \max_{O_{ij} \in O} \underline{C}(O_{ij}, s_{t+1}) \quad (2)$$

Over the whole scheduling process,

$$\sum_{t=0} r_t = \max_{O_{ij} \in O} \underline{C}(O_{ij}, s_0) - C_{\max} \quad (3)$$

where $C_{\max}$ is the final makespan. Since the first term is constant, minimizing the makespan is equivalent to maximizing the cumulative reward.

3.2.2 Hierarchical Attention Network

The scheduling problem can be formulated as a disjunctive graph, as in the FJSP. This representation makes it possible to extract operation-level and machine-level features through attention-based graph learning. However, operations and machines are heterogeneous nodes with different attributes, and their interactions cannot be adequately modeled by standard graph

attention alone [8]. To address this issue, a Hierarchical Attention Network is introduced to jointly encode operations, machines, and their cross-type dependencies.

The HAN structure is illustrated in Fig. 2. It contains two modules: an operation attention module for extracting operation features and a machine attention module for extracting machine features. The operation module contains two attention mechanisms, while the machine module contains one. The resulting operation and machine embeddings are then fused into a global representation, where ψ denotes all trainable parameters of the HAN.

Operation attention

The operation attention module consists of two parts. Attention Mechanism 1 evaluates the setup cost between each operation and its compatible machines, producing an operation representation that explicitly includes setup-time information. Attention Mechanism 2 then models dependencies between adjacent operations within the same job to obtain the final operation embedding.

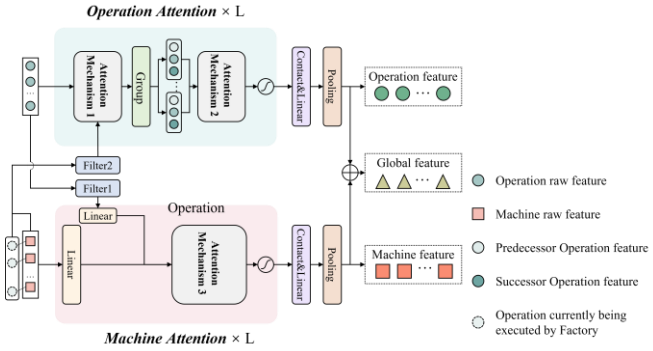


Figure 2: Hierarchical Attention Network

In Attention Mechanism 1, cross-attention biases each operation toward compatible machines with lower setup costs. Let the raw features of operation O_{ij} and machine M_k be $h_{O_{ij}}$ and h_{M_k}, respectively. The set of idle and compatible machines for O_{ij} is denoted by M_{ij}. For each candidate machine $M_k \in M_{ij}$, let O^k be the most recent operation processed on that machine, and let U_{ij}^k denote the setup time required to switch from O^k to O_{ij}. The attention coefficient is computed as

$$e_{ijk} = \text{softmax}\left(\frac{1}{\sqrt{d}} Q_{ij} K_k - \tau U_{ij}^k\right) \tag{4}$$

$$h_{ijk} = \sum_{k=1}^{M} e_{ijk} V_k \tag{5}$$

$$h_{O_{ij}}^M = h_{O_{ij}} + h_{ijk} \tag{6}$$

where $Q_{ij} = h_{O_{ij}} W_Q$, $K_k = h_{M_k} W_K$, and $V_k = h_{M_k} W_V$. Here, τ controls the influence of setup time. In this way, machines with lower switching cost and better feature matching receive higher attention, and the operation embedding is enriched accordingly through a residual connection.

Attention Mechanism 2 further refines each operation by attending only to its immediate predecessor and successor within the same job. This local design captures key job-level dependencies while keeping the computation efficient. Let the predecessor and successor of O_{ij} be $O_{i(j-1)}$ and $O_{i(j+1)}$. The attention coefficients over these neighboring operations and the operation itself are computed by

$$e_{ijp} = \text{LeakyReLU}\left(\vec{b}^{\,T}\left[Wh^{M}_{O_{ij}} \parallel Wh_{O_{ip}}\right]\right) \tag{7}$$

$$e'_{ijp} = \text{softmax}\left(e_{ijp}\right) \tag{8}$$

$$h'_{O_{ij}} = \sigma\left(\sum e'_{ijp} W h_{O_{ip}}\right) \tag{9}$$

where $p = j-1, j, j+1$.

Machine Attention

The machine attention module models competition among machines. Inspired by Wang *et al.*, 2024, Attention Mechanism 3 introduces competing operations as an additional signal so that the network can focus on the most critical machines. Let C_{km} denote the set of operations that may compete for machines M_k and M_m, and let N_k be the set of machines competing with M_k. The competition intensity between M_k and M_m is defined as

$$c_{km} = \sum_{O_{ij} \in C_{km}} h_{O_{ij}} \tag{10}$$

The attention score from M_k to $M_m \in N_k$ is then calculated as

$$u_{km} = \text{LeakyReLU}\left(\vec{d}^{\,T}\left[Z^1 h_{M_k} \parallel Z^1 h_{M_m} \parallel Z^2 c_{km}\right]\right) \tag{11}$$

$$u'_{km} = \text{softmax}\left(u_{km}\right) \tag{12}$$

$$h'_{M_k} = \sigma\left(u'_{kk} Z^1 h_{M_k} + \sum_{M_m \in N_k} u'_{km}\, Z^1 h_{M_m}\right) \tag{13}$$

This mechanism captures machine-level competition induced by shared candidate operations, enabling the model to identify the most influential machines under the current scheduling state.

4 Experiment

4.1.1 Experiment setup

We conduct comparative experiments on synthetic data and on two well-known FJSP benchmarks. The synthetic data comprises two types of instances:
SD1 from (Song *et al.*, 2023), these instances allow different orders to contain different numbers of steps.
SD2 (Wang *et al.*, 2024), in every $n \times m$ instance each order contains exactly m steps.

Note that, unlike the original data sets, we augment every instance with SDST. Setup times are generated from a uniform distribution $\text{U}(0.05\bar{p}_{ij}, 0.2\bar{p}_{ij})$. Additionally, 10 % of the values are drawn from a wider uniform range $\text{U}(0.8\bar{p}_{ij}, 1.2\bar{p}_{ij})$ to mimic real production scenarios require lengthy machine resetting to the extent that the setup time may exceed the processing time. Therefore, the distribution of setup time is:

$$F = \rho\text{U}(\ell_1\bar{p}_{ij}, \upsilon_1\bar{p}_{ij}) + (1-\rho) \times \text{U}(\ell_2\bar{p}_{ij}, \upsilon_2\bar{p}_{ij}) \tag{14}$$

where, $\rho = 0.1, \ell_1 = 0.05, \upsilon_1 = 0.2, \ell_2 = 0.8, \upsilon_2 = 1.2$.

For each instance size, 100 validation cases are generated. During training, the data are generated on the fly from the same distribution to prevent the model from overfitting. Models trained on small-to-medium instances are then evaluated on large instances using an additional 100 previously unseen test cases, in order to assess robustness.

4.1.2 Results on Synthetic Data

Table 1 reports the average makespan and runtime over 100 validation instances drawn from the same distribution as the training data. Across all data and problem sizes, HANRL achieves the shortest makespan, outperforming all priority dispatching rules (PDRs) and the state-of-the-

Table 1: Performance on two synthetic data of small to medium training size

	Size		PDRs				DRL	
			FIFO	MOR	SPT	MWKR	Wang *et al.*, 2024	Ours
SD1	10×5	$C_{\max}$	158.88	153.34	173.95	153.35	154.40	**143.87**
		Time	0.07	0.09	0.06	0.05	0.59	0.88
	20×5	$C_{\max}$	291.89	289.76	314.42	286.03	295.43	**284.1**
		Time	0.25	0.40	0.20	0.42	5.60	4.04
	15×10	$C_{\max}$	248.93	234.45	271.71	232.87	257.73	**225.57**
		Time	0.48	0.51	0.53	0.44	3.30	4.92
	20×10	$C_{\max}$	317.30	298.50	350.87	298.59	323.61	**281.08**
		Time	0.61	0.40	0.52	0.53	5.65	8.57
SD2	10×5	$C_{\max}$	817.48	815.89	800.68	850.71	642.56	**574.97**
		Time	0.07	0.08	0.07	0.08	0.78	0.90
	20×5	$C_{\max}$	1487.71	1540.54	1428.71	1528.64	1148.04	**1023.40**
		Time	0.15	0.14	0.15	0.15	1.20	2.27
	15×10	$C_{\max}$	1416.12	1336.50	1301.72	1352.30	1102.55	**872.50**
		Time	0.23	0.23	0.22	0.23	3.18	4.82
	20×10	$C_{\max}$	1697.98	1654.25	1615.65	1655.76	1238.52	**956.48**
		Time	0.32	0.39	0.44	0.45	4.57	7.65

art DRL baseline. In general, PDRs are faster than DRL methods: their runtime stays below one second and increases only slightly with problem size. However, this speed comes at the cost of much poorer schedules, and their performance deteriorates sharply as instance size grows. For example, on the 20 × 10 SD2 set, HANRL reduces the best PDR makespan by about 40%.

Among the PDRs, MWKR performs best and SPT worst on SD1, whereas the opposite holds on SD2. This difference arises from the characteristics of the two data sets. In SD1, job route lengths vary greatly, so MWKR gives priority to jobs with many remaining operations and helps avoid long tailing steps in later stages. By contrast, SPT tends to postpone long operations, causing severe tailing. In SD2, however, jobs have similar route lengths but varying processing times, so SPT improves machine turnover and reduces waiting time.

Compared with the advanced DRL method of Wang *et al.*, 2024, HANRL achieves better makespan on every small- and medium-scale data set under both sampling and greedy strategies, and the advantage becomes larger as problem size increases. On SD2 20 × 10 instances, for example, HANRL with sampling improves makespan by 22.8% over Wang *et al*. This gain is mainly because, when route lengths are similar as in SD2, sequence-dependent setup times become the dominant factor affecting makespan. By explicitly modeling setup costs through hierarchical attention, HANRL is more sensitive to setup penalties, and this advantage grows with the number of jobs and machines.

Table 2: Performance on two synthetic data of large size

	Size		DRL							
			Greedy				Sampling			
			Wang *et al.*, 2024		Ours		Wang *et al.*, 2024		Ours	
			15×10	20×10	15×10	20×10	15×10	20×10	15×10	20×10
SD1	30×10	$C_{\max}$	463.87	463.86	422.79	**412.29**	422.57	428.12	404.28	**400.58**
		Time	1.21	2.02	4.08	2.77	10.00	13.00	15.70	18.49
	40×10	$C_{\max}$	600.40	600.08	557.92	**543.49**	556.24	564.08	539.82	**532.34**
		Time	1.65	2.38	6.76	3.85	14.44	22.46	34.86	36.46
SD2	30×10	$C_{\max}$	1751.26	1688.67	**1340.09**	1358.43	1450.43	1412.35	**1226.55**	1230.37
		Time	2.45	1.82	3.84	2.79	11.24	13.61	20.28	18.19
	40×10	$C_{\max}$	2218.76	2091.47	**1677.60**	1696.41	1896.88	1809.05	1575.25	**1566.13**
		Time	1.84	4.92	7.14	3.68	15.25	23.49	35.11	36.84

Table 2 further compares the algorithms on larger synthetic instances after training on small- and medium-sized problems; among PDRs, only the two best-performing rules from Section 5.2.1 are included. Even when trained on smaller problems, HANRL still achieves the shortest makespan as job size increases. On SD2, the 20 × 10 model improves over Wang *et al* by 18% with sampling and 13% with greedy, indicating strong generalization to much larger job sets, which is valuable in real production. Using the greedy policy already yields a substantial improvement over the same policy in Wang *et al* while running faster, whereas sampling provides additional quality gains at the cost of some speed. Even so, all variants keep runtime within seconds. By contrast, obtaining exact solutions with OR-Tools often requires hundreds of seconds on medium- to large-scale instances. Therefore, practitioners can choose either strategy according to their needs for speed or solution quality.

References

- Bagheri, A. and Zandieh, M. (2011) "Bi-criteria flexible job-shop scheduling with sequence-dependent setup times—Variable neighborhood search approach," *Journal of Manufacturing Systems*, 30(1), pp. 8–15. Available at: https://doi.org/10.1016/j.jmsy.2011.02.004.
- Cecil, J. *et al.* (2019) "An Advanced Cyber Physical Framework for Micro Devices Assembly," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 49(1), pp. 92–106. Available at: https://doi.org/10.1109/TSMC.2017.2733542.
- He, S. *et al.* (2024) "Pricing and revenue allocation mechanisms for regional logistics networks in Cyber-Physical Internet," *Transportation Research Part C: Emerging Technologies*, 165, p. 104714. Available at: https://doi.org/10.1016/j.trc.2024.104714.
- Qu, X. *et al.* (2024) "Routing protocols for B2B e-commerce logistics in cyber-physical internet (CPI)," *Computers & Industrial Engineering*, 193, p. 110293. Available at: https://doi.org/10.1016/j.cie.2024.110293.
- Shen, L., Dauzère-Pérès, S. and Neufeld, J.S. (2018) "Solving the flexible job shop scheduling problem with sequence-dependent setup times," *European Journal of Operational Research*, 265(2), pp. 503–516. Available at: https://doi.org/10.1016/j.ejor.2017.08.021.
- Song, W. *et al.* (2023) "Flexible Job-Shop Scheduling via Graph Neural Network and Deep Reinforcement Learning," *IEEE Transactions on Industrial Informatics*, 19(2), pp. 1600–1610. Available at: https://doi.org/10.1109/TII.2022.3189725.
- Vallada, E. and Ruiz, R. (2011) "A genetic algorithm for the unrelated parallel machine scheduling problem with sequence dependent setup times," *European Journal of Operational Research*, 211(3), pp. 612–622. Available at: https://doi.org/10.1016/j.ejor.2011.01.011.
- Wang, D. *et al.* (2025) "Flexible robotic cell scheduling with graph neural network based deep reinforcement learning," *Journal of Manufacturing Systems*, 78, pp. 81–93. Available at: https://doi.org/10.1016/j.jmsy.2024.11.010.
- Wang, R. *et al.* (2024) "Flexible Job Shop Scheduling via Dual Attention Network-Based Reinforcement Learning," *IEEE Transactions on Neural Networks and Learning Systems*, 35(3), pp. 3091–3102. Available at: https://doi.org/10.1109/TNNLS.2023.3306421.
- You, Y. and Li, M. (2025) "Limited data-driven router bandwidth configuration for cyber physical internet," *Expert Systems with Applications*, 265, p. 126003. Available at: https://doi.org/10.1016/j.eswa.2024.126003.
- Zhao, F., Jiang, T. and Wang, L. (2023) "A Reinforcement Learning Driven Cooperative Meta-Heuristic Algorithm for Energy-Efficient Distributed No-Wait Flow-Shop Scheduling With Sequence-Dependent Setup Time," *IEEE Transactions on Industrial Informatics*, 19(7), pp. 8427–8440. Available at: https://doi.org/10.1109/TII.2022.3218645.
- Zhao, F., Zhou, G. and Wang, L. (2023) "A Cooperative Scatter Search With Reinforcement Learning Mechanism for the Distributed Permutation Flowshop Scheduling Problem With Sequence-Dependent Setup Times," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 53(8), pp. 4899–4911. Available at: https://doi.org/10.1109/TSMC.2023.3256484.
- Zhao, Z. *et al.* (2024) "Spatial-temporal traceability for cyber-physical industry 4.0 systems," *Journal of Manufacturing Systems*, 74, pp. 16–29. Available at: https://doi.org/10.1016/j.jmsy.2024.02.017.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

LLM-Augmented Strategic Framing for PI-Oriented Supply Chain Networks

Ghada Ben Meriem[1,3], Shenle Pan[1], Walid Klibi[2] and Mohamed Haythem Selmi[3]

1. Centre de Gestion Scientifique (CGS) – Mines Paris – PSL, Paris, France
2. The Centre of Excellence in Supply Chain (CESIT) – Kedge Business School, Bordeaux, France
3. Square Research Center, Neuilly-sur-Seine, France

Corresponding author: ghada.ben_meriem@minesparis.psl.eu

Keywords: *Physical Internet; Supply Chain Network Design; Decision Framing; Decision Support; Generative AI; Large Language Models.*

Conference Domain Fitness: *This contribution aligns with the conference theme "Supply Chain and Logistics Intelligence in the Era of the Physical Internet: Bridging High-Tech and Low-Tech Solutions." It proposes a novel methodological framework for the upstream phase of network design, problem framing, leveraging Generative AI to strengthen decision preparation in supply chain and logistics systems. By focusing on how decision problems are structured before formal modeling, the paper contributes to digital decision support for the design and transformation of interoperable Physical Internet-oriented networks.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☐ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☒ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☐ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

Research Contribution Abstract

Research problem:

Supply chain network design (SCND) concerns strategic decisions regarding the configuration and reconfiguration of supply chain networks, including the number, location, and capacity of production, distribution, and storage facilities; the selection of suppliers, partners, and markets to serve; the allocation of flows between network nodes; and the assignment of sourcing, production, and distribution policies across multiple echelons. These long-term decisions determine the structural architecture of the supply chain and involve substantial investments in facilities, technologies, and resource deployment. As such, SCND is not limited to the construction and solution of formal optimization models. It constitutes a structured strategic decision process that encompasses the definition of the design problem, the development of formal analytical representations, the evaluation of alternative network configurations, and the implementation of the selected design (Janjevic and

Klibi 2025). This decision process can be conceptualized through a general decision-making pipeline. Figure 1, adapted from (Powell 2026), illustrates this pipeline and situates SCND within it, highlighting the framing phase as a structuring interface between unstructured inputs and a formal problem description that underpins subsequent modeling.

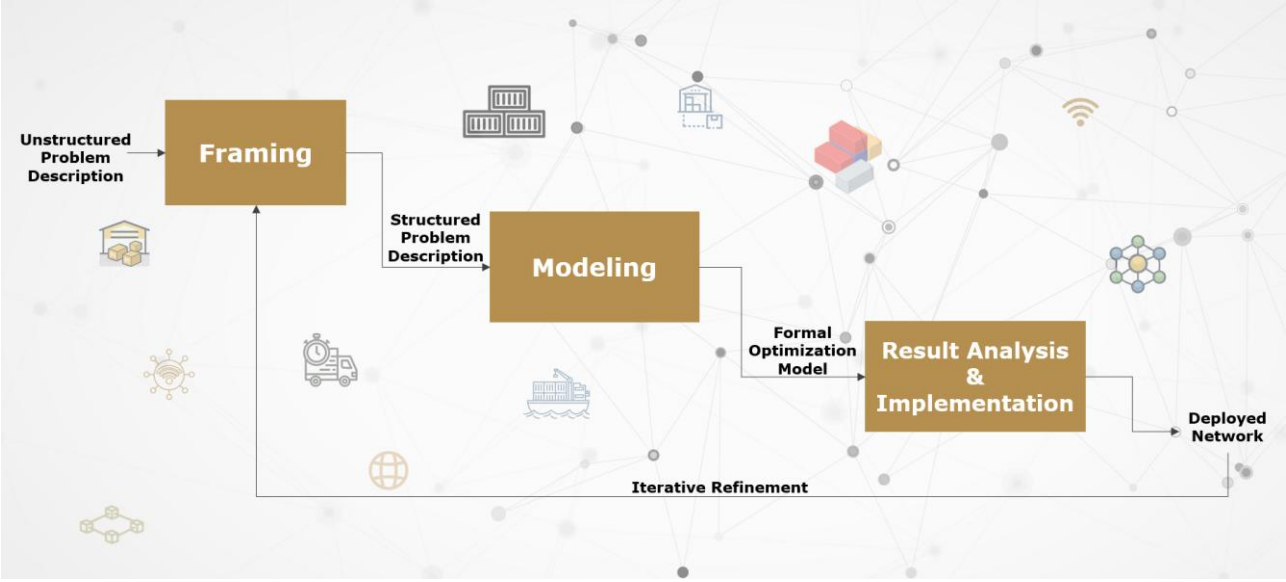


Figure 1: Decision-Making pipeline for SCND

A critical but often overlooked stage in this process is problem framing, defined as the stage in which the SCND problem is explicitly structured prior to modeling. This involves determining the scope of the network configuration, identifying the long-term resource deployment decisions, specifying performance metrics, objectives, and constraints, and making explicit the assumptions and uncertainties that shape the decision context. Framing establishes the problem statement that conditions the coherence and managerial relevance of subsequent modeling (Chevallier et al. 2023).

Recent advances in large language models (LLMs) have expanded the range of applications of generative AI in analytical and decision-support contexts. In operations management, including SCND, existing studies have primarily examined their ability to translate textual descriptions into formal optimization formulations and solver-ready code. Their potential role in supporting upstream problem structuring remains largely unexplored (Jackson et al. 2024).

The integration of LLMs into the framing phase of SCND is particularly pertinent in the context of the Physical Internet (PI). PI defines logistics as a system of interconnected logistics networks structured around coordination, modularity, and digital interoperability across multiple actors. It relies on shared infrastructures, coordinated freight routing, and common operational protocols to enable network-wide collaboration (Montreuil 2011). Designing such configurations requires that structural choices, capacity-sharing arrangements, and coordination mechanisms be explicitly defined at the problem formulation stage. Supporting the systematic articulation of these elements during SCND framing can therefore contribute to greater consistency between strategic network design decisions and the foundational principles of the PI (Pan et al. 2025).

In this context, LLM-augmented framing can support this process by acting as a semantic integration layer, enabling the synthesis of fragmented, multi-source strategic inputs into a coherent and

structured problem representation. Rather than replacing decision makers, the LLM facilitates the explicit articulation and alignment of key modeling elements, thereby improving the consistency and completeness of the framed problem prior to mathematical formulation.

Research gaps:

Effective decision-making in SCND depends on the quality of the problem framing that precedes modeling. Weaknesses in problem framing constitute a recurrent barrier to effective decision-making.

In real industrial settings, framing is neither straightforward nor given. It is a time-consuming and judgment-intensive activity that requires interpreting fragmented organizational information, reconciling multiple stakeholder perspectives, and translating domain-specific language into forms suitable for formal analysis. In practice, omissions, ambiguities, and implicit assumptions frequently persist at this stage and may undermine the robustness of subsequent models (Simchi-Levi et al., 2025). In parallel, existing research on LLMs in optimization has primarily focused on translating already specified problem descriptions into formal mathematical formulations or solver-ready code (AhmadiTeshnizi et al. 2024; Lu et al. 2025; Xiao et al. 2024; Huang et al. 2025). These approaches generally assume that the decision problem is clearly articulated and modeling-ready. This creates a gap between current LLM-based approaches and the realities of decision preparation in complex supply chain environments, where problem structuring itself remains a central challenge. The use of LLMs to support the construction, clarification, and validation of SCND problem statements prior to modeling remains largely unexplored (Jannelli et al., 2025).

Research questions and contributions:

This research addresses two related questions.

The first examines how large language models can support the construction of an explicit and structured framing of a supply chain network design problem from unstructured enterprise information. To address this question, we propose an approach in which LLMs generate a formalized intermediate representation of the SCND problem, enabling AI-assisted knowledge structuring and expert knowledge augmentation. Rather than generating a mathematical model, the output consists of a structured and human-readable representation that captures performance metrics, key decisions, and sources of uncertainty within a PI-oriented supply chain context. This representation supports a coherent understanding of the network design problem prior to formal modeling and promotes explainable and interpretable AI-supported decision preparation.

The second question examines how LLMs can support deeper analytical assessment during the framing phase by evaluating the internal coherence of the proposed network configuration and its temporal and strategic implications. To address this, we investigate how LLMs can help surface issues that often remain implicit in manual framing efforts. These include the interactions between performance metrics, decisions, and sources of uncertainty across PI components, enabling the identification of ambiguities in the problem description, inconsistencies across documents, and missing critical information required for modeling. This supports more informed discussions among decision-makers and reduces the risk of propagating flawed assumptions into subsequent modeling.

Taken together, this research develops a structured framework and evaluation protocol for LLM-assisted framing in SCND. The objective is to strengthen the rigor and clarity of upstream problem

definition, particularly through the explicit articulation of assumptions and sources of uncertainty. LLMs are positioned as decision-support tools that assist human judgment during framing, supporting expert-in-the-loop knowledge integration, while responsibility remains with the decision-maker. In the context of the PI, this support is especially relevant, as interoperable logistics networks rely on clearly defined and shared problem representations prior to formal modeling, enabling more explainable and collaborative human–AI decision-making.

Research methodology:

To address the identified gaps, this research adopts a four-phase methodology, as illustrated in Figure 2:

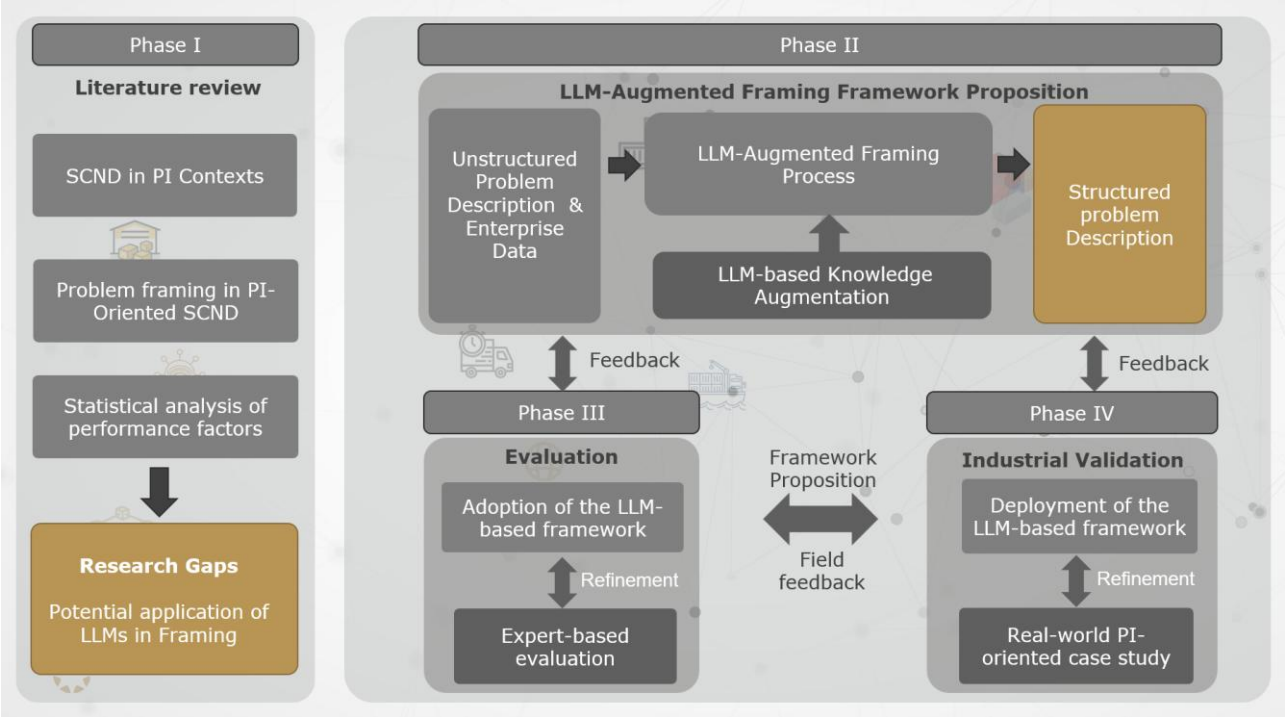


Figure 2: Decision-Making pipeline for PI-Oriented SCND

1. **Phase I: Literature review:** We analyze existing LLM-based approaches in Operations Research, categorizing them into prompting-based, agent-based, and learning-based methods. This analysis highlights a key limitation: most approaches assume that problem descriptions are already well-defined and ready for modeling. As a result, ambiguous natural language inputs are directly translated into mathematical models or solver code, bypassing the upstream problem-framing phase. This limitation is particularly critical in Physical Internet (PI) contexts, where multi-actor settings, shared infrastructures, and interoperability constraints must be explicitly structured prior to formal modeling.
2. **Phase II: Framework Development:** Building on Powell's conceptualization of problem framing, we propose an LLM-augmented framing framework for PI-oriented SCND. The framework structures heterogeneous inputs into a human-readable problem description centered on performance metrics, decisions, and sources of uncertainty. This intermediate layer supports upstream framing before formal optimization modeling and prepares the problem for subsequent translation into representation aligned with the Universal Modeling

Framework (UMF) (Powell 2026). The technical implementation explores document ingestion approaches, structured prompting strategies to guide LLM outputs, and validation protocols to mitigate hallucination risks and ensure factual grounding against enterprise data.

3. **Phase III: Evaluation:** The proposed framework is evaluated against traditional manual approaches through expert feedback, focusing on clarity, completeness, and decision-making relevance. The feedback is used to refine the framework and improve the framing process.
4. **Phase IV: Industrial Validation:** The proposed framework is implemented and applied to a real-world PI-oriented case study using unstructured enterprise data. This phase evaluates the ability of LLMs to structure heterogeneous, multi-actor inputs into a coherent problem description.

Phases II, III, and IV are currently in progress. Initial results from the framework development will be presented at the conference, including a real-world example demonstrating the transformation from unstructured problem descriptions to structured problem representations.

References:

AhmadiTeshnizi, Ali, Wenzhi Gao, and Madeleine Udell. 2024. "OptiMUS: Scalable Optimization Modeling with (MI)LP Solvers and Large Language Models." arXiv:2402.10172. Preprint, arXiv, February 15. https://doi.org/10.48550/arXiv.2402.10172.

Chevallier, Arnaud, Albrecht Enders, and Jean-Louis Barsoux. 2023. "Become a Better Problem Solver by Telling Better Stories." *MIT Sloan Management Review* 64 (3). https://sloanreview.mit.edu/article/become-a-better-problem-solver-by-telling-better-stories/.

Huang, Chenyu, Zhengyang Tang, Shixi Hu, et al. 2025. "ORLM: A Customizable Framework in Training Large Models for Automated Optimization Modeling." *Operations Research* 73 (6): 2986–3009. https://doi.org/10.1287/opre.2024.1233.

Jackson, Ilya, Dmitry Ivanov, Alexandre Dolgui, and Jafar Namdar. 2024. "Generative Artificial Intelligence in Supply Chain and Operations Management: A Capability-Based Framework for Analysis and Implementation." *International Journal of Production Research* 62 (17): 6120–45. https://doi.org/10.1080/00207543.2024.2309309.

Janjevic, Milena, and Walid Klibi. 2025. "Fostering Robustness in Supply Chain Network Design : A Critical Review and Future Opportunities." *Under Review*.

Lu, Hongliang, Zhonglin Xie, Yaoyu Wu, Can Ren, Yuxuan Chen, and Zaiwen Wen. 2025. "OptMATH: A Scalable Bidirectional Data Synthesis Framework for Optimization Modeling." arXiv:2502.11102. Preprint, arXiv, February 21. https://doi.org/10.48550/arXiv.2502.11102.

Montreuil, Benoit. 2011. "Toward a Physical Internet: Meeting the Global Logistics Sustainability Grand Challenge." *Logistics Research* 3 (2–3): 71–87. https://doi.org/10.1007/s12159-011-0045-x.

Pan, Shenle, Eric Ballot, and George Q. Huang. 2025. "Operations with Physical Internet." In *Reference Module in Social Sciences*. Elsevier. https://doi.org/10.1016/b978-0-443-28993-4.00171-2.

Powell, Warren. 2026. *Framing the Problem*. Warren B Powell.

Xiao, Ziyang, Dongxiang Zhang, Yangjun Wu, et al. 2024. "Chain-of-Experts: When Llms Meet Complex Operations Research Problems." *The Twelfth International Conference on Learning Representations*, March 15. https://openreview.net/forum?id=HobyL1B9CZ.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Optimization in Multimodal Hubs under the Physical Internet Paradigm Through Environmental KPIs

Jihane Lamri[1], Monica-Juliana Perez[2] and Tarik Chargui[2]

1. Institut National des Sciences Appliquées Hauts-de-France (INSA), University Polytechnic Hauts-de-France, 59313, Mont Houy, France
2. LAMIH, UMR CNRS 8201, University Polytechnic Hauts-de-France, 59313, Mont Houy, France

Corresponding author: Jihane.Lamri@uphf.fr

Keywords: *Physical Internet; Multimodal Hubs; Network Optimization; Environmental KPIs; Sustainable Logistics; multimodal transportation.*

Conference Domain Fitness: *This contribution aligns with the IPIC 2026 objective of advancing Supply Chain and Logistics Intelligence in the era of the Physical Internet by developing decision support methods that bridge advanced optimization techniques with practical logistics applications. The work contributes to hyperconnected, interoperable, and sustainable logistics ecosystems by proposing a multimodal network optimization model that supports open asset sharing, flow consolidation, and environmentally responsible transport decisions. It addresses key IPIC domains including multimodal transportation, collaborative logistics, and sustainable green logistics while integrating quantitative environmental performance indicators into operational decision-making.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☒ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☒ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

Research Contribution Abstract

The globalization of exchanges has deeply transformed logistics systems, leading to a significant increase in commodity flows and an intensification of transport operations. Within this context, logistics performance can no longer be limited to economic and temporal criteria; it must integrate major environmental stakes, as freight transport is a primary contributor to greenhouse gas emissions. The Physical Internet (PI or π) paradigm offers an innovative vision based on the interconnection, pooling, and standardization of logistics resources. This study proposes a multi-objective optimization model applied to a regional logistics network inspired by the PI, seeking a compromise between total cost, which integrates CO_2 emissions, and global transport time.

The logistics network is modeled as a multimodal graph where nodes represent hubs and arcs represent transport links, including road, rail, and water. Unlike traditional models that treat hubs as abstract points, this work includes a detailed representation of transshipment operations, wait times, and equipment capacities. The model optimizes two conflicting objectives: the minimization of total

cost, which incorporates direct transport, transshipment, and environmental costs, and the minimization of total time, which comprises travel, operational, and wait times at hubs. A fundamental contribution is the inclusion of specific PI-inspired constraints, namely a mandatory minimum vehicle fill rate of 90% to enforce the PI consolidation principle of maximizing vehicle utilization and thereby reducing CO_2 emissions per transported ton, and the estimation of hub wait times using the Erlang-C queuing formula to capture realistic congestion dynamics based on arrival rates, service rates, and handling equipment availability at each hub. Furthermore, since transshipment costs and times in multimodal systems strongly depend on the local characteristics of each hub, a Simple Additive Weighting (SAW) method is employed to compute a hub-specific efficiency factor based on multimodality, throughput capacity, equipment availability, and network specialization, allowing the model to reflect the real heterogeneity of hub infrastructures without requiring a direct estimation of local parameters for each site.

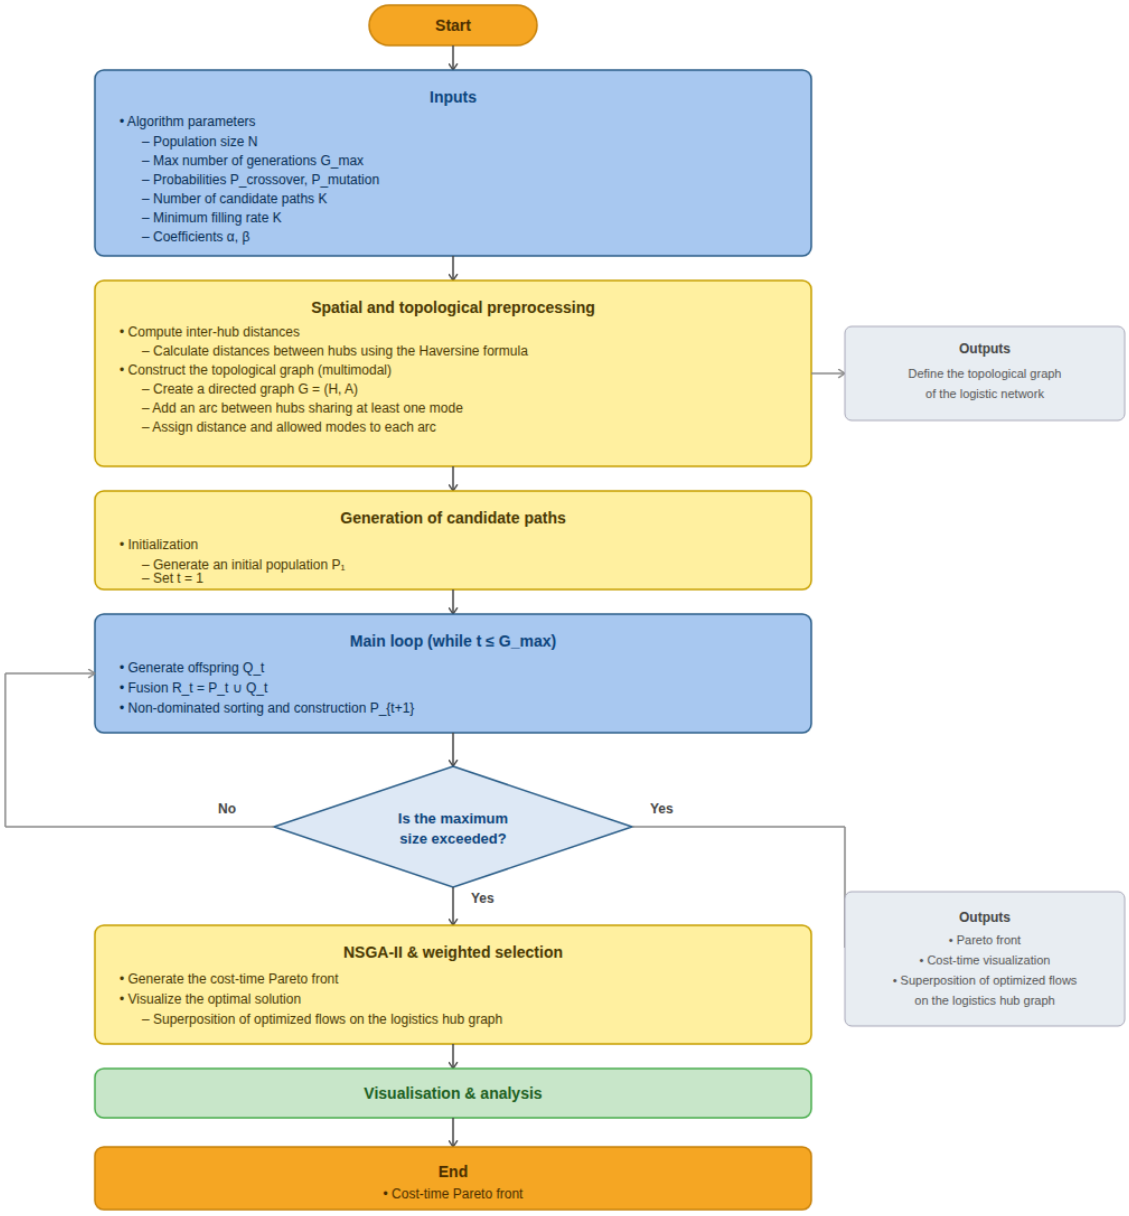


Figure 1:Mathematical Model Solution Approach

The proposed optimization model is formulated as a bi-objective mixed-integer non-linear program who's combinatorial and NP-hard nature, stemming from the simultaneous determination of routing sequences, modal choices, transshipment operations, and flow quantities under multiple logistic and environmental constraints, renders exact solution methods computationally intractable at realistic network scales. The Non-dominated Sorting Genetic Algorithm II (NSGA-II) is therefore adopted as the solution framework, chosen for its proven effectiveness in sustainable transport optimization, its capacity to explore large solution spaces while handling conflicting objectives without any arbitrary a priori aggregation, and its well-established balance between convergence and solution diversity. Each candidate solution is encoded as a chromosome representing a complete logistics chain, capturing the sequence of visited hubs, the selected transport modes, and the quantities routed on each network arc. Through its non-dominated sorting mechanism, the algorithm identifies successive Pareto fronts prioritizing solutions that are not outperformed on any objective, while the crowding distance operator ensures a well-distributed exploration of the cost-time trade-off space, providing decision-makers with a broad and meaningful spectrum of operational compromise solutions.

This optimization framework was implemented in Python and applied to a regional network comprising ten strategic hubs within the Hauts-de-France region. The process generated a series of non-dominated solutions forming the Pareto front, each representing a distinct trade-off between total cost, transport time, and environmental impact. Since the Pareto front alone does not yield a unique operational decision, a weighted aggregation function combining both objectives through a weighting factor α was introduced to select a single compromise solution. A value of $\alpha = 0.9$ was retained, reflecting a deliberate prioritization of economic and environmental performance, given that CO_2 emissions are monetized and embedded within the cost function, so that minimizing cost implicitly penalizes environmentally inefficient routing decisions. The resulting optimized flow, which emphasizes the strategic use of rail and fluvial modes to achieve massification and pooling in line with PI principles, is projected onto the regional topology as illustrated in Figure 2.

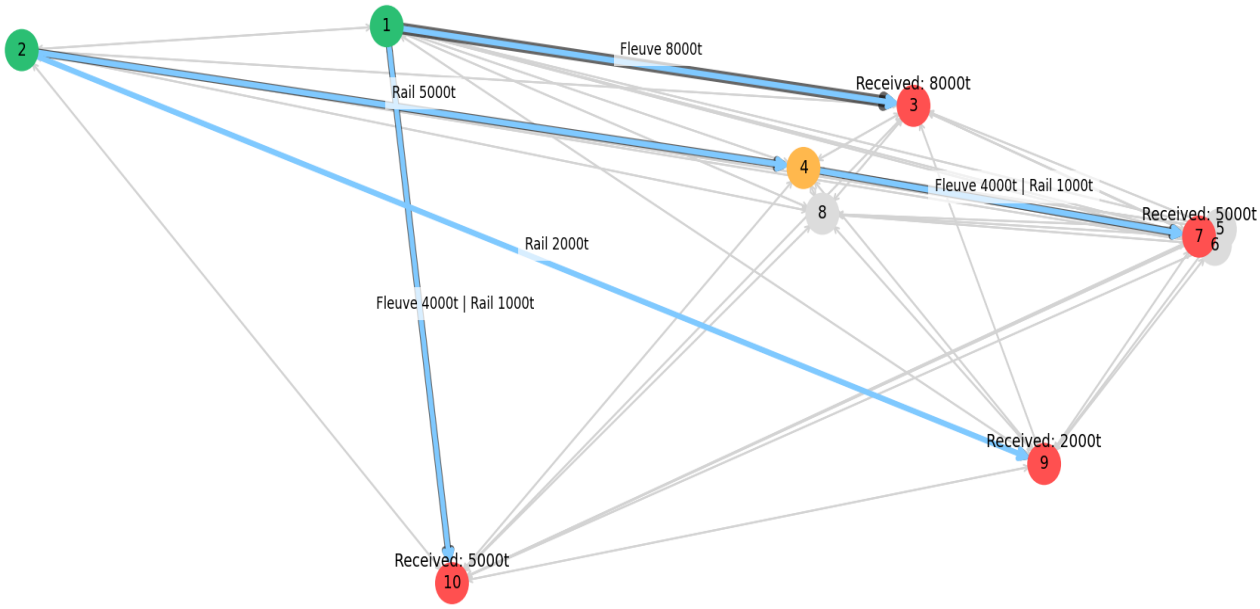


Figure 2: Projection of the optimized operational solution onto the regional multimodal network topology.

The transition from static optimization to a dynamic environment was achieved through a discrete-event simulation in Arena, which serves a purpose that goes beyond mere validation of the static model. While the optimization framework provides theoretically optimal routing and modal

decisions, it necessarily relies on analytically estimated wait times that are capped at saturated hubs, introducing a conservative bias in cost and emission estimates that does not reflect the actual operational behavior of the network. The simulation addresses this fundamental limitation by endogenously generating wait times from actual resource utilization dynamics, explicitly accounting for temporal synchronization between flows, real-time equipment availability, and concurrent competition for handling resources such as port cranes and locomotives, phenomena that no static formulation can fully internalize. Beyond correcting this analytical bias, the dynamic simulation also reveals emergent system behaviors that are invisible to the optimizer, including the effective utilization rates of hub infrastructure, the distribution of value-added versus non-value-added operations along the intermodal chain, and the actual degree of congestion experienced under the optimized flow configuration. This dynamic perspective demonstrated that system-wide congestion under the optimized load is in fact lower than analytically predicted, confirming that the static model tends to overestimate hub saturation under realistic operating conditions. The simulation therefore does not simply confirm the optimization results but actively enriches them, providing operational insights on resource dimensioning and flow synchronization that substantially strengthen the credibility, robustness, and practical deployability of the proposed decision-support framework for sustainable regional logistics under the Physical Internet paradigm.

This study demonstrates that coupling a PI-inspired multi-objective optimization model with a discrete-event simulation constitutes a robust and operationally credible decision-support framework for sustainable multimodal logistics. By explicitly modeling hub dynamics through transshipment operations, congestion-based wait times, and equipment-dependent efficiency factors, the proposed approach bridges the gap between theoretical network optimization and practical logistics management. The integration of environmental KPIs directly within the optimization objectives rather than as post-hoc evaluation metrics ensures that sustainability considerations are embedded in the routing and modal allocation decisions from the outset. The resulting framework provides logistics operators and policymakers with a structured methodology for evaluating the economic, temporal, and environmental trade-offs inherent in Physical Internet network design, paving the way for more informed and environmentally responsible regional logistics planning.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Incorporation of a Social Fatigue KPI in Physical Internet Hub Operations: A Comparative Study between FIFO and Social Policies

Pedro Nunes Maia[1], Tarik Chargui[2] and Monica Perez[3]

1. INSA Hauts-de-France, Valenciennes, France
2. INSA Hauts-de-France, Valenciennes, France
3. INSA Hauts-de-France, Valenciennes, France

Corresponding author: pedro.nunes@uphf.fr

Keywords: *Physical Internet, Social KPI, Fatigue Index, Logistics Hub, Discrete Event Simulation, Human-Centric Logistics.*

Conference Domain Fitness: *This research is primarily positioned within Sustainable and green logistics by addressing the critical social dimension of driver fatigue in Physical Internet operations. It extends the scope of Hyperconnected supply chain & logistics systems by demonstrating how human centric protocols can be integrated into automated hubs without compromising efficiency thereby supporting Transition strategies and implementation roadmaps for socially responsible logistics.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☒ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☐ *PI Networks,* ☐ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☒ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☐ *In-Person presentation*

Research Contribution Abstract

The Physical Internet (PI) paradigm envisions a hyperconnected logistics system capable of satisfying the Triple Bottom Line of sustainability: economic, environmental, and social. However, despite extensive optimization efforts regarding cost and environmental performance, the social dimension remains critically underrepresented in PI operations. Human resources are still frequently treated as infinite operational capacities, particularly in cross-docking and dispatching processes, where standard First-In First-Out (FIFO) logic ignores driver physiological conditions and cumulative fatigue risks. This issue is particularly relevant considering that fatigue-related factors are associated with a significant proportion of road accidents (MacLean et al., 2003; Zhang et al., 2020).

This research addresses this gap by extending human-centric logistics and safety-aware dispatching approaches toward a fatigue-aware dispatching mechanism for Physical Internet hubs, through the integration of a Social Fatigue KPI directly into dispatching decisions. The study focuses on a single simulated PI Hub located in Lille and connected to five predefined destinations (Paris, Le Havre, Strasbourg, Lyon, and Marseille). These destinations include both short-haul and long-haul

transportation missions, generating heterogeneous fatigue accumulation profiles. The problem addressed is formulated as a dispatching and driver allocation problem, distinguishing it from classical vehicle routing problems (VRP): when a consolidation node reaches capacity, the system determines which parked truck and driver should be assigned to the next transportation mission.

The study employs a Discrete Event Simulation (DES) model developed in FlexSim, based on a multimodal hub architecture adapted from the RITMEA project. The simulation parameters were calibrated using operational assumptions commonly adopted in cross-docking and PI Hub simulation studies. Truck arrivals follow an exponential distribution (λ = 1.5 minutes), while driver sleep duration is randomly distributed between 4 and 9 hours to represent variability in real operating conditions. The simulation horizon corresponds to a daily operational window from 06:00 to 21:00. Each scenario was evaluated through 100 simulation iterations, and statistical significance was assessed using ANOVA tests with a 95% confidence level ($\alpha = 0.05$). Additional sensitivity analyses confirmed that the relative performance advantage of the Social Policy remained stable under moderate variations in fatigue parameter assumptions.

The proposed Fatigue Index (FI) is adapted from Zhang, Wu and Zhang (2020) and aggregates four physiological dimensions: fatigue associated with prior sleep duration (Fts), cumulative driving fatigue (Ftd), circadian rhythm effects (Ftc), and restorative recovery during parking periods (Ftr). The index is computed through the additive formulation:

$$FI(t) = Ftd + Fts - Ftr + Ftc$$

where the recovery component is subtracted because it represents physiological restoration. A fatigue value above 100 is considered a critical risk threshold according to Zhang et al. (2020).

Two dispatching policies were comparatively evaluated:

1. **Baseline FIFO Policy**: trucks are dispatched according to accumulated parking waiting time, without considering physiological conditions.
2. **Social Fatigue Policy**: when a consolidation node reaches capacity, the algorithm evaluates the predicted fatigue score associated with the destination route for every available truck and assigns the driver with the lowest projected fatigue level.

The quantitative results demonstrate that the Social Policy significantly improves safety-related indicators without degrading operational performance. The maximum fatigue level observed in the network decreased by 8.92%, while average fleet fatigue decreased by 3.65% ($p < 0.001$). In parallel, no statistically significant differences were observed regarding the total number of consolidated containers shipped by the hub or the average congestion level across the hub's operational nodes, suggesting that social constraints can be integrated into PI Hub operations without compromising productivity.

However, the study also reveals an important emergent phenomenon referred to as the "Equity Paradox." While the Social Policy reduces dangerous fatigue peaks, it simultaneously polarizes waiting times inside the hub. The maximum time spent by a driver inside the hub increased by 121.04%, while the shortest observed waiting duration decreased by 93.07%, generating highly

unequal waiting conditions among drivers. The standard deviation of drivers' waiting times increased from 66 minutes under FIFO to 171 minutes under the Social Policy, revealing substantial operational inequity. To mitigate this effect, the study proposes a hybrid derogation mechanism in which drivers exceeding a waiting threshold of 376 minutes are automatically assigned to the next mission regardless of their fatigue score, thereby balancing safety objectives with fairness constraints.

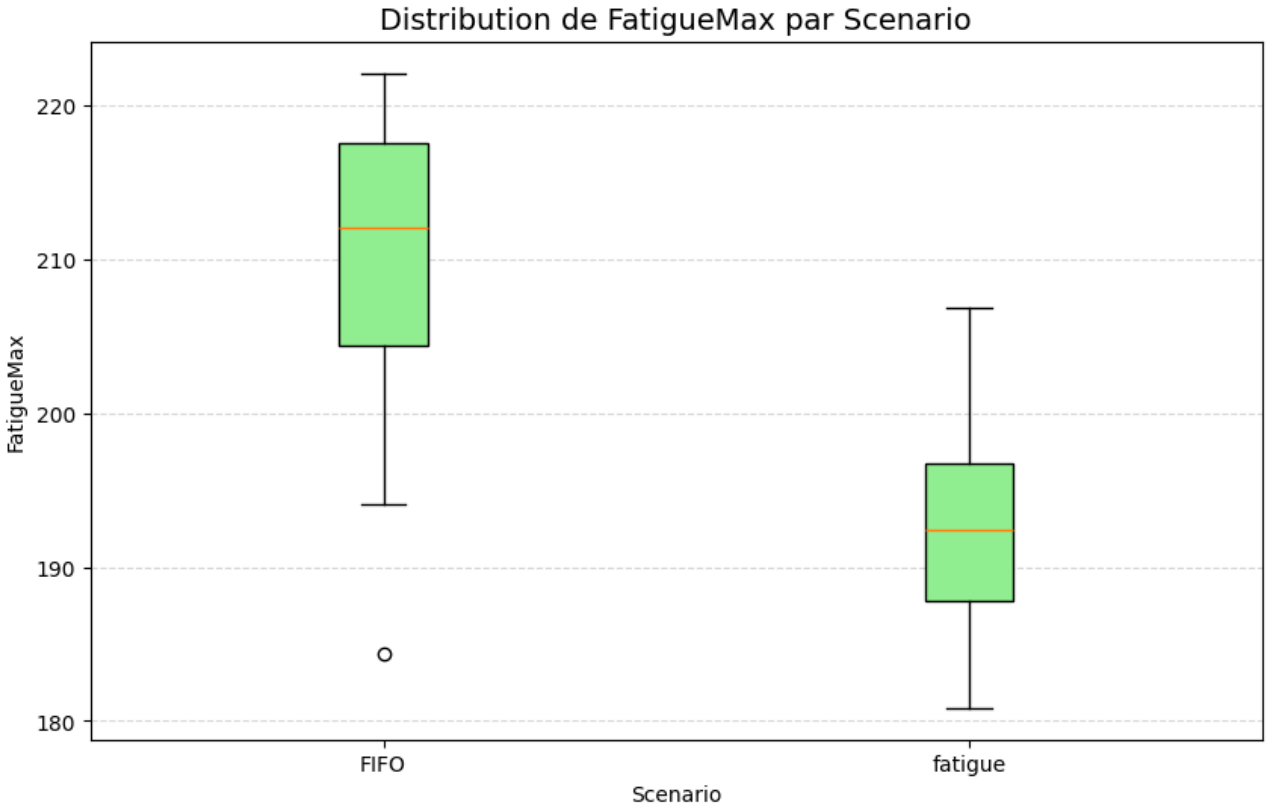


Figure 9 - Distribution comparée des FatigueMax et analyse de la dispersion entre les scénarios FIFO et social. (Source: Auteur)

This research contributes to the operationalization of social sustainability within the Physical Internet by providing a concrete fatigue-aware dispatching framework for PI nodes. Although the proposed mechanism currently operates as a local, hub-level policy, the approach is designed to remain compatible with broader PI governance perspectives and future inter-hub implementations. The results demonstrate that human-centric PI operations can improve safety outcomes while preserving logistics efficiency, while also highlighting the need to integrate fairness constraints into future fatigue-aware dispatching protocols.

References

A. W. MacLean, D. R. T. Davies, and K. Thiele, "The hazards and prevention of driving while sleepy," Sleep Medicine Reviews, vol. 7, no. 6, pp. 507–521, 2003

ZHANG, Qiang, WU, Chao et ZHANG, Hui. *Driving fatigue prediction model considering schedule and circadian rhythm*. 2020.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Optimizing Multi-Modal Physical Internet Hubs: A Mathematical Model for Collision-Free Container Routing

Tarik Chargui[1], Gabriel Zambrano-Rey[2], Monica-Juliana Perez[1]
1. LAMIH, UMR CNRS 8201, Université Polytechnique Hauts-de-France, 59313, Mont Houy, Valenciennes, France
2. LISPEN, Arts et Métiers Institute of Technology, F-59046, Lille, France

Corresponding author: tarik.chargui@uphf.fr

Keywords: *Physical Internet, PI-Hubs, Collision Avoidance, Multi-Modal Logistics, Mixed-Integer Programming.*

Conference Domain Fitness: *This paper proposes a mixed-integer programming (MIP) model for Physical Internet (PI) container routing. The model addresses road-road, road-rail, and water-road PI-hub configurations. It integrates collision avoidance constraints to ensure safe container transfers. The objective is to minimize total routing and transfer time. The formulation is generic and adaptable to multiple PI-hub architectures. Adaptation is achieved through parameter configuration without structural changes. The model is implemented using IBM ILOG CPLEX. Several test instances were evaluated to assess model performance. Results demonstrate efficient and collision-free routing solutions. The proposed approach provides a versatile operational tool for PI-container managements.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☒ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☐ *PI Networks,* ☐ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☐ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

Research Contribution Abstract

The Physical Internet (PI) paradigm represents a transformative vision for freight transportation, promoting modularization, standardization, and interconnectivity of logistics operations through the use of standardized PI-containers and interoperable PI-hubs. These hubs enable seamless intermodal transfers across road, rail, and inland waterway systems, forming the operational backbone of next-generation logistics networks. Despite the growing interest in PI systems, their practical deployment remains constrained by the lack of unified optimization tools capable of addressing the heterogeneous architectures, safety requirements, and operational complexities inherent to modern PI-hubs (Ballot 2012, Ballot 2019, Chargui et al. 2019).

This paper proposes a comprehensive mixed-integer programming (MIP) model for optimizing the routing of PI-containers across heterogeneous PI-hub configurations, including road–road, road–rail, and water–road hubs. Unlike existing approaches that focus on specific hub types or neglect

operational safety aspects, the proposed formulation provides a generic and adaptable framework that can be tailored to diverse PI-hub architectures through parameter configuration alone. A distinctive feature of the model is the explicit integration of collision avoidance constraints, ensuring safe and fluid container transfers within shared conveying networks while minimizing total routing time.

The addressed problem, referred to as the container recomposition problem in PI-hubs, involves the simultaneous routing of multiple containers with heterogeneous dimensions through a common network of conveying units. Containers arrive at different times from inbound vehicles such as trucks, trains, or ships and must be transferred to designated outbound docks without spatial–temporal conflicts. The optimization challenge is compounded by discrete space–time constraints, non-overlapping requirements for containers sharing conveying resources, and multi-modal compatibility constraints linked to vehicle interfaces. To capture these realities, the model explicitly considers the initial and final positions of each PI-container within inbound and outbound vehicles, enabling automatic adaptation to varying vehicle sizes and configurations (Figure 1).

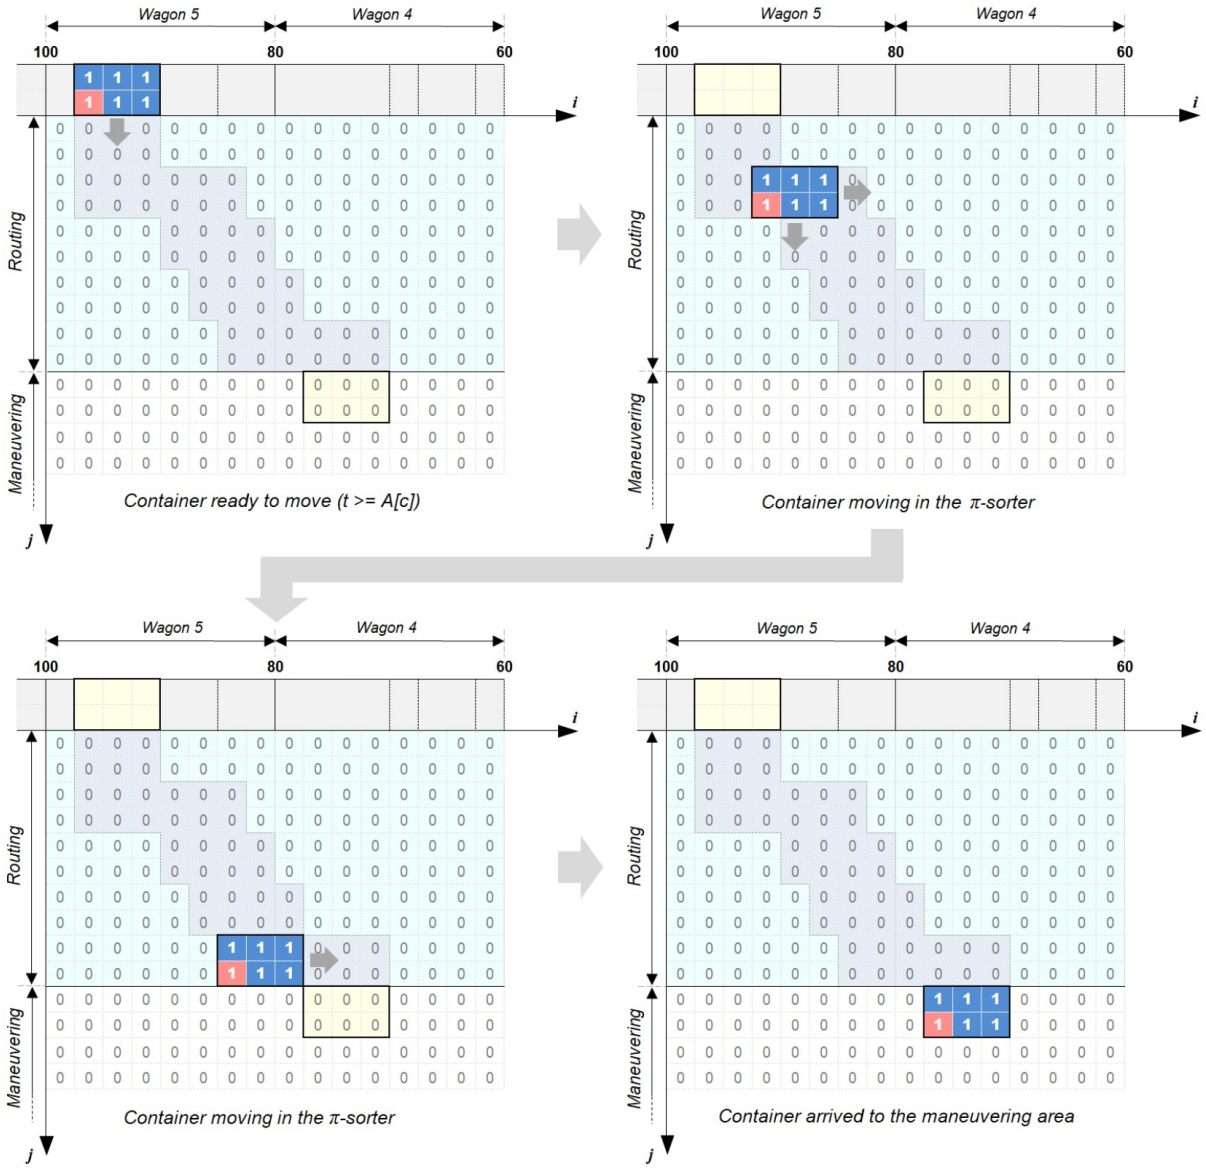


Figure 1 - An illustrative example of the routing of a container in the PI-sorter from an inbound train

Mathematically, the problem is formulated as a binary integer linear program that determines collision-free container trajectories across the sorting zone of the PI-hub. The objective function minimizes the total transfer distance, expressed as the cumulative number of conveying units traversed by all containers from their initial unloading positions to their assigned docking locations. Collision avoidance is enforced through spatial–temporal constraints that prevent simultaneous occupation of the same conveying unit by multiple containers, thereby guaranteeing operational safety. In addition, energy-related metrics are incorporated into the optimization framework, aligning the model with the sustainability objectives of Physical Internet systems.

The proposed model uses a binary decision variable to describe the position of each container over time and the utilization of conveying resources. In this abstract, only the objective function is presented the parameters are described in Table 1.

The objective of the proposed model is to minimize the total transfer distance of all containers from their initial inbound positions to their assigned outbound docking locations. This is achieved by minimizing the cumulative number of conveying units traversed during the transfer operations (Eq.1).

$$\sum_{i=0}^{n+1} \sum_{j=-1}^{m+2} \sum_{c=1}^{C} \sum_{t=1}^{T} p_{ij}^{ct} \tag{1}$$

Table 1 - An illustrative example of the routing of a container in the PI-sorter from an inbound train

Parameter	Description
p_{ij}^{ct}	This binary variable is equal to 1 if container *c* occupies position *(i,j)* at time *t*, where *(i,j)* represents the coordinates of the bottom-left corner of the container. Otherwise, the variable takes the value 0.
n	Number of conveying units along the horizontal axis (dock-facing dimension) of the PI-sorter, with the number of dock positions satisfying (d < n).
m	Number of conveying units along the vertical axis (train-to-dock dimension).
C	Total number of containers requiring routing.
T	Planning horizon expressed in discrete time periods.
w_c	Width of container (c), expressed as an integer multiple of the conveying unit width (a). All containers share a uniform length equivalent to two conveying units.

The proposed MIP model is implemented using IBM ILOG CPLEX 12.8 and evaluated on a set of ten generated test instances. These instances vary key parameters such as the number of containers, time horizon, container widths, arrival times, initial and final positions, and PI-sorter dimensions. Computational experiments demonstrate the robustness and effectiveness of the model across different PI-hub configurations and operational scenarios. Results show that all instances are solved with collision-free solutions, with solution times suitable for operational planning. Small to medium-sized instances exhibit particularly strong performance, while larger instances remain solvable despite increased computational effort (Table 2).

Table 2 - Numerical Results of CPLEX Implementation

#	*C*	w_c	*T*	*(n,m)*	*Objective*	*Time (s)*
1	2	[1,3]	25	(8,8)	29	2.89
2	2	[1,3]	25	(8,8)	28	1.13
3	2	[2,5]	25	(8,8)	30	2.17
4	3	[2,5]	25	(8,8)	39	3.80
5	3	[1,4]	25	(8,8)	45	52.47
6	3	[1,4]	25	(8,8)	46	36.36
7	3	[1,4]	25	(8,12)	48	27.33
8	3	[1,8]	25	(8,12)	53	529.69
9	4	[1,8]	35	(8,16)	76	1800.00*

** Pre-defined time limit*

Overall, this work contributes a unified, operationally robust optimization framework for PI-container routing that bridges methodological rigor with practical applicability. By accommodating multi-modal operations, enforcing collision avoidance, and enabling seamless adaptation to diverse hub architectures without structural modifications, the proposed model provides logistics operators and system designers with a versatile tool for efficient, safe, and sustainable PI-container management.

References

Ballot, E.: The physical internet p. 719 – 734 (2019). DOI 10.1007/978-3-319-92447-2 31

Ballot, E., Gobet, O., Montreuil, B.: Physical internet enabled open hub network design for distributed networked operations 402, 279 – 292 (2012).

Chargui, T., Bekrar, A., Reghioui, M., Trentesaux, D.: Tabu search robustness for cross-dock and pi-hub scheduling under possible internal transportation breakdowns 803, 295 – 307 (2019).

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

A Hybrid Hyper-Heuristic Reinforcement Learning Agent for Disruption-Resilient Synchromodal Transport

Thomas Kroiss[1], Felix Kamhuber[1], Sandra Stein[1] and Fazel Ansari[2]

1. Fraunhofer Austria Research GmbH, Center for Sustainable Production and Logistics, Vienna, Austria
2. TU Wien, Chair of Production and Maintenance Management, Vienna, Austria

Corresponding author: thomas.kroiss@fraunhofer.at

Keywords: *Include 3 to 10 keywords, including Physical Internet. Reinforcement Learning, Synchromodal Transport Re-planning, Agent-based Simulation, Hyper-Heuristics, Disruptions, Vehicle Routing Problem, Physical Internet*

Conference Domain Fitness: *Compose a paragraph positioning your contribution within the IPIC 2026 contribution domain as outlined in the call for contributions, using as directly as possible the provided text and cutting, assembling, and extending it as best fitting. Maximum of 10 lines.*

This contribution addresses the Physical Internet vision of open, interconnected, and resilient logistics networks by proposing an adaptive decision-making framework for synchromodal transport re-planning under uncertainty. Positioned within the IPIC domain of PI Nodes, PI Networks, and Systems of Logistics Networks, the paper demonstrates how hybrid reinforcement learning can enhance real-time orchestration across multimodal transport hubs and networked logistics infrastructures. By focusing on disruption management, operational resilience, and scalable decision support, the work aligns with the conference's emphasis on intelligent, adaptive, and sustainable logistics systems enabling the deployment and evolution of Physical Internet principles in real-world transport networks.

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☒ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☒ *Transportation Equipment,* ☒ *PI Networks,* ☒ *System of Logistics Networks,* ☒ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

Research Contribution Abstract

The implementation of synchromodal transport is one of the key enablers for the realization of the Physical Internet (PI). Despite significant benefits, synchromodal transport introduces a high-dimensional, stochastic control problem as it is confronted with e.g. time-varying demand, uncertain travel times, shared capacities across modes, multiple decision horizons, conflicting objectives (cost, CO_2, service level). In this context, classical optimization approaches often struggle, as large-scale models can become computationally intractable, repeated re-optimization may be too slow for real-time control, and uncertainty is difficult to represent explicitly within deterministic formulations. Reviewing the body of knowledge reveals that reinforcement learning (RL) is a promising approach

for supporting dynamic decision-making in synchromodal transport. Its practical integration, however, remains challenged inter alia due to scalability, robustness under disruption, and acceptance by operational stakeholders. Following this line of research, this paper focuses on an in-depth investigation of a hybrid reinforcement learning agent based on hyper-heuristics.

The proposed hybrid-RL approach combines learning-based adaptation with a portfolio of domain-specific heuristics, allowing the agent to dynamically select re-planning strategies rather than learning low-level control policies from scratch. This design targets operational realism, computational efficiency, and transparency. The agent operates within an agent-based simulation modeling a European synchromodal transport network, which integrates logistics hubs across multiple transport modes under time-dependent constraints.

In terms of evaluating resilience, our agent-based simulation approach introduces stochastic disruptive events such as transport delays, hub capacity reductions, and route unavailability. Results indicate that the hyper-heuristic RL agent significantly improves recovery speed, service reliability, and overall network performance compared to static or single-heuristic approaches. The paper concludes by discussing limitations and outlining future research directions, including different heuristic portfolios and possible exploration of alternative reward formulations and decision structures inspired by PI principles.

Cyber-Physical Internet-Enabled Logistics and Assembly Synchronization for Modular Integrated Construction in the Greater Bay Area

Wanyi Song, Zhiheng Zhao and George Q.Huang

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026
Bordeaux, FRANCE

The Hong Kong Polytechnic University
香港理工大學

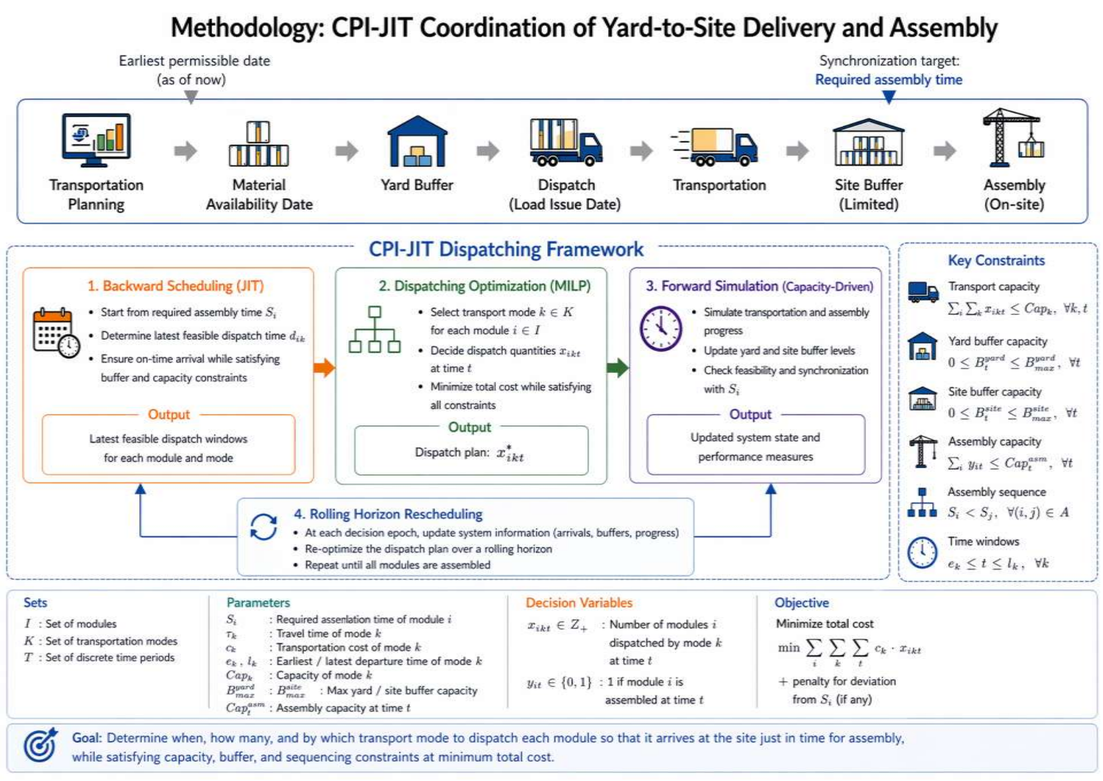
Introduction

- MiC projects in the Greater Bay Area require JIT delivery to space-constrained construction sites in Hong Kong.
- Existing logistics approaches optimize transportation but fail to align delivery with on-site assembly processes.
- This study proposes a CPI-enabled framework to achieve synchronized JIT delivery under buffer and assembly constraints.

Objectives

- To develop a CPI-enabled framework for coordinating yard-to-site delivery of MiC modules.
- To formulate a JIT dispatching model that integrates transportation decisions with buffer and assembly constraints.
- To design a rolling-horizon optimization approach for achieving synchronized delivery and installation.

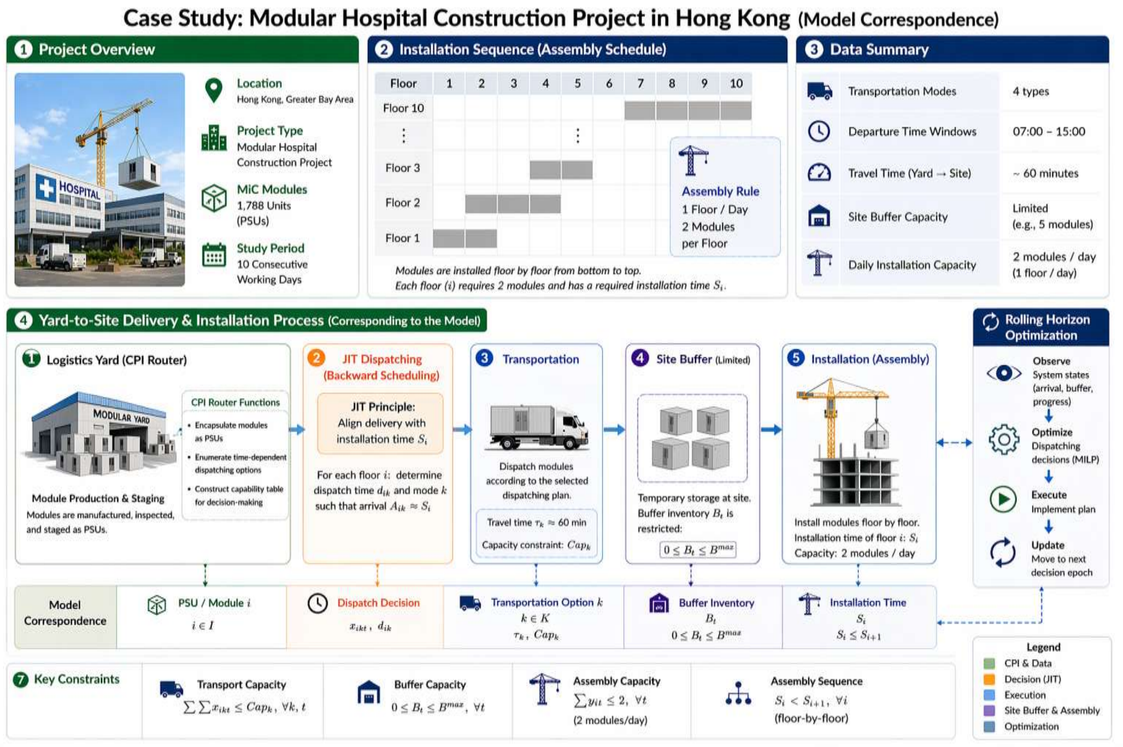
Methodology



- We develop a CPI-enabled framework for coordinating MiC delivery and assembly. Modules are dispatched from the yard (CPI router) using JIT scheduling to match installation time, and decisions are optimized under capacity and buffer constraints with rolling updates.

Case Study

- This case study applies a CPI-enabled framework to a modular hospital project in Hong Kong, where MiC modules are delivered to a constrained site. Due to limited buffer and strict assembly sequences, the problem is modeled as a JIT dispatching problem to synchronize delivery and assembly.



Contributions

- Proposes a CPI-enabled framework to coordinate yard-to-site delivery of MiC modules under construction constraints.
- Formulates a JIT dispatching model that integrates transportation decisions with buffer limitations and assembly scheduling.
- Develops a rolling-horizon optimization approach for dynamic and synchronized dispatching decisions.

Future Work

- Extend the framework to multi-stage logistics systems, including upstream cross-border transportation.
- Incorporate real-time data and uncertainty to enhance robustness and adaptability of dispatching decisions.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE
amU Aix Marseille Université

An urban circular delivery service for the implementation of proximity kilometer logistics

Fouzia Ounnar[1] and Patrick Pujo[1, 2]
1. CRET LOG – Aix-Marseille Université, Aix-en-Provence, France
2. UTOPII - UAR 2049 CNRS Ingénierie, Paris, France
Corresponding author: fouzia.ounnar@univ-amu.fr

Abstract: *The Physical Internet (PI) is a transformative logistics concept that aligns with the goals of the circular economy by increasing efficiency, resource sharing, and sustainability. Built on three pillars, Physical Web, Information Web, and Ecosystem Business Model, PI promotes open, standardized, and interconnected logistics systems that reduce emissions, empty mileage, and costs. Standardized modular containers and advanced data exchange support both forward and reverse logistics, enabling large-scale circular flows and greener business ecosystems. In urban contexts, PI operates through networks of urban and peri-urban hubs connected by low-impact mobility solutions like electric vehicles and cycle logistics. Real time data systems coordinate demand, capacity, and routing across these networks. The circular economy reshapes logistics flows through feedback loops that preserve product value and minimize resource extraction, while also introducing new logistical challenges such as unpredictable demand and irregular reverse flows. To manage this complexity, a holonic, self-organizing control approach, based on autonomous, cooperative entities, is proposed. This enables adaptive, decentralized management of urban circular supply chains. Ultimately, the Urban Physical Internet aims to mutualize assets and data, increase resilience, improve energy performance, and enhance overall efficiency in urban freight systems.*

Keywords: *Circular Economy, Physical Internet, Urban Logistics System, Self Organized Holonic Control, Circular Supply Chain Pyramid.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☒ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

1 Introduction

The Physical Internet (PI) (Montreuil, 2011) contributes to the circular economy by enhancing resource utilization efficiency, enabling intensive sharing of logistics infrastructures, significantly reducing the environmental impacts of transportation, and fostering the development of collaborative business ecosystems oriented toward sustainability. This paper contributes to ongoing research on the Physical Internet as a framework for urban logistics in the transition toward a circular economy. It examines how container standardization, resource pooling and networked configurations of urban and peri-urban hubs can support the deployment of circular supply chains in cities. The analysis highlights the specific uncertainties associated with circular flows, the increasing complexity of controlling urban logistics systems, and the

need for agile, decentralized and holonic control architectures. The paper is organized into four main sections: (i) the conceptual foundations and operational principles of the Physical Internet, (ii) its deployment in urban contexts, (iii) the reconfiguration of logistics flows induced by the Circular Supply Chain Pyramid and (iv) the resulting uncertainty, complexity and control issues, addressed through a holonic and self-organized Digital Twin of urban circular delivery services.

2 Physical Internet: a lever for urban circular economy

2.1 Physical Internet: conceptual foundations and underlying purpose

The Physical Internet represents a major conceptual and operational opportunity for promoting the circular economy (Geisendorf and Pietrulla, 2018). Its relevance stems from its capacity to enhance the efficiency and sharing of resources, substantially mitigate environmental impacts, and foster the emergence of new sustainable business ecosystems (Rajahonka et al., 2019). The Physical Internet enhances logistical efficiency through standardized modular containers, higher load factors and integrated open transport networks, thereby reducing transport distances, empty mileage and overall resource consumption. It enables intensive sharing of storage, hubs and transport capacity across firms in an open logistics web, turning underused assets into common infrastructure that supports both forward and reverse flows required by circular economy systems. These features lead to substantial reductions in CO_2 emissions and road freight externalities, while fostering collaborative ecosystems with new business models, measurable performance standards and green service offerings that underpin more sustainable and circular supply chains.

2.2 Structuration of the Physical Internet and benefits

The PI framework is structured around three fundamental pillars: the “Physical Web”, the “Information Web” and the “Ecosystem Business Model”. The pooling of flows and resources (vehicles, hubs, data) is a key issue for urban logistics because it allows to consolidate shipments from multiple actors, thereby reducing the number of tours, cutting congestion and emissions, and optimizing scarce urban space. Inspired by the architecture of the digital Internet, the Physical Internet introduces standardized modular containers, collaborative routing hubs and shared communication protocols that enable different logistics networks to interconnect and share capacity. By making infrastructures and information systems interoperable, it unlocks large-scale pooling, allowing cities to handle the growth of e-commerce and last-mile deliveries without a proportional increase in noise, pollution and traffic.

2.3 Operational principles of the Physical Internet

The standardization and modularization of containers and logistics units facilitate consolidation, multimodality, and resource reuse, while reducing empty mileage and redundant infrastructures. For an equivalent service level, the PI framework delivers significant performance gains in terms of GHG emissions, travel distances, and costs, thereby aligning with and operationalizing the sustainability objectives of circular economy. Furthermore, the combination of open networks, advanced traceability mechanisms, and innovative business models promotes the establishment of return loops and the emergence of “green” logistics ecosystems that integrate environmental and social dimensions into holistic performance metrics. Standardizing container dimensions and interfaces simplifies loading, unloading, and cross-docking operations, thereby reducing break of bulks and processing times at urban logistics platforms. Modularity (small building-block units that can be combined into larger

loads) enables easy grouping of heterogeneous flows toward consolidation hubs and their fine reconfiguration for last-mile delivery. This facilitates the consolidation and massification of flows while maintaining a high level of flexibility. By providing a common standard for multiple operators, these modular containers serve as a foundation for integrated control of delivery tours, vehicle loading, and dock capacity, a necessary condition for the systemic optimization of urban freight flows.

3 Deployment of the Physical Internet within urban environments

3.1 Architecture of urban logistics

In the specific context of urban deployment, PI builds upon a networked configuration of urban and peri-urban hubs (comprising urban distribution centers, micro-hubs, and PI platforms) interconnected through low impact mobility solutions such as electric vehicles and cycle-logistics, and equipped with modular containers of suitable sizes (Matusiewicz, 2024). This urban logistics system enables systematic transport consolidation through container grouping, rapid cross-docking, and multimodality between motorized and non-motorized modes within urban π-hubs. In peri-urban areas, large regional or national logistics platforms act as high-capacity hubs, often multimodal, that concentrate flows from long-distance corridors and redistribute them toward metropolitan areas (Pan *et al.*, 2017). At the urban fringe, city distribution centers and logistics hotels perform cross-docking and consolidation between multiple carriers and shippers before goods enter dense inner-city areas. Within city centers, micro-hubs and proximity logistics spaces (Paché, 2024) ensure pre-haulage in small batches and redistribution by cargo bikes or low-emission vehicles within a dense network of short-range break of bulk points. The whole forms an interconnected hub-and-spoke network, in which each hub type (peri-urban consolidation hub, urban distribution center, micro-hub) fulfills a specific function (massification, cross-docking, proximity), contributing to the systemic optimization of urban freight flows (Taniguchi *et al.*, 2016). (Papaioannou *et al*., 2023) advocate the use of low environmental impact vehicles, particularly electric cargo bikes, citing environmental, health, operational and urban arguments. The authors emphasize that increasing urban density and car dependence exacerbate congestion, pollutant emissions and health risks, thereby making more sustainable last-mile transport systems a necessity. In this context, electric cargo bikes are presented as a credible alternative to conventional vehicles, able to replace up to 10% of small trucks without compromising operational efficiency while reducing CO_2 emissions by about 73%, with an estimated 40 - 50% of urban delivery tours being theoretically substitutable by this mode. These climate benefits are complemented by economic and spatial advantages, including lower investment and operating costs and high maneuverability in dense urban fabrics (historic centers, narrow streets, pedestrian areas) where access for motor vehicles is restricted. Finally, advances in battery technology enable reliable operation under previously challenging conditions (gradients, headwinds), reinforcing the view that such low-impact vehicles constitute a central lever for the decarbonization of urban logistics. With regard to economic aspects, (Robichet *et al*., 2022) demonstrates that, based on operational data from a major Parisian logistics operator, the use of cargo bikes and micro-hubs for proximity kilometers (first and last mile) allows to reduce emissions whilst remaining economically competitive (Tongiani *et al.*, 2022). This demonstrates the competitiveness of the cycle logistics model in areas with a high density of deliveries.

3.2 Operation and key features of urban freight flows

The management of these urban flows requires real time control systems able of aggregating and processing real time data in order to dynamically coordinate demand, capacities, and

routing decisions. In addition to the use of modular containers suited to urban environments, an urban logistics system must be organized to facilitate the consolidation and grouping of freight flows. This requires break of bulk points that enable multimodality between motorized and non-motorized modes of distribution: urban consolidation centers. (Allen *et al*., 2012) systematically describes the concept of urban distribution center, flow organization schemes, economic models, as well as the observed effects on traffic, emissions and logistics performance across 117 case studies in Europe. At the same time, (Hübner *et al.*, 2016) provides a detailed analysis of the growing demands of customers regarding the speed and reliability of delivery. The rise of e-commerce has been accompanied by a continuous shortening of the delivery times deemed "acceptable" by customers, who have shifted from expecting delivery within a few days to expecting delivery within 24 hours, or even on the same day, whilst remaining very price sensitive. This research describes this phenomenon as a widespread adoption of the "everything, right now" standard, driven by major online actors who have structured their logistics networks to offer very short delivery times and have, in effect, shifted the expectations of all consumers, with direct consequences for the intensity of logistics flows and the pressure on the last mile. This necessitates real time control of urban logistics flows, combining the coordination of demand, capacity and routes. To all this, we shall see that new constraints are superimposed due to the implementation of circularity.

4 Circular economy and the restructuring of logistics flows

4.1 "Circular Supply Chain Pyramid" Model

The circular economy profoundly redefines the logistics flows that ensure the operation of the ten circularity levers (Kirchherr *et al.*, 2017), while integrating three interdependent circular feedback loops (Pujo *et al.*, 2024). The resulting framework, referred to as the "Circular Supply Chain Pyramid" (figure 1), primarily aims to preserve the use value of products and components, reduce dependence on natural resources, and lower the environmental impacts associated with resource extraction and transformation. Increasing the intensity of use at the upper levels of the pyramid necessitates stronger consumer engagement, as the deployment of certain circularity levers generates new logistics flows while amplifying others. This model aims to preserve use value as far as possible by extending product lifespan, increasing utilization rates and prioritizing reuse over premature replacement. In doing so, it seeks to reduce dependence on primary resources by limiting the need for new material extraction and processing through maintenance, repair, refurbishment and sharing. Ultimately, this framework contributes to lowering the environmental impacts associated with production–consumption cycles (emissions, energy use, waste generation) by aligning economic strategies with trajectories of material and energy sobriety.

Consumer interactions with the urban logistics system are steadily increasing in both frequency and intensity, driven by the widespread diffusion of e-commerce and by the rise of functional economy models (subscriptions, on-demand services, access-based and sharing schemes). Individuals no longer act solely as one-off buyers of goods; they become users of logistics services, actively selecting among service configurations (express vs. standard delivery, subscription-based delivery, pickup points, lockers, click-and-collect) and optimizing the frequency and mode of access to the desired function rather than to the mere ownership of products. This heightened involvement, mediated by platforms that monetize recurrent access to logistics services (unlimited delivery passes, premium membership, rental or sharing services), is reshaping flows, time patterns and organizational forms in urban logistics, which must now closely incorporate users' service preferences, temporal constraints and economic trade-offs into planning and operational decisions.

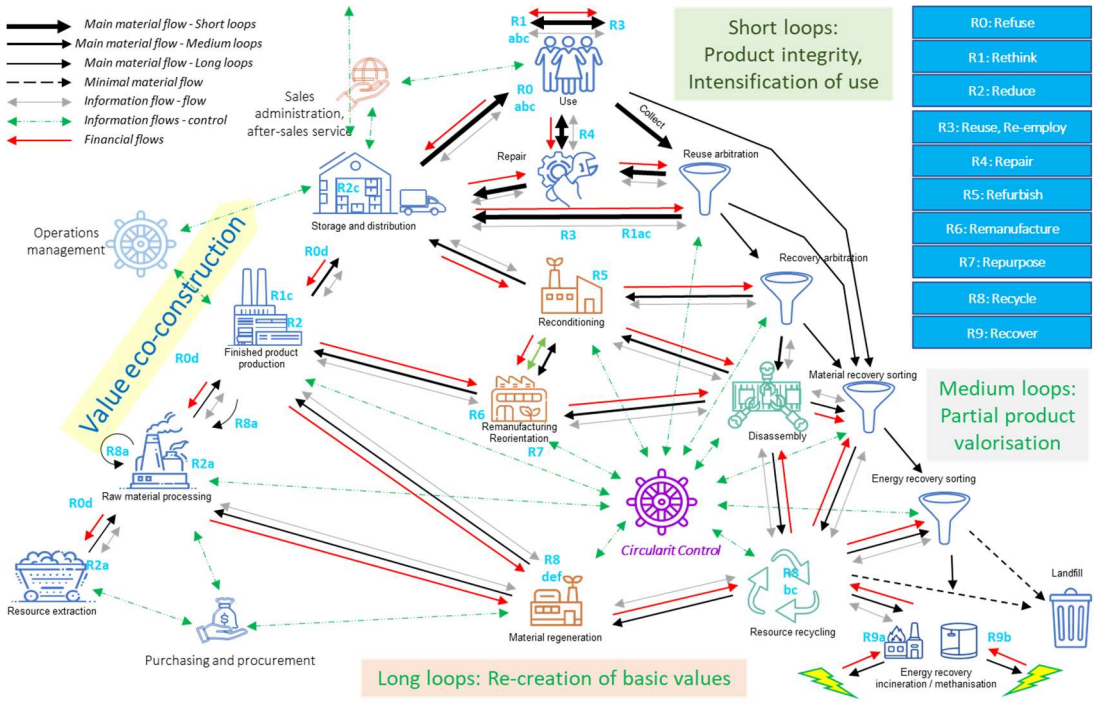


Figure 1: Circular Supply Chain Pyramid

4.2 Circular material flows and urban logistics systems

The transition towards a circular economy also introduces distinctive features in the approach to urban logistics. In a linear economy, urban logistics is generally considered from the perspective of last-mile delivery; in contrast, a circular model also requires the explicit integration of the first mile. This local logistics system (encompassing both the first and last mile) is characterized by greater dynamics in circular flows compared to traditional linear systems (Ounnar *et al.*, 2025a). In this context, the uncertainty inherent in consumption is amplified by the uncertainty regarding the volumes of flows from circular supply chains, all of which is made unpredictable by the geographical dispersion of users and their propensity to generate more or less waste and/or to consume responsibly.

In general, interactions with urban logistics users are concentrated mainly in the upper loops of the CSC Pyramid, that is, at the level where the functional economy is located and where strategic trade-offs and service design choices are formalized. At these higher levels, users are not merely passive recipients, but stakeholders whose expectations, preferences and feedback help shape the system's development trajectory. This concentration of interactions within the higher-level loops indicates that the user's voice is considered primarily when defining overall strategic directions, service delivery configurations and trade-offs between performance, sustainability and service quality.

5 Uncertainties and complexity of circular control in urban areas

This suggests that the operation of urban CSC is subject to a greater degree of uncertainty.

5.1 Structural uncertainties inherent to the circular economy

An Urban Physical Internet therefore confront two major categories of uncertainty. First, demand becomes more difficult to forecast (no longer governed by standardized consumption patterns but rather by the interplay among new product competition, the diffusion of sharing practices, the growth of repair and reuse services, the evolution of product durability (and

consequently reduced renewal needs), and shifting consumer expectations shaped by socio-economic and marketing factors. Second, supply becomes erratic and heterogeneous, as it largely depends on reverse circular flows that are irregular, spatially dispersed, and diversified by use effects, aging, and deterioration, organized into nested loops of value conservation across different spatial and temporal scales (short, medium, and long). This second dimension of uncertainty intensifies the complexity of circular flows' control, thereby calling for agile, decentralized logistics organization able to adjust in real time to fluctuations in demand and return flows. In this regard, a holonic, self-organized, and multicriteria approach has been proposed to control circular supply chains (Ounnar *et al.*, 2025b). It relies on autonomous and cooperative entities that interact in a coordinated manner to continuously optimize the overall performance of inter-organizational circular flow control systems.

5.2 Organizational implications for logistics systems

Urban logistics systems are characterized by a high degree of heterogeneity among actors (carriers, retailers, digital platforms, public authorities), strong variability in demand (seasonality, daily peaks, promotional events, weather-related disruptions) and multi-level, partner-based network structures, which makes purely centralized, hierarchical control schemes largely inadequate and complex. In this context, the need for agile and decentralized control arises from the superior ability of such architectures to absorb local disturbances (delays, congestion, breakdowns, demand shifts) through self-organization, cooperation and distributed coordination among autonomous entities, rather than through slow and cumbersome global re-planning. Research on holonic and isoarchic control architectures shows that, by relocating part of the decision-making capability to the nodes themselves (resources, logistics hubs, vehicles), it becomes possible to reconcile local adaptability, global robustness and dynamic optimization, notably via distributed negotiation mechanisms, local task generation and on-line reconfiguration of client-supplier relations. Within this framework, continuous adaptability of the control system to variations in logistics flows (volumes, origin-destination patterns, time constraints) is a key property: it enables ongoing adjustment of routes, resource assignments and service priorities without locking the system into a rigid configuration. A static control scheme designed for an "averaged" or assumed stable environment tends to amplify desynchronization, local overload and inefficiency (bullwhip effects, break of bulk, under-utilization of certain resources), whereas agile and decentralized control, grounded in local perception-decision loops and horizontal coordination mechanisms, enhances resilience, service continuity and the ability of the urban logistics network to reconfigure itself in near real time.

6 Towards a self-organized Urban Circular Delivery Service

Several scientific studies advocate the same approach to agile, decentralized and adaptive control for logistics systems, including urban ones, due to their volatile and uncertain nature.

6.1 A holonic and self-organized approach

Distributed control and self-organization (holonic, multi-agent) are particularly well suited to systems exposed to high demand variability and frequent disturbances. He highlights the central role of continuous adaptative capability in maintaining overall system performance. (Leitão *et al.*, 2012) show, in the field of reconfigurable manufacturing, that decentralized, multi-agent architectures inspired by biological systems enable online adaptation and rapid reconfiguration in response to disturbances, thereby improving agility and robustness compared to conventional centralized control. (Leitão *et al.*, 2016) emphasize the importance of distributed control structures that can react locally to changes in the environment and in flow patterns, and

dynamically reconfigure themselves, which supports the need for agile, decentralized control in complex and evolving systems such as those required in urban logistics. Finally, (Perboli and Rosano, 2019) confirm that cooperative and flexible organizational forms, based on distributed decision-making among actors, are better able to cope with uncertainty in urban demand and variations in logistics flows.

We propose to adapt this approach to the control of an Urban Physical Internet, as a means of implementing an urban logistics delivery service aligned with the principles of an "Ecosystem Business Model", rooted in resource pooling, inter-organizational cooperation, and collective value creation. The self-organized approach of Physical Internet's control is particularly well-suited to supporting the circular economy, as it profoundly transforms the organization of physical, informational, and economic flows.

6.2 Implementation of an Urban Circular Delivery Service

The proposed control system is based on an isoarchic self-organized holonic architecture as defined by (Pujo, 2012; Ounnar and Pujo, 2005)). A holon is an autonomous entity that is a "whole that is part of a whole": this enables a recursive and hierarchical approach to a system. It consists of a physical component, the M_holon, and an immaterial component, the I_holon, which itself contains an information system managing data and variables, an internal state management system and a decision-making system. The latter enables real time management decisions to be made based on multicriteria analyses carried out using data from the field or through collaboration with other holons. This approach is enhanced, on the one hand, by the addition of a Geographic holon, which provides access to a Geographic Information System describing the characteristics of the urban area, and, on the other hand, by a Digital Twin, which enables simulation-based projections into the near future, thereby allowing certain events to be anticipated and better decisions to be made (Kritzinger *et al.,* 2018). In all cases, a holon's Digital Twin possesses exactly the same properties as the holon and shares exactly the same algorithms that enable self-organized control. The only difference lies in the source of the data feeding these algorithms: in an operational holonic system, a large proportion of the data comes from the M_holon, whereas in a Digital Twin during the study and design phase, it comes from a simulated M_holon, and in a Digital Twin during the operational phase, it comes from the M_holon but is used by the I_holon for real time operational control and the prospective analysis of scenarios over a given time horizon. The Digital Twin model of a Holon x can be expressed according to the following definition:

$$M_{DTx} = (M_Hx(PE); I_Hx(VE; Ss; DD; CN) \quad (1)$$

Where: PE refers to the physical entity and constitutes the physical part of the holon, VE is the virtual equipment, Ss stands for services for PE and VE, DD refers to Digital Twin data, and CN is the connection between PE, VE, Ss and DD, using the work of (Tao *et al.*, 2018).

The model of an Urban Circular Delivery Service is in fact made up of various holarchies (sets of holons), such as:

- a set of π_containers $\{C_PH_i\}$ to be transported within the city or to a suburban area, containing products involved in a given circular flow, described by an Origin and a Destination;
- two distinct sets of resources: a set of urban logistics hubs $\{H_RH_j\}$ and a set of vehicles $\{V_RH_k\}$ (cargo bikes, three-wheelers, light-duty vehicles (LDVs)) belonging to different logistics service providers;
- a set of circular flows $\{F_OH_l\}$ relating to the aforementioned C_PH_i;

- a GIS interacting with all holons located or moving within the relevant geographical area.

The Digital Twin model of an Urban Circular Delivery Service (UCDS) results from the combination of the Digital Twin models of these holons. The class diagram of this UCDS Digital Twin is shown in Figure 2.

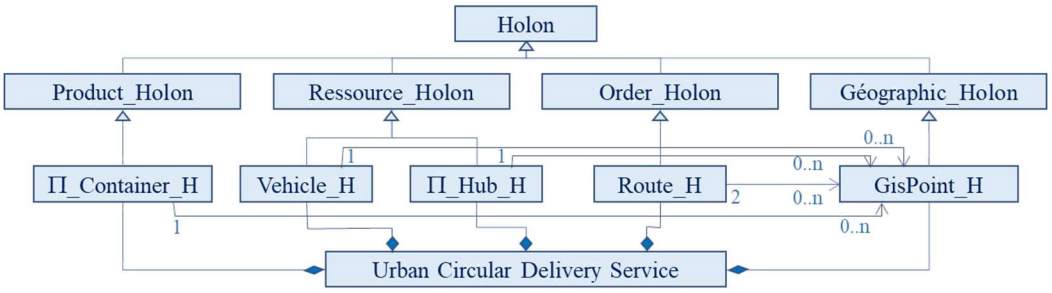


Fig. 2. Class diagram of the UCDS Digital Twin.

Each holon is equipped with self-management capabilities as described by state graphs in (Dubromelle *et al.*, 2012). When faced with multiple options, each holon utilizes a multicriteria decision support service based on the Analytic Network Process (Saaty, 2004; Dubromelle *et al.*, 2012). These capabilities are inherited from a Holon entity and enable each holon to participate contextually in decision-making. The same applies to interactions between holons, and particularly to decision-making involving an auction-based decision mechanism (Ounnar and Pujo, 2005; Pujo, 2012). Thus, any holon may issue a Call For Proposal to a holarchy with a view to take a decision on a given task. Each holon in this holarchy analyses this request and attempts to provide the best response, including, if necessary, launching a multicriteria analysis to comparatively evaluate different alternative solutions. If the response provided by this holon is better than the best solution already published (or if it is the first), this holon takes the lead by publishing its Proposal. Upon expiry of the deadline, the holon holding the lead wins the auction.

6.3 Illustration of self-organized Urban Circular Delivery Service

The study aims to assess the progress made towards greater speed and circularity in the home delivery system of a fast-food chain located in the very heart of a megacity. The demonstrator developed (figure 3) compares the performance of a conventional control system for a fleet of electric delivery tricycles (each statically assigned to two of the various restaurant outlets and delivering only its own orders) with a self-organizing control system for a shared delivery fleet, where assignments are dynamic, based on changes in delivery demand and various technical, economic and urban criteria. The ANP multicriteria analysis used to take decisions on the allocation of delivery tasks is based on a number of indicators, such as quantities (total number of orders, number of pending orders), times (average time per order, average trip time), utilization rates (use of delivery tricycles), status (vehicle energy and battery levels), the urgency of the task (delivery of hot/cold goods (urgent), or recovery of dirty dishes or organic waste collection (less urgent)) and customer satisfaction rates.
The results demonstrate the feasibility of the operation and an improvement in performance (meeting delivery deadlines, better utilization of the fleet) thanks to the dynamic, real time re-routing of the delivery tricycles. Beyond this performance, a shift in the geographical center of gravity of the location of all delivery tricycles was observed, in line with changes in the center of gravity of all the households from which orders were placed and where dirty dishes must be collected after use.



Fig. 3. Partial snapshot of the UCDS Digital Twin.

7 Conclusion

The deployment of the Physical Internet (PI) in urban environments represents a systemic approach aimed at optimizing the management of urban logistics flows through the pooling of material and informational resources across stakeholders. This initiative seeks to minimize operational inefficiencies and negative externalities associated with traditional logistics systems. The concept relies on the use of standardized physical containers, interconnected through networks of urban hubs and governed by dynamic flow routing. Combined with a self-organized holonic control framework, this architecture enables real time responsiveness to the inherent uncertainties of urban logistics, such as volume variability, service diversification, circular economy requirements, and temporal constraints. The shared use of urban logistics infrastructures and capacities promotes a more rational and sustainable utilization of these resources, while providing a coherent techno-organizational architecture suited to contemporary challenges of resilience and sustainability. Expected outcomes include enhanced network robustness, improved ecological performance, reduced greenhouse gas emissions, and overall optimization of urban freight transport efficiency. In the long term, this approach would strengthen the sustainability of urban freight by mitigating congestion, logistical costs, and environmental impacts, without compromising the level of service delivered to users.

References

- Allen J., M. Browne, A. Woodburn, J. Leonardi (2012): *The role of urban consolidation centres in sustainable freight transport*. Transport reviews, v32, no4, 473-490.
- Dubromelle Y., T. Louati, F. Ounnar, P. Pujo, (2010): AHP/ANP a Decision Making Service in PROSIS Model. **Proceedings of IMS10 10th IFAC Workshop on Intelligent Manufacturing Systems**, Lisbonne.
- Dubromelle Y., F. Ounnar, P. Pujo (2012): *Service Oriented Architecture for Holonic Isoarchic and Multicriteria Control*. In Service Orientation in Holonic and Multi-Agent Manufacturing Control, vol. 402, ed. Borangiu, T., Thomas, A., Trentesaux, D. Studies in Computational Intelligence. Springer
- Geisendor S., F. Pietrulla (2018): *The circular economy and circular economic concepts—a literature analysis and redefinition*. Thunderbird Int. Business Review, v60, no5, 771-782.

- Hübner A., H. Kuhn, J. Wollenburg (2016): *Last mile fulfilment and distribution in omni-channel grocery retailing: A strategic planning framework.* Int. Journal of Retail & Distribution Management, v44, no3, 228–247.
- Kirchherr J., D. Reike, M. Hekkert (2017): *Conceptualizing the circular economy: An analysis of 114 definitions*. Resources, conservation and recycling, v127, 221-232.
- Kritzinger W., M. Karner, G. Traar, J. Henjes, W. Sihn. (2018): *Digital Twin in Manufacturing: A Categorical Literature Review and Classification.* IFAC-PapersOnLine, v51, no11, 1016–1022
- Leitão P., J. Barbosa, D. Trentesaux (2012): *Bio-inspired multi-agent systems for reconfigurable manufacturing systems*. Engineering Applications of Artificial Intelligence, v25, no5, 934–944.
- Leitão P., A. W. Colombo, S. Karnouskos (2016): *Industrial automation based on CPS technologies: Prototype implementations & challenges*. Computers in Industry, v81, 11-25.
- Matusiewicz M. (2024): *Framework for Physical Internet deployment in cities*. Urban, Planning and Transport Research, v12, no1, 2303341.
- Montreuil B. (2011): *Towards a Physical Internet: Meeting the Global Logistics Sustainability Grand Challenge*, Logistics Research, v3, no2-3, 71-87,
- Ounnar F., P. Pujo (2005): *Supplier evaluation process within a self-organized logistical network* Int. Journal of Logistics Management, v16, 159-172.
- Ounnar F., S. Boissée, F-M. Lambert, P. Pujo (2025)a. Pilotage multicritère des flux dans la Pyramide de la Supply Chain Circulaire. **Proceedings of CIEC2025 2ème Congrès Interdisciplinaire sur l'Economie Circulaire.** Amiens, France.
- Ounnar F., P. Pujo, F-M. Lambert (2025)b. Holonic, self-organized and multi-criteria management of the Circular Supply Chain. **Proceedings of SOHOMA2025 15th International Workshop on Service-Oriented, Holonic and Multi-Agent Manufacturing Systems for Industry of the Future**. Aix-en-Provence, France.
- Paché G. (2024). *Warehousing facilities in urban space: the challenges of "proximity logistics"*. Journal of Supply Chain Management: Research & Practice, 18 (1), pp.1-18. ⟨hal-04722385⟩
- Pan S., E. Ballot, G.Q. Huang, B. Montreuil (2017): *Physical Internet and interconnected logistics services: research and applications*. Int Journal of Production Research, v55, no9, 2603-2619.
- Papaioannou E., C. Iliopoulou, K. Kepaptsoglou (2023): *Last-mile logistics network design under e-cargo bikes*. Future Transportation, v3, no2, 403-416.
- Perboli G., M. Rosano (2019): *Parcel delivery in urban areas: A cooperative strategy to overcome the last mile problem*. European Journal of Operational Research, v273, no3, 1150–1162.
- Pujo P., F. Ounnar, A. Leblanc, F-M. Lambert (2024). Structuration de la Supply Chain Circulaire et positionnement vis-à-vis des filières REP. **Proceedings of CIEC 2024 1er Congrès Interdisciplinaire sur l'Economie Circulaire**, Montpellier, France.
- Rajahonka M., A. Bask, S. A. Yawar, M. Tinnilä (2019): *The physical internet as enabler of new business models enhancing greener transports and the circular economy*. Sustainable and efficient transport, 70-94.
- Robichet A., P. Nierat, F. Combes (2022): *First and last miles by cargo bikes: ecological commitment or economically feasible? The case of a parcel service company in Paris*. Transportation research record, v2676, no9, 269-278.
- Saaty T.-L. (2004): *Decision Making - the Analytic Hierarchy and Network Processes (AHP/ANP)*. Journal of Systems Science and Systems Engineering, v13, no1, 1-35
- Taniguchi E., R. G. Thompson, T. Yamada (2016): *New opportunities and challenges for city logistics*. Transportation research procedia, v12, 5-13.
- Tao F., M. Zhang, Y. Liu, A. Y. Nee (2018): *Digital twin driven prognostics and health management for complex equipment*. Cirp Annals, Manufacturing Technology, v67, no1, 169-172.
- Tongiani, M.-G., E. Ferrari and A. Alaimo (2022): *Fast digitalization in the pandemic era and the urgency to discover new business model opportunities for Italian small grocery retailers*. In L. Reis, L. Cagica Carvalho, C. Silveira and D. Xara-Brasil (Eds.), Digitalization as a driver for smart economy in the post-Covid-19 era. Hershey (PA): IGI Global, 55-77.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Multi-Stakeholder Governance Model for the Physical Internet

Attallah Qweider, Shafagh Alaei and Koen Mommens
Mobilise Mobility and Logistics Research Group, House of Sustainable Transitions (HOST),
Vrije Universiteit Brussels, Belgium
Corresponding author: attallah.qweider@vub.be

Keywords: *Governance; Trust and Collaboration; Physical Internet; MAMCA; Stakeholder Engagement.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☐ *PI Networks,* ☐ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☒ *Governance.*

Targeted Delivery Mode-s*:* ☐ *Paper,* ☒ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

1 Introduction

The Physical Internet (PI) constitutes a revolutionary change in traditional logistics, shifting from fragmented, bilateral, hierarchical coordination structures to open, interconnected networks inspired by the Digital Internet (DI). In this envisioned ecosystem, modular containers move across heterogeneous networks of hubs, transport modes, and handling facilities through dynamic routing protocols, without dependence on a single provider controlling end-to-end delivery (Spanuth et al., 2025). The potential benefits are substantial: improved asset utilization, reduced empty kilometers, and enhanced resilience through shared infrastructure.

However, realizing the PI vision requires more than technical and economic feasibility. It demands broad acceptance and active participation from a diverse set of industry and public actors. Successful implementation is expected to be achieved through collaborative design efforts involving both researchers and practitioners, and hinges on progressively validated solutions and standards that guide the implementation, adoption, and continuous evolution of the PI ecosystem (Ballot et al., 2021).

A central barrier to this participation is trust. Many logistics companies are unwilling to participate in such networks because they are concerned about exposing sensitive business data to competitors or to a central orchestrator (Cassan et al., 2023). Therefore, trust plays a key role in enabling collaboration within PI, as it involves interactions among diverse stakeholders such as shippers and logistics service providers. Addressing this requires adopting governance architecture that is decentralized yet coherent. Sun et al. (2024) argue that integrating trusted third-party mechanisms within a decentralized architecture can facilitate effective collaboration across diverse logistics actors.

Cassan et al. (2023) conceptualize PI as a decentralized architecture of interoperable software systems that enable collaboration and data exchange across different logistics nodes within the entire network. In line with this conceptualization, Heimburg et al. (2025) describe decentralized platform ecosystems as collectively governed by multiple participating

organizations rather than a single owner and may rely on blockchain technology or consortium-based decision-making, where no single authority controls rules and incentives.

As the implementation of PI networks progresses from conceptual design to operational deployment, the coordination and alignment of diverse actors' activities become increasingly critical to ensure seamless integration across jurisdictions and industry sectors. Yet governance considerations remain insufficiently explored, with existing contributions primarily focusing on the applicability of principles such as net neutrality and on the theoretical development of governance structure, leaving a research gap on how governance models can be practically designed to ensure trust among participants in PI-enabled platform ecosystem. We address this gap by developing a multi-stakeholder governance model using structured workshops and Multi-Actor Multi-Criteria Analysis (MAMCA). This participatory approach systematically integrates perspectives from different stakeholders, ensuring alignment with operational realities, incentive structures, and trust expectations in real-world networks

2 Methodology

This study adopts a multi-phase research design combining conceptual development with participatory evaluation in order to develop a governance model for the PI. It combines literature analysis, scenario development, stakeholder identification, and a participatory MAMCA approach. First, a targeted literature review is conducted to identify key governance dimensions relevant to PI, including coordination mechanisms, data governance, and decision structures. Second, based on these dimensions, a set of governance scenarios are developed to represent alternative institutional setups varying in centralization, regulation, interoperability, and data-sharing approaches. Third, relevant stakeholders are identified and grouped into five categories: technology providers, research and academia, logistics service providers, infrastructure operators, and public authorities, selected based on their expected operational and strategic roles within PI-enabled collaborative logistics ecosystems. Within each category, participants were selected as knowledgeable representatives of their respective stakeholder groups. Fourth, a single MAMCA workshop is conducted, in which representatives from each stakeholder group participate in a structured evaluation process to define governance objectives, derive evaluation criteria, and assess governance scenarios according to their operational priorities and trust requirements. This participatory process enables comparison of stakeholder perspectives and highlights areas of convergence and conflicting interests. Finally, results are synthesized into a validated governance prototype, defining roles, decision rights, and coordination mechanisms aligned with stakeholder input and theoretical insights.

3 Proposed Governance Scenarios for PI

This section presents five distinct governance scenarios varying in centralization, decision rights, coordination, and standardization approaches:
1. **Shared Governance Model:** This fully decentralized model relies on distributed leadership and consent-based decision-making, with authority shared among actors through consensus processes. While promoting inclusiveness, it may face challenges in scalability and coordination.

2. **Lead Organization Model:** This model centralizes governance in a single dominant actor responsible for coordination and key decisions. This model ensures strong control, but limited decentralization.

3. **Global Federation Model:** This model defines high-level interoperability standards globally, while regional hubs adapt and implement these standards according to local or sectoral requirements.

4. **PI Community Model:** This model is inspired by the bottom-up standardization process of the Digital Internet, in which technical standards and protocols emerge through open, collaborative, and transparent processes driven by the collective needs and contributions of the community.

5. **Common Process Space Model:** This hybrid model adopts a two-level governance structure. The common level ensures interoperability through identity management and service catalogues. The individual level consists of regional, national, or sector-specific networks that directly engage stakeholders.

4 Preliminary Results

This section discusses the preliminary results of the stakeholders' evaluations in the workshop. The results reveal significant variation in stakeholder preferences across the proposed governance scenarios (Figure 1). Overall, the results emphasize that no single governance model fully satisfies the priorities of all stakeholder groups, confirming the inherent trade-offs between decentralization, scalability, and decision-making legitimacy in designing a multi-stakeholder PI governance model. Across all stakeholder categories, Common Process Space and Global Federation governance models outperform purely centralized or decentralized approaches. The Common Process Space model ranks highest overall, particularly among technology providers and academia, while also obtaining relatively high support from public authorities and LSPs. In contrast, logistics service providers show a preference for more centralized approaches emphasizing trust and clear decision-making. The Shared Governance Model ranks lowest overall, mainly due to concerns related to scalability and coordination complexity.

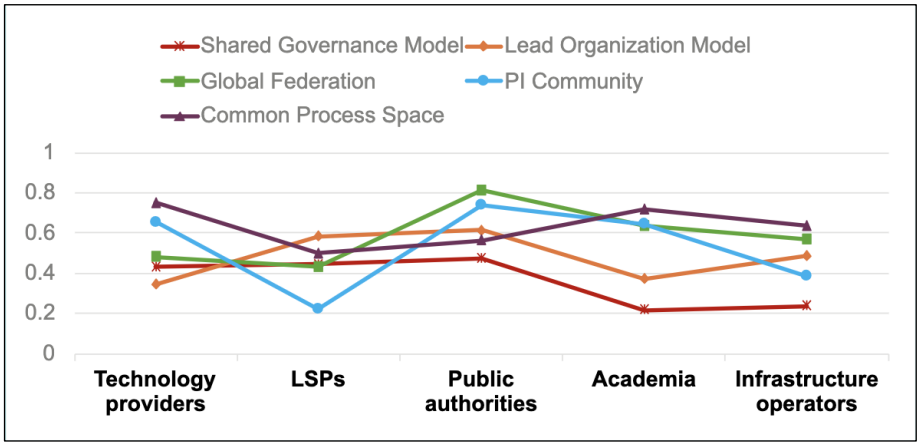


Figure 1: Overall evaluation of scenarios by stakeholder groups

5 Conclusion

This study sets out to address a critical yet underexplored challenge in the development of PI: how to design governance structures that enable trust, coordination, and scalability within PI ecosystem. We identified five stakeholder groups; and by combining scenario development with participatory evaluation through MAMCA, the results demonstrate that: Stakeholder

preferences vary significantly depending on their roles, priorities and trust requirements. Each governance scenario presents distinct strengths and trade-offs. These findings suggest that pure governance archetypes are insufficient, and that hybrid models are necessary to balance flexibility, control, and trust, and are best suited for governing PI, as they reconcile diverse stakeholder needs.

Future work will focus on completing stakeholder evaluations, formalizing governance mechanisms, and validating the model in real-world pilot settings.

References

- *Spanuth, M., Franklin, R., Ludwig, A., & Niemann, G. (2025). Governing the Physical Internet: Insights from Internet Governance and Future Research Directions. Georgia Tech Digital Repository. https://doi.org/10.35090/gatech/13225*
- *Ballot, E., Montreuil, B., & Zacharia, Z. G. (2021). Physical Internet: First results and next challenges. Journal of Business Logistics, 42(1), 101–107. https://doi.org/10.1111/jbl.12268*
- *Cassan, C., Michiels, P., Sun, S., Lemos, V., Van Bever, D., Cant, A., Fink, J., & Macharis, C. (2023). Data Sharing in the Physical Internet: A Capability-Based Approach for Trustless Logistic Networks. In IPIC 2023 Conference papers and posters Contributions Proceedings (Vol. 9, pp. 161-173). Alice.*
- *Sun, S., Michiels, P., Macharis, C., Cant, A., Dries, V. B., & Mommens, K. (2024). Unlocking the Potential of the Physical Internet: a Trust-enabling Decentralized Process Sharing Connector. VUBIR (Vrije Universiteit Brussel). https://doi.org/10.35090/gatech/11062*
- *Heimburg, V., Alakbarov, H., & Wiesche, M. (2025). Governance of decentralized platform ecosystems. In Lecture notes in information systems and organisation (pp. 97–112). https://doi.org/10.1007/978-3-031-80125-9_6*

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE

Toward a Precursor System for the Physical Internet: Minimal Objectives and Requirements for Sustained Interconnection

Malte Spanuth[1], Rod Franklin[1], Sahrish Jaleel Shaikh[2], Tiankuo Zhang[2], Paria Nourmuhammadi[2]

1. Kühne Logistics University, Hamburg, Germany
2. Georgia Institute of Technology, Atlanta, USA

Corresponding author: malte.spanuth@klu.org

Abstract: *The Physical Internet (PI) envisions open, protocol-based interconnection between independent logistics networks, yet its realization remains limited to isolated mechanisms and bounded pilot implementations. This study addresses the lack of understanding of how initial interconnection can become operationally viable and persist under real-world conditions. Drawing on Design Science Research, we synthesize insights from PI literature, pilot projects, and the historical evolution of the Digital Internet to identify minimal design objectives and requirements for a precursor system. The analysis reveals a structural gap between conceptual maturity and its practical realization and defines five core objectives enabling sustained interconnection. Consistent with the first two echelons of the echelon Design Science Research approach, this paper deliberately focuses on establishing the conceptual foundation and design objectives for a precursor system rather than its full implementation or empirical validation, which represent necessary next steps. The findings contribute by introducing the concept of a precursor system and providing a structured foundation for incremental PI implementation beyond project-based environments.*

Keywords: *Physical Internet, Interconnection, Design Science Research*

Physical Internet (PI) Roadmap Fitness*: ☒ Access and Adoption, ☒ Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐ Flash Video, ☐ In-Person presentation*

1 Introduction

Inspired by the success of the Digital Internet (DI), the concept of the Physical Internet (PI) envisions a transformation of logistics systems in which independent logistics networks interconnect and exchange shipments through standardized protocols. Accordingly, a growing body of research has focused on developing individual components required to enable the PI, including routing mechanisms, containerization concepts, and digital coordination platforms, as well as pilot implementations that test these components in real-world settings (e.g., ICONET, 2018; Landschützer et al., 2015; Sarraj et al., 2014; URBANE, 2022). However, the PI represents a highly complex socio-technical system whose realization is unlikely to occur all at once. In contrast to the evolutionary development observed in the DI, existing research often focuses on fully developed system visions without explicitly addressing how such systems can emerge incrementally.

The development of the Internet suggests a different perspective. The roots of today's DI trace back to ARPANET, a research network launched in 1969 and funded by the U.S. Defense Advanced Research Projects Agency, which initially connected four universities in the United States (Alexander et al., 2025; Leiner et al., 2009). Importantly, ARPANET did not constitute the DI itself but served as an experimental precursor that enabled researchers to test networking concepts and establish protocols and coordination structures. Rather than emerging fully

formed, the DI evolved through successive stages in which technical mechanisms and governance arrangements were gradually refined. Adopting a similar perspective, this study does not aim to design a fully realized PI, but instead focuses on an intermediate stage that enables initial interconnection between independent logistics networks.

Building on this perspective, the analysis reveals a critical gap in existing PI research. While prior work has explored individual mechanisms and envisioned comprehensive PI architectures, there is limited understanding of the conditions under which initial interconnection between independent logistics networks becomes feasible and sustainable in practice. In particular, it remains unclear which mechanisms are strictly necessary without assuming a fully developed PI system. To address this gap, this study adopts the notion of a precursor system, defined as a minimal[1] but sufficient configuration of mechanisms that enables sustained and viable interconnection between independent logistics networks:

RQ: What minimal design objectives and requirements enable sustained and viable interconnection between independent logistics networks under real-world operational conditions?

To answer this question, we follow the echelon Design Science Research (eDSR) approach by Tuunanen et al. (2024). We first analyze existing PI research and pilot implementations to identify structural limitations. We then derive design objectives and requirements grounded in PI research, pilot implementations, and the evolution of the DI, providing a foundation for a precursor system that enables incremental progression toward the PI. In doing so, this study contributes to PI research by introducing the concept of a precursor system, identifying a structural gap between conceptual research and pilot implementations, and deriving design objectives and requirements that define the minimal conditions for sustained interconnection.

2 Theoretical Background

To ensure the precursor system concept is not conflated with existing coordination paradigms such as Horizontal Logistics Collaboration (HLC), groupage networks, or digital freight platforms, this section identifies the structural characteristics that distinguish the PI and delineates the boundary conditions the precursor system must satisfy.

2.1 The Physical Internet as an Interoperability Paradigm

The PI can be understood as an interoperability paradigm that enables independent logistics networks to interconnect through shared protocols (Ballot et al., 2014; Montreuil, 2011). In contrast to traditional collaboration models, interconnection in the PI is not achieved through negotiated agreements between specific actors, but through standardized interaction mechanisms that allow independent logistics networks to operate jointly. It is characterized by several interrelated structural properties. First, it promotes interoperability through standardized protocols (Niemann & Franklin, 2024; Sarraj et al., 2014). Second, coordination is decentralized, with routing and allocation decisions emerging from distributed decision-making mechanisms rather than centralized control (Crainic et al., 2020). Third, it incorporates standardized and modular containers that facilitate seamless transfer across actors and transport modes (Ballot et al., 2014; Montreuil, 2011). Fourth, the PI enables shared accessibility and utilization of logistics assets, allowing flows to be dynamically recombined across organizational boundaries (Chen et al., 2022; Montreuil, 2011). These properties imply that the PI is not merely a coordination mechanism, but it defines how logistics systems interconnect.

[1] Here, "minimal" refers to the smallest set of conditions required for interconnection to function under real-world operational constraints, rather than a simplified or incomplete system.

2.2 Related Collaborative Logistics Paradigms

The following approaches illustrate how existing paradigms structure coordination and where they diverge from the PI's open, protocol-based interconnection.

Horizontal Logistics Collaboration, structures coordination through predefined coalitions of participating firms, relying on joint planning, negotiated agreements, and profit-sharing mechanisms (Cruijssen et al., 2007; Schmoltzi & Wallenburg, 2012; Simatupang & Sridharan, 2002; Verdonck et al., 2013). While enabling efficiency gains, participation remains restricted to predefined coalitions and interconnection depends on ex ante alignment, making HLC structurally incompatible with open, cross-network interoperability.

Groupage Networks like Cargoline improve transport utilization by consolidating shipments through hub-and-spoke structures and standardized handling processes (Rodrigue, 2020). However, central operators define routing, access rules, and standards, making interoperability internal to the network and interconnection centrally controlled rather than distributed.

Digital Freight Platforms like Flexport coordinate logistics services through algorithmic marketplaces, enabling dynamic pricing, real-time visibility, and digital contracting (Heinbach et al., 2022; Herold et al., 2023). Although facilitating flexible matching, participation is governed by the platform operator through proprietary interfaces, keeping interoperability platform-bound and coordination effectively centralized.

2.3 Synthesis: Structural Differences Across Paradigms

To systematically differentiate the PI from existing collaborative logistics paradigms, key structural dimensions of coordination, rules, decision-making, participation, interoperability, standardization, and scalability are synthesized and compared across the identified approaches (Table 1). These dimensions reflect the structural conditions under which interconnection between logistics networks is enabled, constrained, or scaled.

Table 1: Comparison Logistics Paradigms

Dimension	Logistics Collaboration / HLC	Groupage Networks	Digital Freight Platforms	Physical Internet
Coordination Structure	Peer-to-peer coalition	Integrator-led network	Platform-mediated marketplace	Distributed network-of-networks
Who Defines the Rules?	Participating firms jointly	Network operator	Platform owner	Shared open standards
Who Executes Operational Decisions?	Joint planning among partners	Central operator	Platform algorithms under owner control	Distributed routing under shared protocols
Participation Model	Restricted coalition membership	Restricted network membership	Open but platform-controlled	Open under interoperability standards
Interoperability Scope	Limited to coalition agreements	Internal to network	Limited to proprietary interface	Cross-network interoperability
Standardization Requirement	Not structurally required	Internal standards	Digital interface standards	Protocols, Interface standards
Scalability Logic	Expansion of coalition	Expansion of network coverage	Growth of platform user base	Open globally

The comparison highlights that existing collaboration approaches structure coordination through predefined relationships, centralized control, or platform-mediated interaction. As a result, interconnection remains bounded, either to a fixed set of participants, a centrally

governed system, or a proprietary platform environment. At the same time, the PI defines interconnection as an open and protocol-based capability that enables independent logistics networks to interact without prior alignment.

3 Research Methodology

This study follows a Design Science Research (DSR) approach, which aims to develop prescriptive knowledge in the form of design artifacts that address identified problems (Hevner et al., 2004; Peffers et al., 2007). In contrast to explanatory research, DSR focuses on deriving actionable design knowledge, such as algorithms, methods, or design principles, to guide the development of novel systems. To structure this process, the study adopts an eDSR approach (Tuunanen et al., 2024), which decomposes complex design problems into sequential echelons, each producing an intermediate artifact that is independently validated.

Accordingly, we present two echelons. In the problem analysis echelon, we derive a problem statement. We validate it against two criteria: (1) the degree to which the identified problem has been solved, and (2) its solvability (Tuunanen et al., 2024). To assess these criteria, we synthesize PI literature and practitioner-oriented pilot implementations, through which we identify structural limitations that prevent sustained interconnection under real-world conditions. Building on this, in the objectives and requirements echelon, we derive a set of design objectives and requirements. Consistent with the validation criteria proposed by Tuunanen et al. (2024), we develop these objectives and requirements to ensure (1) fit to the validated problem statement and (2) applicability (to the problem).

4 Problem Analysis Echelon

This section presents the problem analysis echelon. Following the methodological approach outlined in Section 3, we derive a problem statement by synthesizing insights from PI research and pilot implementations. The analysis focuses on identifying structural limitations that prevent sustained interconnection between independent logistics networks under real-world conditions.

4.1 Overview of PI research

Table 2: Analysis of Conceptual Maturity of the PI

Layers	Studies
Application	**Service platforms** *(Betti et al., 2019);* **Information sharing & track- & - tracing** *(Fahim et al., 2021a; Frendi et al., 2024)*; **Blockchain** *(Betti et al., 2019; Narzt et al., 2021)*
Transport	**Auction mechanisms** *(Guo et al., 2021; Pan et al., 2014)*; **Dynamic pricing** (He et al., 2024; Qiao et al., 2019); **Inventory deployment** *(Kim et al., 2023; Pan et al., 2015);* **Smart Contracts** *(Betti et al., 2020)*
Network	**Routing protocols** *(Gontara et al., 2018; Muthukrishnan et al., 2025)*; **Hub location & networks** (Ballot et al., 2012; Faugère et al., 2020)*;* **Smart nodes** *(Leung et al., 2022; Rebuglio et al., 2023)*
Link	**Link Protocol** *(Niemann & Franklin, 2024);* **Cross-docking** *(Chargui et al., 2022; Sallez et al., 2015)*; **Hub scheduling** *(Chargui et al., 2020; Essghaier et al., 2023)*; **Intermodal transfer** *(Ambra et al., 2019; Lemmens et al., 2019)*; **Port hubs** *(Caldeirinha et al., 2022; Fahim et al., 2021b)*
Infrastructure	**Modular smart containers** *(Lin et al., 2014; Sallez et al., 2016);* **Container composition & loading** *(Krommenacker et al., 2016; Sternberg & Denizel, 2021)*

PI research has developed a comprehensive body of work across all functional layers of a layered system architecture analogous to the DI, following the 5-layer DoD model (application, transport, network, link, physical layers) (Dong & Franklin, 2021; Wu et al., 2025). Building

on this perspective, Table 2 synthesizes representative research streams, illustrating how existing studies operationalize key functions required for PI operation.

Taken together, these research streams demonstrate that PI research has reached a high level of conceptual maturity. Across all layers, a wide range of mechanisms has been proposed and analytically explored, including routing protocols, pricing and allocation mechanisms, modular containers, and digital coordination platforms (e.g., Crainic et al., 2020; Gontara et al., 2018; Pan et al., 2021; Qiao et al., 2019).

4.2 PI Pilots and Implementation Projects

In parallel, a number of EU-funded pilot initiatives test selected PI mechanisms under real-world constraints. Table 3 provides an overview of representative EU projects. These pilots demonstrate that individual components, digital service platforms, information-sharing tools, and shared logistics hubs, can operate across organizational boundaries within controlled environments. At the same time, all reviewed pilot initiatives remain structurally bounded: participation is limited to predefined actor sets, confined to specific regions and fixed project durations, with interconnection not extending beyond the consortium environment. This indicates that, while the problem of enabling inter-organizational coordination has been partially addressed through individual mechanisms and pilot implementations, it remains unresolved with respect to sustained interconnection between independently operated logistics networks under real-world conditions.

Table 3: Overview Pilot Projects

Project	Mechanisms Addressed	Actors	Dur.	Scope
ICONET	**Application:** Service platforms, Information sharing	**>10**	2018 – 2022	EU logistics network digital infrastructure pilots
URBANE	**Application:** Service platforms, Information sharing **Link:** Urban logistics hubs	**>40**	2021 – 2025	Urban logistics pilots in multiple European cities
DISCO	**Application:** Service platforms, Information sharing	**>40**	2019 – 2023	Digital supply chain orchestration pilots across logistics actors
ePIcenter	**Application:** Service platforms, Information sharing Collaborative Governance	**>30**	2020 – 2024	Global logistics corridors with focus on air cargo

4.3 Synthesis and Problem Statement

The preceding sections reveal a disconnect between conceptual development and practical implementation of the PI. While existing research demonstrates substantial conceptual maturity, and pilot initiatives validate selected mechanisms in real-world settings, implementations remain bounded and do not extend into sustained interconnection between independent logistics networks. This indicates that a central challenge of PI realization does not lie in the absence of technical solutions, but in the lack of clearly defined conditions under which initial interconnection becomes operationally viable and persists over time. Following the evolutionary development of the DI, this suggests the need for a precursor stage, like the ARPANET and follow-up work prior to the TCP/IP deployment in 1983, that enables initial interconnection before full-scale deployment.

The analysis therefore reveals the following core problem: *despite substantial conceptual development and practical experimentation, the PI lacks clearly articulated conditions that enable initial interconnection to become stable and self-sustaining beyond pilot environments*.

5 Objectives and Requirements Echelon

In the following section, we define the objectives and requirements for a precursor system of the PI to address the validated problem identified in Section 4. Following Tuunanen et al. (2024), the resulting objectives and requirements were developed to ensure (1) fit to the validated problem and (2) completeness. To derive and validate these objectives and requirements, we build on three sources of justificatory knowledge: (1) research on operational mechanisms within the PI (e.g., Ballot et al., 2012; Gontara et al., 2018; Sarraj et al., 2014), (2) pilot implementations that provide insights into how these mechanisms are realized under practical conditions (e.g., ICONET, 2018; URBANE, 2022), and (3) the historical evolution of the DI (e.g., Dong & Franklin, 2021; Kurose & Ross, 2017; van Eeten & Mueller, 2013). The objectives were derived through iterative abstraction and triangulation across these sources, translated into minimal design objectives and requirements for sustained and viable interconnection, grounded in the PI principles outlined in Section 2. Following eDSR logic, validation at this echelon is ensured through the justificatory grounding and internal consistency of the derived objectives and requirements.

First, both the PI literature and the historical development of the DI emphasize that large-scale systems evolve through the interconnection of autonomous systems via shared protocols rather than through centralized integration. In the DI and its precursors, this was enabled by packet switching and the separation of networks into independently operated domains, allowing data to be forwarded across heterogeneous infrastructures without requiring internal standardization (Clark, 1988; Kurose & Ross, 2017; McQuillan & Walden, 1977). The PI similarly conceptualizes logistics networks as loosely coupled entities interacting through standardized protocols (Montreuil, 2011; Sarraj et al., 2014). Pilot implementations such as ICONET and URBANE demonstrate that interconnection is feasible in practice, though interoperability relies on predefined interfaces and consortium structures rather than open protocol-based interaction.

Second, the DI shows that end-to-end transmission across independently operated networks is enabled through inter-domain routing protocols such as the Border Gateway Protocol, which allow networks to exchange routing information and dynamically determine paths across network boundaries (Gao, 2001; Kurose & Ross, 2017). PI research adopts this logic by proposing routing mechanisms for distributed logistics networks (Arnau et al., 2022; Crainic et al., 2020), including Border Gateway Protocol-inspired approaches (Gontara et al., 2018), complemented by mechanisms for coordinating transfers between logistics segments such as cross-docking and intermodal operations (Ambra et al., 2019; Chargui et al., 2022). Pilot implementations such as URBANE demonstrate that multi-leg transport is feasible in practice, while routing remains based on predefined or centrally coordinated decisions within bounded environments. To align more closely with the DI principles, routing mechanisms that can operate across network boundaries without relying on predefined coordination structures must be enabled to facilitate continuous flow across independently operated networks.

Third, both the PI vision and the evolution of the DI highlight the importance of incremental and scalable growth. The DI expanded through the continuous addition of autonomous systems, which could join the network independently while maintaining control over their internal operations (Kurose & Ross, 2017; van Eeten & Mueller, 2013). PI research similarly emphasizes open and decentralized structures enabling logistics actors to participate while retaining local autonomy (Ballot et al., 2012; Faugère et al., 2020). Pilot implementations like DISCO and URBANE remain tied to predefined actor groups, resulting in stable but bounded collaboration. Consistent with DI principles, scalable system growth depends on mechanisms that enable the dynamic inclusion of new participants without requiring prior alignment or disrupting existing operations.

Fourth, the DI demonstrates that sustained adoption requires more than technical feasibility: early uptake by influential actors creates network effects that drive broader diffusion. . In the DI, early interconnection expanded organically from a small group of participants as utility became evident (Abbate, 2000; Russell, 2014). In the PI context, roadmaps emphasize adoption, but research provides limited guidance on how diffusion across independently operated logistics networks can be achieved in practice (ALICE, 2020; METI, 2022). Pilot implementations such as DISCO and URBANE demonstrate feasibility, yet remain confined to predefined consortia. Consistent with DI principles, sustained interconnection depends on mechanisms that enable initial uptake by influential early adopters, whose participation encourages subsequent actors to join.

Fifth, both PI research and the evolution of the DI highlight that coordination and governance structures emerge and adapt over time rather than being fully predefined. In the DI, early network coordination relied on centralized arrangements, which gradually evolved into more decentralized governance structures as the network expanded (van Eeten & Mueller, 2013). PI research proposes pricing, incentive, and capacity allocation mechanisms for such coordination (He et al., 2024; Qiao et al., 2019). Pilot implementations such as URBANE and DISCO rely on predefined consortium agreements that remain limited in their ability to adapt to changing participation conditions. Consistent with DI principles, enabling sustained and scalable interconnection requires governance mechanisms that can evolve over time, supported by adaptive coordination rules, pricing and incentive mechanisms, and contractual safeguards that facilitate reliable interaction between independent actors. The resulting minimal conditions for sustained and viable interconnection are summarized in Table 4.

Table 4: Design Objectives and Requirements

Design Objective	Design Requirements	Justificatory Knowledge
DO1: Enable interconnection between heterogeneous logistics networks without requiring prior integration	DR1: Provide standardized interaction protocols for service exchange between heterogeneous logistics networks DR2: Enable minimal onboarding requirements without bilateral agreements	PI: Open interoperability *(Sarraj et al., 2014)* DI: TCP/IP interconnection *(Cerf & Kahn, 1974; Clark, 1988)* Pilots: reliance on predefined consortia
DO2: Enable continuous shipment flow across network boundaries	DR3: Implement decentralized routing mechanisms across networks DR4: Ensure handover mechanisms between networks (multi-leg coordination)	PI: Routing literature *(Gontara et al., 2018; Muthukrishnan et al., 2025)* DI: Packet switching *(McQuillan & Walden, 1977)* Pilots: lack of end-to-end flows
DO3: Enable scalable network growth through dynamic participation	DR5: Allow dynamic entry of new participants without system-wide reconfiguration DR6: Maintain local autonomy of participating networks	PI: Decentralized system logic *(Ballot et al., 2012)* DI: BGP scaling *(Gao, 2001; Rekhter & Li, 1994)* Pilots: fixed actor sets
DO4: Enable initial uptake and diffusion beyond early adopter groups	DR7: Support participation of influential early adopters DR8: Support self-sustaining operational processes (beyond pilots)	DI: Early adoption and diffusion *(Abbate, 2000; Russell, 2014)* Pilots: lack of long-term adoption
DO5: Enable the emergence and adaptation of coordination and governance mechanisms over time	DR9: Allow adaptive coordination mechanisms (not fully predefined) DR10: Support rule evolution and gradual standardization DR11: Enable pricing and incentive mechanisms for participation DR12: Provide contractual safeguards for inter-organizational trust	PI: Pricing literature *(He et al., 2024; Qiao et al., 2019),* Trust literature *(Spanuth et al., 2025)* DI: Governance evolution *(van Eeten & Mueller, 2013)* Pilots: reliance on predefined agreements

6 Conclusion and Outlook

This study clarifies a central challenge in the realization of the PI. While existing research has developed a wide range of mechanisms and pilot implementations have demonstrated their feasibility in controlled environments, it remains unclear how interconnection between independently operated logistics networks can become operationally viable and persist beyond these bounded settings. The analysis shows that addressing this challenge requires a shift in perspective. Rather than focusing on the further development of isolated mechanisms or directly pursuing fully realized system architectures, it is necessary to understand the structural conditions under which interconnection can function under real-world constraints. Drawing on the historical evolution of the DI, the study highlights that such systems emerge through intermediate stages in which initial forms of interconnection are established and gradually stabilized. The notion of a precursor system provides a conceptual basis for this transition by defining a minimal set of conditions that enable sustained interconnection. These insights reposition the focus of PI research toward the problem of stabilizing interconnection across independently operated systems. They also provide a structured foundation for moving beyond project-based experimentation toward more persistent forms of interconnection. Future research should translate these conditions into an implementable precursor system and evaluate its performance under realistic operational settings. This includes, for example, simulation-based analyses and extended pilot implementations that examine whether interconnection can persist beyond predefined consortia and controlled environments. In this regard, the precursor system represents a critical step in enabling the gradual transition from experimental PI implementations to sustained, system-level interconnection.

References

Abbate, J. (2000). *Inventing the internet*. MIT press.

Alexander, F. M., DeNardis, L., Levinson, N. S., & Musiani, F. (2025). *Geopolitics at the Internet's Core*. Springer.

ALICE. (2020, November 13). *ALICE Roadmap to Physical Internet released!* https://www.etp-logistics.eu/alice-physical-internet-roadmap-released/

Ambra, T., Caris, A., & Macharis, C. (2019). Towards freight transport system unification: Reviewing and combining the advancements in the physical internet and synchromodal transport research. *International Journal of Production Research*, *57*(6), 1606–1623. https://doi.org/10.1080/00207543.2018.1494392

Arnau, Q., Barrena, E., Panadero, J., de la Torre, R., & Juan, A. A. (2022). A biased-randomized discrete-event heuristic for coordinated multi-vehicle container transport across interconnected networks. *European Journal of Operational Research*, *302*(1), 348–362. https://doi.org/10.1016/j.ejor.2021.12.035

Ballot, E., Gobet, O., & Montreuil, B. (2012). Physical internet enabled open hub network design for distributed networked operations. In *Service orientation in holonic and multi-agent manufacturing control* (pp. 279–292). Springer.

Ballot, E., Montreuil, B., & Meller, R. D. (2014). *The Physical Internet* (Predit, Ed.). La Documentation Française. https://hal-mines-paristech.archives-ouvertes.fr/hal-01113648

Betti, Q., Khoury, R., Hallé, S., & Montreuil, B. (2019). Improving Hyperconnected Logistics With Blockchains and Smart Contracts. *IT Professional*, *21*(4), 25–32. IT Professional. https://doi.org/10.1109/MITP.2019.2912135

Betti, Q., Montreuil, B., Khoury, R., & Hallé, S. (2020). Smart Contracts-Enabled Simulation for Hyperconnected Logistics. In M. A. Khan, M. T. Quasim, F. Algarni, & A. Alharthi (Eds.), *Decentralised Internet of Things: A Blockchain Perspective* (pp. 109–149). Springer International Publishing. https://doi.org/10.1007/978-3-030-38677-1_6

Caldeirinha, V., Nabais, J. L., & Pinto, C. (2022). Port Community Systems: Accelerating the Transition of Seaports toward the Physical Internet—The Portuguese Case. *Journal of Marine Science and Engineering*, *10*(2), Article 2. https://doi.org/10.3390/jmse10020152

Cerf, V., & Kahn, R. (1974). A Protocol for Packet Network Intercommunication. *IEEE Transactions on Communications*, *22*(5), 637–648. IEEE Transactions on Communications. https://doi.org/10.1109/TCOM.1974.1092259

Chargui, T., Bekrar, A., Reghioui, M., & Trentesaux, D. (2020). Proposal of a multi-agent model for the sustainable truck scheduling and containers grouping problem in a Road-Rail physical internet hub. *International Journal of Production Research*, *58*(18), 5477–5501.

Chargui, T., Ladier, A.-L., Bekrar, A., Pan, S., & Trentesaux, D. (2022). Towards designing and operating physical internet cross-docks: Problem specifications and research perspectives. *Omega*, *111*, 102641.
Chen, S., Su, L., & Cheng, X. (2022). Physical Internet deployment in industry: Literature review and research opportunities. *Industrial Management & Data Systems*, *122*(6), 1522–1540. https://doi.org/10.1108/IMDS-07-2021-0416
Clark, D. (1988). The design philosophy of the DARPA internet protocols. *Symposium Proceedings on Communications Architectures and Protocols, SIGCOMM '88*, 106–114. https://doi.org/10.1145/52324.52336
Crainic, T. G., Gendreau, M., & Jemai, L. (2020). Planning hyperconnected, urban logistics systems. *Transportation Research Procedia, 22nd EURO Working Group on Transportation Meeting, EWGT 2019, 18th – 20th September 2019, Barcelona, Spain*, *47*, 35–42. https://doi.org/10.1016/j.trpro.2020.03.070
Cruijssen, F., Cools, M., & Dullaert, W. (2007). Horizontal cooperation in logistics: Opportunities and impediments. *Transportation Research Part E: Logistics and Transportation Review*, *43*(2), 129–142. https://doi.org/10.1016/j.tre.2005.09.007
Dong, C., & Franklin, R. (2021). From the Digital Internet to the Physical Internet: A Conceptual Framework With a Stylized Network Model. *Journal of Business Logistics*, *42*(1), 108–119.
Essghaier, F., Chargui, T., Hsu, T., Bekrar, A., Allaoui, H., Trentesaux, D., & Goncalves, G. (2023). Fuzzy multi-objective truck scheduling in multi-modal rail–road Physical Internet hubs. *Computers & Industrial Engineering*, *182*, 109404.
Fahim, P. B. M., An, R., Rezaei, J., Pang, Y., Montreuil, B., & Tavasszy, L. (2021). An information architecture to enable track-and-trace capability in Physical Internet ports. *Computers in Industry*, *129*, 103443. https://doi.org/10.1016/j.compind.2021.103443
Fahim, P. B. M., Martinez de Ubago Alvarez de Sotomayor, M., Rezaei, J., van Binsbergen, A., Nijdam, M., & Tavasszy, L. (2021). On the evolution of maritime ports towards the Physical Internet. *Futures*, *134*, 102834. https://doi.org/10.1016/j.futures.2021.102834
Faugère, L., White III, C., & Montreuil, B. (2020). Mobile access hub deployment for urban parcel logistics. *Sustainability*, *12*(17), 7213.
Frendi, M., Nachet, B., & Adla, A. (2024). Physical Internet Based Ontology for Supporting Traceability in Logistic IoT. *International Journal of Computing and Digital Systems*, *15*(1), 1–13. https://doi.org/10.12785/ijcds/150133
Gao, L. (2001). On inferring autonomous system relationships in the Internet. *IEEE/ACM Transactions on Networking*, *9*(6), 733–745.
Gontara, S., Boufaied, A., & Korbaa, O. (2018). Routing the Pi-Containers in the Physical Internet using the PI-BGP Protocol. *2018 IEEE/ACS 15th International Conference on Computer Systems and Applications (AICCSA)*, 1–8. https://doi.org/10.1109/AICCSA.2018.8612885
Guo, C., Thompson, R. G., Foliente, G., & Kong, X. T. (2021). An auction-enabled collaborative routing mechanism for omnichannel on-demand logistics through transshipment. *Transportation Research Part E: Logistics and Transportation Review*, *146*, 102206.
He, S., Zhang, M., Wang, S., & Huang, G. Q. (2024). Pricing and revenue allocation mechanisms for regional logistics networks in Cyber-Physical Internet. *Transportation Research Part C: Emerging Technologies*, *165*, 104714. https://doi.org/10.1016/j.trc.2024.104714
Heinbach, C., Beinke, J., Kammler, F., & Thomas, O. (2022). Data-driven forwarding: A typology of digital platforms for road freight transport management. *Electronic Markets*, *32*(2), 807–828. https://doi.org/10.1007/s12525-022-00540-4
Herold, D. M., Fahimnia, B., & Breitbarth, T. (2023). The digital freight forwarder and the incumbent: A framework to examine disruptive potentials of digital platforms. *Transportation Research Part E: Logistics and Transportation Review*, *176*, 103214. https://doi.org/10.1016/j.tre.2023.103214
Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design Science in Information Systems Research. *MIS Quarterly*, *28*(1), 75–105. https://doi.org/10.2307/25148625
ICONET. (2018). *New ICT infrastructure and reference architecture to support Operations in future PI Logistics NETworks*. CORDIS | European Commission. https://doi.org/10.3030/769119
Kim, N., Montreuil, B., Klibi, W., & Babai, M. Z. (2023). Network inventory deployment for responsive fulfillment. *International Journal of Production Economics*, *255*, 108664.
Krommenacker, N., Charpentier, P., Berger, T., & Sallez, Y. (2016). On the usage of wireless sensor networks to facilitate composition/decomposition of physical internet containers. In *Service orientation in Holonic and multi-agent manufacturing* (pp. 81–90). Springer.
Kurose, J. F., & Ross, K. W. (2017). *Computer networking: A top-down approach* (Seventh edition). Pearson.
Landschützer, C., Ehrentraut, F., & Jodin, D. (2015). Containers for the Physical Internet: Requirements and engineering design related to FMCG logistics. *Logistics Research*, *8*(1), 8. https://doi.org/10.1007/s12159-015-0126-3
Leiner, B. M., Cerf, V. G., Clark, D. D., Kahn, R. E., Kleinrock, L., Lynch, D. C., Postel, J., Roberts, L. G., & Wolff, S. (2009). A brief history of the internet. *SIGCOMM Comput. Commun. Rev.*, *39*(5), 22–31. https://doi.org/10.1145/1629607.1629613
Lemmens, N., Gijsbrechts, J., & Boute, R. (2019). Synchromodality in the Physical Internet–dual sourcing and real-time switching between transport modes. *European Transport Research Review*, *11*(1), 19.
Leung, E. K. H., Lee, C. K. H., & Ouyang, Z. (2022). From traditional warehouses to Physical Internet hubs: A digital twin-based inbound synchronization framework for PI-order management. *International Journal of Production Economics*, *244*, 108353. https://doi.org/10.1016/j.ijpe.2021.108353

Lin, Y.-H., Meller, R. D., Ellis, K. P., Thomas, L. M., & Lombardi, B. J. (2014). A decomposition-based approach for the selection of standardized modular containers. *International Journal of Production Research*, *52*(15), 4660–4672. https://doi.org/10.1080/00207543.2014.883468
McQuillan, J. M., & Walden, D. C. (1977). The ARPA network design decisions. *Computer Networks (1976)*, *1*(5), 243–289.
METI. (2022). *Roadmap for Realizing the Physical Internet*. Ministry of Economy, Trade and Industry. https://www.meti.go.jp/shingikai/mono_info_service/physical_internet/pdf/20220308_1.pdf
Montreuil, B. (2011). Toward a Physical Internet: Meeting the global logistics sustainability grand challenge. *Logistics Research*, *3*(2), 71–87. https://doi.org/10.1007/s12159-011-0045-x
Muthukrishnan, P., Sankar, S., Shaikh, S. J., Gaffney, J. C., & Montreuil, B. (2025). Dynamic Freight Routing and Dispatch Protocols for the Physical Internet. *11th International Physical Internet Conference*.
Narzt, W., Schützeneder, P., Dragoev, P., Schatzlmayr-Piekarz, B., & Schwaiger, M. (2021). *Blockchain-Based Load Carrier Management in the Physical Internet*. *12783 LNCS*, 417–432. Scopus. https://doi.org/10.1007/978-3-030-77750-0_26
Niemann, G., & Franklin, R. (2024). A protocol for node-to-node transshipment in the PI analogizing from Internet principles. *10th International Physical Internet Conference*. International Physical Internet Conference.
Pan, S., Nigrelli, M., Ballot, E., Sarraj, R., & Yang, Y. (2015). Perspectives of inventory control models in the Physical Internet: A simulation study. *Computers & Industrial Engineering*, *84*, 122–132.
Pan, S., Trentesaux, D., McFarlane, D., Montreuil, B., Ballot, E., & Huang, G. Q. (2021). Digital interoperability in logistics and supply chain management: State-of-the-art and research avenues towards Physical Internet. *Computers in Industry*, *128*, 103435. https://doi.org/10.1016/j.compind.2021.103435
Pan, S., Xu, X., & Ballot, E. (2014). *Auction-based transport services allocation in Physical Inter-net: A simulation framework*.
Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A Design Science Research Methodology for Information Systems Research. *Journal of Management Information Systems*, *24*(3), 45–77. https://doi.org/10.2753/MIS0742-1222240302
Qiao, B., Pan, S., & Ballot, E. (2019). Dynamic pricing model for less-than-truckload carriers in the Physical Internet. *Journal of Intelligent Manufacturing*, *30*(7), 2631–2643. https://doi.org/10.1007/s10845-016-1289-8
Rebuglio, M., Ferrari, A., Zenezini, G., & Grimaldi, S. (2023). Integrating an AS/RS with Digital Twin in a Physical Internet network. In *Proceedings of the Summer School Francesco Turco*. 28th Summer School Francesco Turco, 2023.
Rekhter, Y., & Li, T. (1994). *RFC1654: A Border Gateway Protocol 4 (BGP-4)*.
Rodrigue, J.-P. (2020). *The geography of transport systems*. Routledge.
Russell, A. L. (2014). *Open standards and the digital age*. Cambridge University Press.
Sallez, Y., Berger, T., Bonte, T., & Trentesaux, D. (2015). Proposition of a hybrid control architecture for the routing in a Physical Internet cross-docking hub. *IFAC-PapersOnLine, 15th IFAC Symposium onInformation Control Problems inManufacturing*, *48*(3), 1978–1983. https://doi.org/10.1016/j.ifacol.2015.06.378
Sallez, Y., Pan, S., Montreuil, B., Berger, T., & Ballot, E. (2016). On the activeness of intelligent Physical Internet containers. *Computers in Industry, Emerging ICT Concepts for Smart, Safe and Sustainable Industrial Systems*, *81*, 96–104. https://doi.org/10.1016/j.compind.2015.12.006
Sarraj, R., Ballot, E., Pan, S., Hakimi, D., & Montreuil, B. (2014). Interconnected logistic networks and protocols: Simulation-based efficiency assessment. *International Journal of Production Research*, *52*(11), 3185–3208. https://doi.org/10.1080/00207543.2013.865853
Schmoltzi, C., & Wallenburg, C. M. (2012). Operational Governance in Horizontal Cooperations of Logistics Service Providers: Performance Effects and the Moderating Role of Cooperation Complexity. *Journal of Supply Chain Management*, *48*(2), 53–74. https://doi.org/10.1111/j.1745-493X.2011.03262.x
Simatupang, T. M., & Sridharan, R. (2002). The Collaborative Supply Chain. *The International Journal of Logistics Management*, *13*(1), 15–30. https://doi.org/10.1108/09574090210806333
Spanuth, M., Franklin, R., Ludwig, A., & Niemann, G. (2025). Governing the Physical Internet: Insights from Internet Governance and Future Research Directions. *11th International Physical Internet Conference*.
Sternberg, H. S., & Denizel, M. (2021). Toward the Physical Internet—Logistics Service Modularity and Design Implications. *Journal of Business Logistics*, *42*(1), 144–166. https://doi.org/10.1111/jbl.12261
Tuunanen, T., Winter, R., & Vom Brocke, J. (2024). Dealing with Complexity in Design Science Research: Using Design Echelons to Support Planning, Conducting, and Communicating Design Knowledge Contributions. *MIS Quarterly*, *48*(2). https://doi.org/10.25300/MISQ/2023/16700
URBANE. (2022). *Upscaling innovative green urban logistics solutions through multi-actor collaboration and PI-inspired last mile deliveries*. CORDIS | European Commission. https://doi.org/10.3030/101069782
van Eeten, M. J., & Mueller, M. (2013). Where is the governance in Internet governance? *New Media & Society*, *15*(5), 720–736. https://doi.org/10.1177/1461444812462850
Verdonck, L., Caris, A., Ramaekers, K., & Janssens, G. K. (2013). Collaborative Logistics from the Perspective of Road Transportation Companies. *Transport Reviews*, *33*(6), 700–719. https://doi.org/10.1080/01441647.2013.853706
Wu, H., Li, M., Yu, C., Ouyang, Z., Lai, K., Zhao, Z., Pan, S., Wang, S., Zhong, R. Y., Kuo, Y.-H., Zhang, F., Huang, W., Shen, Z.-J. M., Ballot, E., & Huang, G. Q. (2025). Towards cyber-physical internet: A systematic review, fundamental model and future perspectives. *Transportation Research Part E: Logistics and Transportation Review*, *197*, 104051. https://doi.org/10.1016/j.tre.2025.104051

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Federated Autonomy and Availability Intelligence in Multi-Organizational Physical Internet Ecosystems for Critical Supply Resilience

Ashwin S. Pothen[1], Benoit Montreuil[1,2], and Sahrish Jaleel Shaikh[1]

1. Physical Internet Center, Supply Chain and Logistics Institute, H. Milton Stewart School of Industrial & Systems Engineering, Georgia Institute of Technology, Atlanta, United States
2. Coca-Cola Chair in Material Handling and Distribution

Corresponding author: ashwin.pothen@gatech.edu

Keywords: *Availability Intelligence, Hyperconnected Logistics, Multi-Organizational, Physical Internet, Societally Critical Products, Resilience*

Conference Domain Fitness: *This research aligns with IPIC 2026's focus on advancing Physical Internet knowledge, theories, models, and technologies enabling interconnected, open, and intelligent logistics systems. It contributes a modeling and methodological framework for hyperconnected supply chain ecosystems, emphasizing interoperability, open asset sharing, and collaborative logistics across independent organizations. The proposed federated autonomy architecture supports horizontal cooperation and platform-based coordination consistent with the Physical Internet paradigm. By integrating digital twin intelligence, modular containerization, and equity-aware allocation mechanisms, the work advances resilient and sustainable supply chain ecosystems, directly supporting the conference theme of Supply Chain and Logistics Intelligence in the Era of the Physical Internet.*

Physical Internet (PI) Roadmap Fitness*: ☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories), ☒ PI Networks, ☒ System of Logistics Networks, ☒ Vertical Supply Consolidation, ☒ Horizontal Supply Chain Alignment, ☒ Access and Adoption.*

Targeted Delivery Mode-s*: ☒ Paper, ☒ Flash Video, ☒ In-Person presentation*

Research Contribution Abstract

The Physical Internet (PI) paradigm promotes hyperconnected, interoperable logistic ecosystems capable of delivering order-of-magnitude improvements in economic, environmental, and societal performance [1]. While extensive research has explored modular containerization, asset sharing, and synchromodal flow consolidation, comparatively less attention has been given to how independent organizations can coordinate availability decisions without centralized control. Building upon our prior work on hyperconnected critical-product distribution and autonomy-driven operations [2], this paper proposes a federated autonomy architecture for availability intelligence across multi-organizational PI ecosystems.

In earlier research, we designed and implemented a multi-tier, demand-driven distribution system for critical consumables during the COVID-19 pandemic. That system integrated dynamic forecasting

[3], autonomy-based inventory targets, shared control tower visibility, and simulation-optimization to achieve seamless distribution with no stockouts and reduced inventory levels. However, the framework operated within a single institutional governance structure. Extending this model to independent organizations such as hospitals, primary care clinics, and educational institutions introduces additional challenges, including heterogeneous objectives, partial visibility, trust constraints, and equity considerations under disruption [4, 5].

We propose a federated autonomy architecture that preserves local decision sovereignty while enabling structured collaboration through interoperable π-containers, shared hyperconnected hubs, and graded visibility regimes. Each organization maintains decentralized forecasting capabilities and autonomy targets representing desired robustness against stochastic demand and supply uncertainty. Collaboration is parameterized through a shareable inventory fraction, while visibility is modeled through progressively richer information exchange protocols. Rather than relying on centralized command or vendor-managed inventory models, coordination occurs via a federated cooperation protocol that allows lateral flows and adaptive redistribution under disruption, consistent with open interoperability principles of the Physical Internet.

To evaluate the proposed framework, we develop a region-level simulation environment representing a heterogeneous ecosystem of healthcare and educational institutions connected through shared π-hubs and common suppliers. Stochastic demand surges and supply-side disruptions are introduced to emulate pandemic-like stress conditions. We systematically assess the impact of collaboration intensity, visibility granularity, and vulnerability-weighted allocation policies on inventory efficiency, service robustness, and time-to-recovery. Equity integration draws upon structured vulnerability considerations to evaluate not only aggregate performance but also socially differentiated service continuity.

The simulation study is designed to test a set of structural hypotheses regarding federated autonomy in multi-organizational PI ecosystems. First, we hypothesize the existence of collaboration threshold effects, whereby limited inventory sharing yields modest improvements, while structured federated autonomy beyond a minimum collaboration level generates disproportionate gains in inventory efficiency and service robustness. Second, we hypothesize that under disruption, federated configurations reduce time-to-recovery relative to siloed systems by enabling lateral redistribution across institutions through shared π-hubs. Third, we examine whether increasing visibility exhibits diminishing marginal returns, such that intermediate visibility regimes capture most resilience benefits without requiring full transparency. Finally, we investigate whether vulnerability-weighted allocation mechanisms can enhance service continuity for institutions serving socially vulnerable populations while maintaining acceptable system-wide efficiency trade-offs. These hypotheses collectively guide the experimental design and evaluation of the proposed federated autonomy framework.

Overall, availability intelligence in PI ecosystems is fundamentally a socio-technical orchestration challenge requiring balanced governance, interoperability, and selective transparency. By integrating autonomy-based decision logic, modular containerization, and equity-aware allocation within a federated coordination framework, this research advances the operationalization of multi-organizational Physical Internet systems for critical supply resilience. The work contributes (1) a formalized federated autonomy architecture extending single-enterprise hyperconnected systems to independent organizations, (2) a systematic simulation-based evaluation of collaboration and visibility trade-offs at a regional scale, and (3) quantitative evidence supporting resilient and socially responsive availability management without centralized control.

References

1. Montreuil, B., Ballot, E., Tremblay, W., 2015. Modular Design of Physical Internet Transport, Handling and Packaging Containers, in: Progress in Material Handling Research: 2014, 13, MHI, 2015, International Material Handling Research Colloquium.
2. Shaikh, S.J., Pothen, A.S., Montreuil, B., 2023. Hyperconnected Critical-Product Supply and Distribution System: Towards Autonomous Operations. IFAC-PapersOnLine 56, 7579–7584. https://doi.org/10.1016/j.ifacol.2023.10.669
3. Yim, J., Montreuil, B., Yang, S., Kim, N., 2021. Data-driven prediction of demand in a pandemic for critical consumables over a shared distribution network. Proceedings of 2021 IISE Annual Conference, pp. 788–793.
4. Tang, C., Tomlin, B., 2008. The power of flexibility for mitigating supply chain risks. International Journal of Production Economics 116, 12–27. https://doi.org/10.1016/j.ijpe.2008.07.008
5. Ergun, O., Hopp, W.J., Keskinocak, P., 2023. A structured overview of insights and opportunities for enhancing supply chain resilience. IISE Transactions 55, 57–74. https://www.tandfonline.com/doi/full/10.1080/24725854.2022.2080892

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Place your logo here

Toward a Physical Internet for Food Aid: Measuring Node Capacity and Need to Improve Cost, Reliability, and Fairness

Abdullah Konak and Sadan Kulturel-Konak
Penn State Berks, Reading, PA, USA
Corresponding author: auk3@psu.edu

Abstract: *Food banks across the United States are under pressure. Demand is unpredictable, supply is inconsistent, and many people who need help cannot access it due to transportation barriers or gaps in service coverage. However, most planning tools treat food distribution as a simple logistics problem, ignoring the reality that outcomes depend on a network of agencies, pantries, volunteers, and community organizations, each with different capacities and local need constraints. This paper reimagines charitable food distribution through a Physical Internet (PI) framework, treating the network as a system of interconnected nodes that can route food across space and time. Such an approach requires analyzing network nodes, which are agencies that distribute food from food banks. We develop a six-step decision-support framework that combines census block-level need estimation, spatial allocation modeling, and vehicle routing optimization to evaluate how effectively each agency serves its surrounding community, relative to the marginal distance to reach it. Applied to a regional food bank in Pennsylvania, the framework reveals a clear pattern: agencies closest to high-need urban areas deliver the greatest impact at the lowest cost, while rural agencies bear a disproportionate share of the routing burden. These findings highlight clear opportunities for PI-inspired redesign, including shared consolidation hubs and smarter last-mile routing, that can improve both fairness and efficiency without requiring major new resources.*

Keywords: *Physical Internet; Node capacity planning; Charitable food distribution; Network fairness*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☐ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☒ *Horizontal Supply Chain Alignment,* ☒ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

1 Introduction

Charitable food distribution programs play a critical role in mitigating food insecurity in the US. A key component of charitable food distribution systems is food banks, which are nonprofit organizations that collect donated food and redistribute it to people experiencing food insecurity. However, persistent inequities in food access remain across neighborhoods and populations. Vulnerable communities, particularly those in geographically isolated rural areas and underserved urban neighborhoods, often face disproportionate barriers to reliable food access due to a combination of socioeconomic factors and structural obstacles, including limited transportation, restricted operating hours, and fragmented service networks. At the same time, food banks must operate under tightening resource constraints and navigate complex operational tradeoffs involving fluctuating supply and donations, volunteer-dependent labor,

storage limitations, and routing and scheduling decisions (Akkerman et al., 2023). These challenges echo broader concerns in humanitarian logistics, emphasizing that distribution systems should not only be cost-effective and responsive but also fair, transparent, and sustainable in the allocation of scarce resources (Iqbal & Ahmad, 2022).

At the most basic level, food banks operate as intermediaries in a supply chain connecting food donors (retailers, manufacturers, farms, government surplus programs, and individual donors) to agencies that directly serve people (food pantries, soup kitchens, shelters, senior housing, and schools). The usual setup is a hub-and-spoke model. A central food bank warehouse receives bulk donations, sorts and stores them, and then distributes food to a network of downstream agencies, which are local pantries and community organizations that interact directly with clients. Some food banks also have kitchens to prepare and distribute hot meals. In a broad sense, food banks aim to balance three often-competing goals (Islam & Ivy, 2022; Sengul Orgut et al., 2016): *equity* (distributing food proportionally to need across their service area), *effectiveness* (maximizing the total amount of food distributed), and *efficiency* (minimizing operational costs). Recent work on food bank operations has formalized equity–efficiency tradeoffs, showing that modest relaxations from perfect equity can substantially improve total food distributed and reduce waste (Sengul Orgut & Lodree, 2023), while also highlighting the constraints imposed by capacity and geography.

The Physical Internet (PI) has been proposed as a revolutionary concept in logistics, in which goods are routed through an open, highly connected network of shared logistics facilities, similar to how data packets move through the Digital Internet. Emerging applications demonstrate how PI-inspired concepts, modular containers, shared hubs, and smart coordination can support more efficient and resilient last-mile services in humanitarian logistics (Grest et al., 2019), including food delivery to vulnerable populations. For example, Lowrey and Boyer (2024) studied a model called “Pantry Direct,” in which local pantry volunteers pick up food directly from retail stores, bypassing the central warehouse. Zou et al. (2025) noted that a larger on-site pantry can help control food waste by providing a direct means of serving clients, thereby facilitating better control based on inventory levels. Similarly, Stauffer (2022) suggested that increased partner agency capacity can help address variability in supply and distribution capacity while accounting for food expiration times. However, the implications of a distributed, PI-style architecture in charitable food systems remain largely unexplored, despite the clear potential to rethink how agencies, mobile distributions, and community sites are connected in an open network that jointly optimizes reach, cost, and fairness. We provide a framework for identifying candidate PI nodes and equity–efficiency tensions in charitable food distribution.

There is a growing need for a decision-support framework that accounts for fairness in network design and operations while embracing open interconnectivity and shared infrastructure. Agency-level decisions need more geographic detail than county-level research typically provides. The framework must map both food needs and actual charitable distribution patterns at the neighborhood scale. From there, it can evaluate whether the current agency network is doing its job fairly and compare service levels to actual need to identify underserved areas, isolated pockets, and overlapping coverage that summary statistics miss. With this detailed picture, we can then test whether consolidating networks, creating shared hubs, or reorganizing in other ways might improve both equity and efficiency in food bank operations. This paper aims to address the gaps identified above. The objectives of the study are threefold: (1) to develop and implement a modeling framework that estimates food need at detailed geographic scales and links it to observed patterns of charitable food distribution; (2) to assess the fairness of the existing agency network in terms of service levels relative to estimated need, highlighting

underserved and isolated areas; and (3) to explore how PI-inspired consolidation, shared hubs, and reorganization strategies might impact both equity and efficiency in food bank operations. The framework is implemented in a regional charitable food distribution network to explore whether a PI-inspired, hyperconnected network design can better align service provision with estimated needs. This implementation follows a participatory research approach, involving food bank staff, partner agencies, community members, and other stakeholders in shaping the analysis and interpreting its implications. The analysis combines operational food distribution data with census-based indicators and publicly available food access datasets to develop block-level estimates of food need. It then links these estimates to observed distribution patterns using spatial analysis. These findings are further analyzed alongside operational measures, including marginal delivery distance, agency isolation, and opportunities for shared hubs and modular flows, to evaluate how different network configurations may impact both equity and efficiency.

2 Methodology

2.1 Defining the Framework

We adopt a six-step framework to evaluate partner agencies within a charitable food distribution network. The framework integrates need estimation, spatial analysis, and operational considerations to support equity-informed decision-making.

1. Define the Objectives: The first step is to clearly define the performance metrics of the analysis. The literature includes several fairness-oriented performance metrics, such as fill-rate service relative to need, access, delivery time regularity, and distribution equity. In charitable food distribution, equity (allocating based on need) is typically preferred over equal distribution to all. In addition to fairness-oriented metrics, our approach considers operational costs (efficiency). The operational costs are not absolute objectives to minimize; therefore, we opt for marginal costs of serving, such as additional travel distance or routing complexity due to vehicle capacity constraints.

2. Data Collection and Integration: Data from multiple sources are collected and combined into a unified dataset. These sources include food distribution ERP records, agency characteristics, geospatial, and publicly available socioeconomic data. Data integration ensures that all relevant details, both operational, demographic, and geospatial, are aligned and usable for analysis.

3. Decide Spatial Unit of Analysis: Unit of Analysis refers to the geographic unit used to represent need and service, such as a county or a census tract. Each analysis unit aggregates population characteristics (e.g., need, access barriers), and typically, the county level is used as the unit of analysis in fairness-oriented literature since many census statistics are available at this level. However, since the objective is to evaluate each agency, it is critical to use a more granular unit of analysis to achieve the precision needed to estimate needs and interpret access to services. In our analysis, we therefore use a census tract or census block as the unit of analysis. Census block groups represent the smallest geographic unit for which detailed demographic and socioeconomic data are publicly available through the American Community Survey (ACS). The selection of a granular unit of analysis also has limitations, including limited available data. In the case study, we discuss how we address data availability at this scale.

Let $CT = \{1, \dots, K\}$ be the set of census tracts (indexed by k), $CB = \{1, \dots, I\}$ the set of census blocks (indexed by i), $CB(k) \subseteq CB$ and the set of census blocks belonging to census tract k. Each census block $i \in CB$ is represented by a population center coordinates (lon_i, lat_i) and its population P_i^{CB}.

4. Estimate Need in Each Unit of Analysis

In charitable food systems, the need is usually conceptualized as the gap between the food required to maintain a minimum level of food security and the food that can be reliably accessed, purchased, or obtained through programs. This definition of food insecurity addresses not only insufficient supply but also economic, physical, and social access constraints (Coleman-Jensen, 2010). Because it is impossible to directly measure the food shortfall without individual-level data, we rely on observable indicators to approximate both the magnitude of need and the barriers that shape it. The following indicators capture distinct but interrelated dimensions of need. Let N_k^{CT} be an indicator of need in CT *k*. N_k^{CT} can be a single indicator (e.g., poverty rate, population below the poverty line) or a composite indicator that aggregates many indicators. Because the indicators at the census block level are rarely available, the census tract-level need indicators are downscaled to the block level by allocating the tract-level need to each block in proportion to the block's share of the tract population as follows:

$$N_i^{CB} = \frac{P_i^{CB}}{P_k^{CT}} N_k^{CT} \; i \in CB(k)$$

where P_i^{CB} is the population of census block *i,* and P_k^{CT} is the population of census tract *k*.

5. Assess the Potential of Each Agency to Capture Need

Each agency's potential to capture need is calculated to develop a network-based understanding of service delivery. Agencies do not operate in isolation. In practice, individuals can access multiple nearby agencies, and agencies share responsibility for serving the same community. Simply measuring how much food an agency distributes does not indicate whether it is serving high-need areas or overlapping with other agencies. If multiple agencies capture the same need, this may indicate overlapping services and redundancies. If certain areas have a low probability of being served by any agency, this reveals underserved areas or gaps in the distribution network. Identifying the needs captured by each agency is also essential for equity, which is not only about total distribution but also about serving high-need populations.

Finally, the potential needs each agency can capture provide a foundation for optimization and policy decisions, such as reallocating resources to better meet those needs, increasing capacity, or consolidating agencies.

Each agency's potential to capture surrounding need is evaluated using a Huff-type spatial allocation model where the attractiveness (w_{ij}) of a census block i being served by agency j decreases with distance (d_{ij}) as follows

$$w_{ij} = \left(1 + \frac{d_{ij}}{\theta(j)}\right)^{-\beta}$$

where β is the distance decay parameter, and $\theta(j)$ is the distance scale over which an agency remains attractive. The parameter β represents behavioral sensitivity to travel burden. A larger value of $\theta(j)$ means that agency j typically has a wider service area (e.g., rural areas), whereas a smaller value indicates that agency j typically has a narrower service area (e.g., urban areas). We assume the need from each block is split across multiple agencies probabilistically:

$$p_{ij} = \frac{w_{ij}}{\sum_{l \in J(i)} w_{il}}$$

where $J(i)$ is set of candidate agencies is accessible from block i. Hence, the potential need captured by agency j is given as follows:

$$N_j = \sum_{i \in CB} p_{ij} N_i^{CB}$$

In turn, N_j is a measure of how much each agency j is likely to capture and reflects its service reach within the network.

6. Estimate the Marginal Cost of Delivery for Each Agency

To evaluate the operational burden of delivering to each agency, we estimate the marginal delivery cost in terms of additional travel distance. This measure captures the additional effort required to include a specific agency in the delivery network. A key challenge in this analysis is the lack of reliable historical routing data. While delivery schedules and the quantity of food delivered to each agency (in pounds) were available, detailed route-level information was either unavailable or inconsistent. As a result, we adopt a mathematical modeling-based approach to approximate delivery costs as follows.

For each day t, we solve a capacitated vehicle routing problem with time windows (CVRPTW) using the known delivery demands and vehicle capacity constraints. This yields a feasible and efficient set of delivery routes. To estimate the marginal contribution of agency j on day t, we use a route-based distance decomposition. Specifically, for each occurrence of agency j in a route on day t, we compute the additional distance associated with visiting that agency as:

$$\Delta_{jt} = d_t(\text{prev}(j), j) + d_t(j, \text{next}(j)) - d_t(\text{prev}(j), \text{next}(j))$$

where prev(j) and next(j) denote the predecessor and successor stops of agency j in the route, and $d_t(\cdot,\cdot)$ is the travel distance between two locations on day t.

This expression measures the extra distance incurred by inserting agency j into the route. The marginal delivery cost of agency j is then computed as the average additional distance over all days:

$$\Delta_j = \text{Avg}_t(\Delta_{jt})$$

This method provides a local approximation of marginal delivery cost based on route structure, rather than a full re-optimization of the routing problem after removing each agency. It is computationally efficient, eliminating the need to solve many extra CVRP instances.

3 Case Study

This case study focuses on two counties in the Northeast US. With around 12% of the population in both counties living in poverty, food banks are essential for providing vital support to those who struggle to afford basic food. Both counties are served by a single food bank. Over 110,000 individuals in both counties qualify for food assistance, creating significant operational and logistical challenges for partner agencies and food banks in providing support and resources. Both counties have a highly diverse population, ranging from urban areas to numerous rural communities. The procurement and distribution of various foods that meet the diverse dietary needs of communities with different cultural backgrounds also pose another challenge. In response to these pressing needs, the local food bank annually distributes nearly 10 million pounds of food valued at over $15,000,000 through nearly 350 charitable food program partners (i.e., food pantries, soup kitchens, shelters, after-school programs, mobile markets, and senior housing facilities).

3.1 Data Sources and Data Preprocessing

The spatial unit of analysis in this study is the 2020 Census Block Group, as defined by the U.S. Census Bureau. Census block groups represent the smallest geographic unit for which detailed demographic and socioeconomic data are consistently available through the American Community Survey (ACS), making them well-suited for capturing localized variation in food need and access barriers at a scale meaningful for agency-level planning. The Food Access Research Atlas (FARA), developed and maintained by the United States Department of Agriculture Economic Research Service (USDA-ERS), is a comprehensive census tract-level mapping and data tool designed to assess food access conditions across the United States. The FARA includes a series of indicators that characterize the extent to which low-income populations and other demographic groups face limited access to affordable, nutritious food, with particular emphasis on supermarket accessibility as a proxy for food environment quality (Service, 2026). In this study, these census tract-level indicators were used to calculate the need at each census block (N_i^{CB}). However, this required a multi-step spatial data preparation process to reconcile differences in geographic boundaries and spatial resolution between the FARA tract-level data and the 2020 census block group geography.

The first challenge was that FARA variables are reported at the 2010 census tract level, but the data were compiled in 2019. This study used the 2020 census tract and block group geographies because population estimates are more recent. To address this inconsistency, areal interpolation was performed to redistribute 2010 tract-level FARA variables onto 2020 tract boundaries. The intersection of 2020 Census block group boundaries with 2010 Census tract polygons was computed, and 2010 Census tract variable values were distributed to 2020 Census tracts proportionally based on the fractional area of overlap.

The second challenge involved classifying each census tract as urban or rural to parameterize the Huff model's distance sensitivity. Although FARA included this distinction, the classification is not available for 2020 Census block group boundaries. Tracts were classified using a spatial overlay approach: each 2020 tract polygon was intersected with the 2020 Census Urban Area boundaries, obtained from the U.S. Census Bureau TIGER/Line shapefile repository. A tract was designated as urban if more than 50% of its land area overlapped a delineated urban area polygon; otherwise, it was designated as rural.

Figure 1 illustrates the result of the process explained above. The spatial distribution of population centers reveals pronounced geographic and socioeconomic contrasts between urban and rural census tracts, which directly inform the parameterization of the distance-decay model.

These tracts exhibit the highest low-income population values in the study area, with several block groups exceeding 2,000 low-income residents, reflecting the well-documented concentration of food insecurity in urban centers. Given the shorter travel distances typical of urban food environments, these tracts are assigned a distance-decay parameter of $\theta(j) = 2.0$ km, reflecting the expectation that urban residents access food assistance providers within a relatively compact service radius. In contrast, the rural tracts are in the western portions of County 1 and the entirety of County 2. These areas have a lower total population and lower absolute counts of low-income residents per block group. However, the geographic dispersion of these demand points across a substantially larger land area poses a distinct access challenge: rural residents travel farther to reach the nearest food pantry or meal program. Therefore, we used $\theta(j) = 5.0$ for rural tracts, acknowledging that rural populations must travel farther to access food assistance. $\beta = 1.5$ was used for both urban and rural areas.

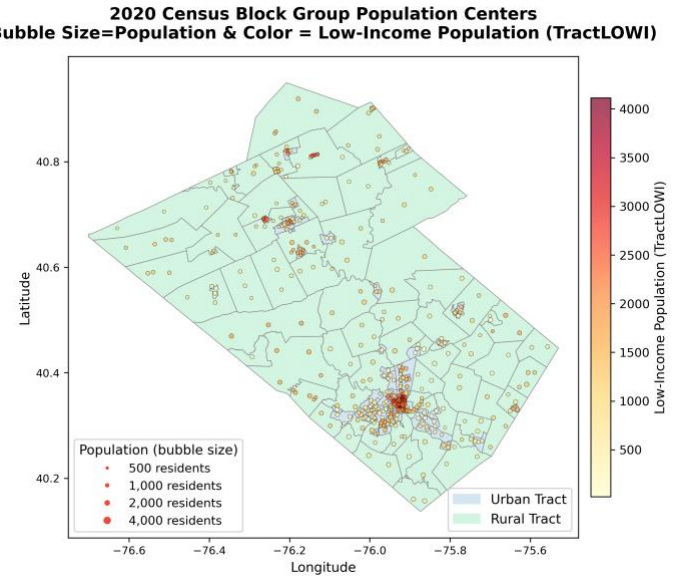


Figure 1: Population-weighted centroids of 2020 Census Block Groups in Counties 1 and 2. Bubble size represents the total block group population, and bubble color indicates the low-income population count.

To estimate the marginal cost of delivery for each agency, the primary operational dataset used in this study was obtained from the food bank's internal order management system, comprising all agency delivery records (weight and date) for 2024. These data provided the demand and delivery schedules for CVRPTW. Agency-level spatial and administrative attributes were sourced from an internal food bank directory and public sources to complement the agencies' locations. Routing distances between all pairs of agencies in the network were calculated using the Google Maps API. This matrix was used to construct day-specific distance matrices for the CVRPTW formulation. The food bank operates a heterogeneous fleet of delivery vehicles, with capacities of either 5,000 or 10,000 pounds. Fixed deployment costs were assigned to each vehicle to reflect the operational cost of activating a truck for a given route, independent of distance traveled. All vehicles were assumed to operate within a 10-hour shift window, with individual route durations capped at 5 hours. Using these data, a CVRPTW was solved for each operation day in 2024 to find routes and calculate Δ_j as discussed in the framework.

3.2 Results and Implications for PI

Figure 2 presents the relationship between marginal routing distance (Δ_j) and need captured (N_j) across all agencies in the study area, where each bubble represents an agency and bubble size represents total weight distributed in 2024 on a log scale. This result shows that agencies requiring greater marginal routing effort tend to capture less need from the surrounding population, consistent with the map in Figure 1. This relationship is statistically significant based on a power-law regression fit estimated on the log-log scale (a slope of -0.98 with an R^2 of 0.19).

Several notable patterns emerge from this analysis. The upper-left quadrant contains agencies with very low marginal routing distances ($\Delta_j < 1$ km) and high need captured ($N_j > 600$). These agencies are predominantly located within or immediately adjacent to high-density, high-need urban block groups in County 1's metropolitan area, where the concentration of low-income populations is greatest. These agencies represent the most spatially efficient service

outlets in the network, delivering high food access value at minimal additional routing cost. Therefore, these are high-need low-distance agencies.

The right portion of the figure ($\Delta_j > 10$ km) contains agencies with substantially higher marginal routing distances, yet moderate need-capture values ($N_j \approx 300$–500). Several of these agencies are among the largest in terms of total weight distributed, suggesting they play important logistical roles in the network despite their higher routing costs. These are rural agencies serving geographically dispersed populations where no closer alternative exists. We will call this group moderate-need, high-distance agencies.

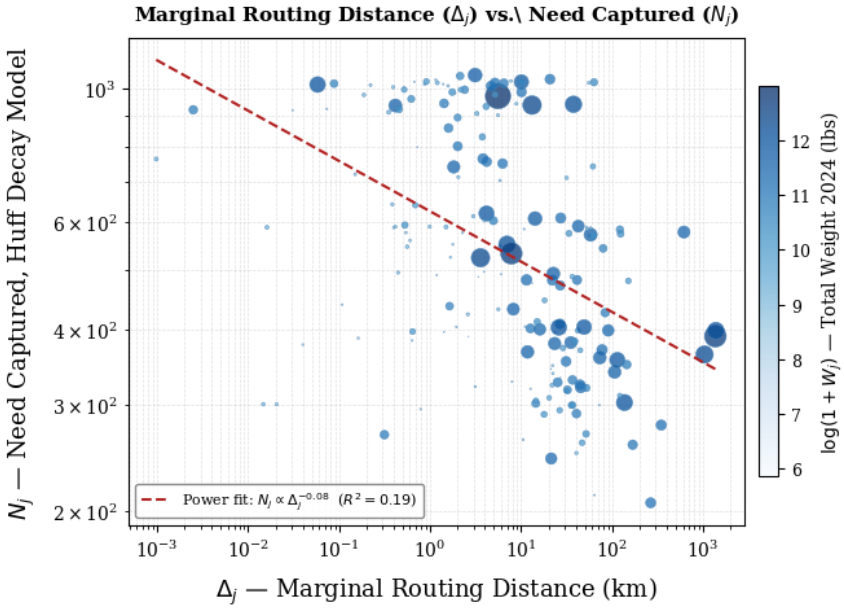


Figure 2: The relationship between marginal routing distance (Δ_j) and need captured (N_j) across all food bank agencies in the study.

A cluster of large, dark bubbles in the center of the figure ($\Delta_j \approx 5$–20 km, $N_j \approx 500$–$1{,}000$) represents high-throughput agencies that both capture substantial need and distribute large volumes of food. These agencies appear to occupy an efficient middle ground in the network, balancing routing cost and service impact.

Under PI principles, these findings directly motivate the design of a denser, more distributed network of transshipment hubs and relay points that would reduce the marginal routing distance for agencies currently operating in the high-Δ_j regime by enabling food to be transferred closer to the point of need rather than routed directly from a single central warehouse. The current distribution model, operating from a single central warehouse, results in inefficient long-haul routes to rural agencies, which could instead be served through intermediate consolidation points or by directly shipping food from donors. In Figure 2, the cluster of high-throughput agencies in the mid-range of both axes ($\Delta_j \approx 5$–20 km, $N_j \approx 500$–$1{,}000$) can serve intermediate nodes that receive consolidated shipments from the central warehouse or donors and redistribute smaller loads to surrounding lower-volume agencies. In addition, low-distance, high-volume agencies are candidates for PI last-mile nodes.

Operational decisions should not be based on efficiency alone. To synthesize the equity and operational dimensions of the analysis, each agency was positioned in a two-dimensional performance space defined by its marginal routing distance (Δ_j) on the vertical axis and the

number of neighboring agencies within $d = 3$ miles (n_j) on the horizontal axis. Agencies were then classified into four quadrants using median splits on both dimensions: *High Efficiency,* Dense / High Cost, *Isolated / Low Cost, and Inefficient*. As Figure 3 shows, agencies classified as "Inefficient," with high marginal routing distance and few neighbors, can exhibit the highest need-to-delivered ratios, where W_j is the actual total pounds of food delivered to agency j. These agencies serve populations with significant unmet needs despite their high operational costs. Removing or consolidating these agencies without accounting for the communities they serve could worsen food access in already underserved areas. This also justifies using PI principles to increase efficiency without compromising equity.

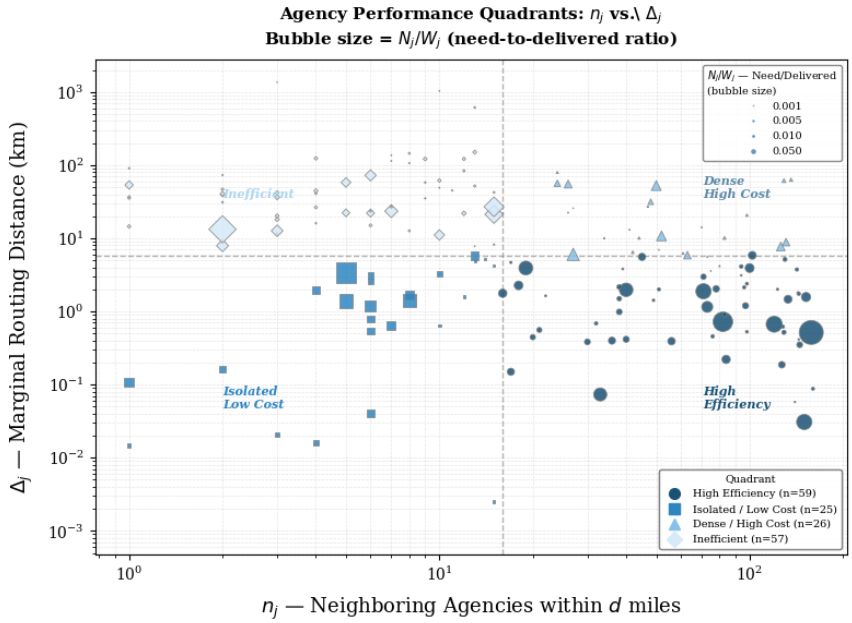


Figure 3: *Agency performance quadrants defined by the number of neighboring agencies within d=3 miles (n_j, x-axis) and marginal routing distance (Δ_j, y-axis). Bubble size represents the need-to-total delivery ratio (N_j/W_j) for agency j.*

4 Conclusion

This study proposed and implemented a six-step decision-support framework for evaluating partner agencies within a charitable food distribution network, integrating need estimation, spatial analysis, and operational cost modeling to support equity-informed planning. The framework was applied in two counties in the Northeast US. The framework combined census block-level need estimates derived from USDA Food Access Research Atlas indicators with a Huff-type spatial allocation model and a capacitated vehicle routing problem with time windows (CVRPTW) to assess both the equity and efficiency of the existing agency network. The results revealed a statistically significant inverse relationship between marginal routing distance and need captured, indicating that agencies located in high-density urban areas serve the greatest need at the lowest operational cost, while rural agencies incur substantially higher delivery burdens relative to the need they capture. These findings carry direct implications for Physical Internet (PI)-inspired network redesign: high-throughput, mid-range agencies emerge as natural candidates for intermediate consolidation nodes, while low-distance, high-volume urban agencies are well positioned as last-mile distribution points in a more open, hyperconnected logistics architecture. By drawing on PI principles of shared infrastructure, modular flows, and open interconnectivity, this work demonstrates how charitable food distribution networks can be restructured to better align service provision with estimated need,

reduce routing inefficiencies, and improve equity across geographically diverse service areas. Despite these contributions, several limitations should be acknowledged. First, the need estimation relies on census-tract-level indicators, downscaled to the block-group level via population-proportional allocation, which assumes uniform within-tract distributions and may obscure localized concentrations of food insecurity. Second, the Huff model parameters were calibrated using general urban-rural classifications rather than empirically observed client travel behavior, introducing uncertainty into the spatial allocation of need. Third, the marginal delivery cost estimation uses a route-based decomposition rather than full re-optimization upon agency removal, providing a computationally efficient yet locally approximate measure of routing burden. Fourth, the analysis is limited to a single food bank service area and a single calendar year of operational data, which may limit the generalizability of findings to other regions or time periods. Future work should address these limitations by incorporating client-level survey data, refining model parameters through empirical calibration, and extending the framework to multi-region or multi-food-bank settings to validate the broader applicability of PI-inspired network redesign in humanitarian logistics contexts.

References

- Akkerman, R., Buisman, M., Cruijssen, F., de Leeuw, S., & Haijema, R. (2023). Dealing with donations: Supply chain management challenges for food banks. *International Journal of Production Economics*, *262*, 108926. https://doi.org/10.1016/j.ijpe.2023.108926
- Coleman-Jensen, A. J. (2010). US food insecurity status: toward a refined definition. *Social Indicators Research*, *95*(2), 215-230.
- Grest, M., Lauras, M., & Montreuil, B. (2019). Toward humanitarian supply chains enhancement by using physical internet principles. 2019 International Conference on Industrial Engineering and Systems Management (IESM),
- Iqbal, T., & Ahmad, S. (2022). Transparency in humanitarian logistics and supply chain: the moderating role of digitalisation. *Journal of Humanitarian Logistics and Supply Chain Management*, *12*(3), 425-448. https://doi.org/10.1108/jhlscm-04-2021-0029
- Islam, M. H., & Ivy, J. S. (2022). Modeling the role of efficiency for the equitable and effective distribution of donated food. *OR Spectrum*, *44*(2), 485-534. https://doi.org/10.1007/s00291-021-00634-z
- Lowrey, J., & Boyer, K. (2024). Pantry direct: A transformative supply chain for reducing food insecurity. *Journal of Business Logistics*, *45*(1), e12341. https://doi.org/10.1111/jbl.12341
- Sengul Orgut, I., Brock Iii, L. G., Davis, L. B., Ivy, J. S., Jiang, S., Morgan, S. D., Uzsoy, R., Hale, C., & Middleton, E. (2016). Achieving Equity, Effectiveness, and Efficiency in Food Bank Operations: Strategies for Feeding America with Implications for Global Hunger Relief. In C. W. Zobel, N. Altay, & M. P. Haselkorn (Eds.), *Advances in Managing Humanitarian Operations* (pp. 229-256). Springer International Publishing. https://doi.org/10.1007/978-3-319-24418-1_11
- Sengul Orgut, I., & Lodree, E. J. (2023). Equitable distribution of perishable items in a food bank supply chain. *Production and Operations Management*, *32*(10), 3002-3021. https://doi.org/https://doi.org/10.1111/poms.14019
- Service, U. S. D. A. E. R. (2025). *Food Access Research Atlas - Download the Data*. Retrieved 2025-06-22 from https://www.ers.usda.gov/data-products/food-access-research-atlas/download-the-data
- Stauffer, J. M., Vanajakumari, M., Kumar, S., & Mangapora, T. (2022). Achieving equitable food security: How can food bank mobile pantries fill this humanitarian need. *Production and Operations Management*, *31*(4), 1802-1821. https://doi.org/10.1111/poms.13663
- Zou, F., Sharma, L., Pekgün, P., & Ahire, S. L. (2025). Reducing Food Waste in U.S. Food Banks: The Role of Distribution Strategies. *Production and Operations Management*, *34*(11), 3667-3685. https://doi.org/10.1177/10591478251345133

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE
Physical Internet Center

Orchestrating Sovereign Logistics Networks Through Performance Commitments in the Physical Internet

Sahrish Jaleel Shaikh[1], Benoit Montreuil[1], Alan Erera[1], Tiankuo Zhang[1], Malte Spanuth[2], Paria Nourmuhammadi[1]

1. Georgia Institute of Technology, Atlanta, USA
2. Kühne Logistics University, Hamburg, Germany

Corresponding author: sahrish.shaikh@gatech.edu

Abstract: *The Physical Internet promises transformative gains in logistics efficiency and sustainability through open, interconnected networks. Yet large-scale adoption remains hindered by a persistent tension: effective coordination appears to demand operational integration, conflicting with the competitive sovereignty that defines logistics markets. This paper argues the impasse arises from a misdiagnosis - the barrier is not that coordination requires operational integration, but that prevailing architectures demand full data transparency that competitive networks are unwilling to provide. Drawing on incomplete contract theory, we propose a performance commitment framework wherein networks commit to verifiable outcomes at network interfaces while retaining complete sovereignty over internal operations. Sovereignty spans three dimensions: data sovereignty, decisional sovereignty and asset sovereignty. The framework directly addresses data sovereignty through graduated disclosure levels and preserves decisional sovereignty as an invariant; asset sovereignty lies outside the protocol's scope. We formally define performance commitments and network data sovereignty as the coordination primitives and introduce a graduated data sovereignty framework supporting three participation levels - each defined by a distinct boundary on what operational data exits the network. The framework establishes a theoretical foundation for commitment-based coordination in the Physical Internet, with direct implications for protocol design, governance structures, and empirical validation.*

Keywords: *Physical Internet; Network sovereignty; Logistics coordination; Performance commitments; Network-of-networks architecture; Governance; Hyperconnected Logistics; Interoperability; Orchestration; Incomplete contracts*

Physical Internet (PI) Roadmap Fitness*: ☒ PI Networks, ☒ System of Logistics Networks, ☒ Vertical Supply Consolidation, ☒ Access and Adoption, ☒ Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☒ In-Person presentation*

1. Introduction

The Physical Internet envisions a fundamental transformation of logistics systems from closed, vertically integrated networks to open, interconnected systems capable of achieving higher efficiency, greater resilience, and reduced environmental impact (Montreuil, 2011). Central to this vision is interoperability: heterogeneous logistics networks operated by competing firms must coordinate effectively to enable seamless, multi-hop freight flows across organizational boundaries (Pan et al., 2021). Despite

widespread recognition of the PI's potential and expanding field pilots and ventures, progress toward large-scale implementation remains hindered by a persistent tension: effective coordination appears to demand operational integration, yet such integration conflicts with the competitive sovereignty that defines logistic markets.

In this paper, we argue that the impasse arises from a conceptual gap in what coordination fundamentally requires. Prevailing PI coordination architectures equate coordination with joint optimization, demanding full data transparency as a precondition. But coordination does not require shared decision-making or transparency into internal operations - it requires only that actors form reliable expectations about externally observable behavior at network interfaces. Early work on networked manufacturing revealed that selective information sharing at firm boundaries can yield substantial scheduling coordination gains without exposing internal operational logic (D'Amours et al., 1999). The barrier to PI adoption is not that coordination demands operational integration. It is that prevailing architectures demand full data transparency (Zheng, 2019; Leung 2022), which networks operating in competitive markets are unwilling to provide.

Building on incomplete contract theory (Hart & Moore, 1990; Hart, 1995), we propose a performance commitment framework wherein networks commit to verifiable outcomes at network boundaries while retaining complete data sovereignty over internal operations. We introduce graduated data sovereignty - three participation levels each defined by a distinct boundary on what operational data exits the network - enabling heterogeneous networks to self-select their coordination depth while staying interoperable through a common commitment language. Hereafter, the term container refers to the family of modular PI packaging, handling, and transport containers introduced by Montreuil et al. (2016).

The paper introduces a foundational set of conceptual and theoretical contributions to Physical Internet literature. First, we reframe the PI coordination problem from joint optimization under shared information to commitment-based coordination under protected data sovereignty. Second, we formally define performance commitment and network data sovereignty as the coordination primitives the PI literature has lacked. Third, we develop a graduated data sovereignty framework showing that heterogeneous disclosure depths within a common protocol structure strengthen rather than fragment the ecosystem.

Sovereignty in logistics networks spans three dimensions: *data sovereignty* (what operational information exits the network), *decisional sovereignty* (who controls internal operational decisions such as routing, consolidation, and pricing), and *asset sovereignty* (who owns and controls physical infrastructure), adapted from a decomposition developed in the digital sovereignty literature, where scholars of cloud computing and digital infrastructure have distinguished analogous concerns: control over what information exits organizational boundaries, control over operational processes and infrastructure availability, and control over digital assets including software and content (Flinders, 2024). Translating this decomposition to physical logistic networks, data sovereignty shifts from protecting organizational data in cloud environments to protecting competitive operational intelligence from exposure through network participation; decisional sovereignty shifts from ensuring digital infrastructure availability to preserving an actor's authority over internal routing, consolidation, and pricing decisions in the face of orchestrator influence; and asset sovereignty shifts from control over software and digital infrastructure to control over tangible physical assets through which logistics flows are executed. The framework directly addresses data sovereignty through graduated disclosure levels and preserves decisional sovereignty as an invariant across all participation levels - the commitment to

outcomes rather than processes is precisely a decisional sovereignty guarantee. Asset sovereignty lies outside the coordination protocol scope; networks keep full ownership and control over physical infrastructure regardless of participation level.

The paper is structured as follows. Section 2 reviews related works. Section 3 develops theoretical foundations. Section 4 presents the graduated commitment framework. Finally, Section 5 discusses implications and limitations.

2. Background and Related Work

Substantial PI research addresses inter-network coordination through optimization-based architectures wherein networks share operational information and coordinate through joint planning, shared routing algorithms, or network-wide consolidation strategies (Hakimi et al., 2012; Sarraj et al., 2014; Fazili et al., 2017). These approaches demonstrate theoretically achievable efficiency frontiers under full coordination. However, they may require networks to reveal cost structures, routing algorithms, and capacity positions to a central optimizer, creating adoption barriers that the literature increasingly recognizes as the primary obstacle to PI implementation (Ballot et al., 2021). The present paper does not contest the efficiency arguments of optimization-based approaches; it addresses their institutional feasibility.

The PI adoption literature consistently identifies data exposure, not loss of operational control, as the central barrier to network participation. Logistic firms guard routing algorithms and cost structures as key competitive assets and resist revealing them to competitors or central orchestrators (Ballot et al., 2021; Pan et al., 2021). Recent work has begun proposing data sovereignty and data space frameworks as enabling mechanisms for PI-style and inter-organizational data sharing, where interoperability rests on usage-controlled, policy-based exchange rather than unrestricted disclosure (Dalmolen et al., 2019). Data spaces and infrastructural sovereignty models explicitly aim to let participants specify and enforce how data may be used via standardized policies and connectors, so that cross-organizational collaboration becomes acceptable without requiring maximal transparency (Jung & Dörr, 2022). Our framework formalizes this intuition for the Physical Internet and extends it into a graduated, protocol-level architecture that articulates PI participation as a hierarchy of minimally required, sovereignty-preserving disclosures rather than an all-or-nothing choice.

Incomplete contract theory demonstrates that when contracts cannot specify all contingencies ex ante, as is inevitable in dynamic logistics environments, efficient governance allocates residual control rights to parties with superior local information while contracting only on verifiable outcomes (Grossman & Hart, 1986; Hart & Moore, 1990; Hart, 1995). Applied to supply chain coordination, this implies contracting on observable performance outcomes rather than internal processes (Hart & Moore, 2007; Aghion & Holden, 2011). We apply this logic to PI network interfaces: commit on verifiable outcomes, protect everything else as data sovereignty.

3. Coordination, Commitments, and Incomplete Contracts

3.1 Coordination as Expectation Formation

Coordination in logistic systems arises from interdependencies among actors whose decisions affect one another through shared infrastructure, temporal coupling of flows, and capacity constraints. When a shipment moves from Network A to Network B to Network C, each network's performance depends on

the others: When a shipment moves from Network A to Network B to Network C, each network's performance is interdependent, but not in a strictly sequential or information-complete sense. Network B does not need to wait for Network A to deliver before making commitments to Network C; rather, it forms expectations about upstream performance based on historical reliability, current system state, and observed signals. These expectations allow B to commit to downstream delivery - typically within a time window or with a specified robustness level - even under uncertainty. In practice, commitments are therefore made ex ante and are inherently probabilistic or buffered. What matters for coordination is not precise knowledge of when upstream actions will occur, but sufficiently reliable expectations about their distribution, so that downstream networks can plan consolidation, allocate capacity, and manage service commitments.

Crucially, coordination does not require joint decision-making or shared optimization. As established in the strategic framework for networked manufacturing (Montreuil et al., 2000), it requires actors to form reliable expectations about the behavior of others. A downstream network does not need to know which route an upstream network will choose, what trucks it will deploy, or how it sequences shipments. It needs confidence about when containers will arrive, what capacity will be available, and how deviations will be handled. This distinction between expectations about outcomes versus knowledge of processes is fundamental to what Physical Internet coordination actually requires.

Bilateral contracts and service-level agreements establish reliable expectations in traditional logistics, but scale poorly as the number of interacting parties grows and flows traverse multiple organizational boundaries. Most critically, bilateral mechanisms do not compose: a contract between A and B and a separate contract between B and C do not automatically create reliable end-to-end expectations from A to C. The Physical Internet addresses this through protocol-based coordination: each network conforms to a common standard, enabling any network to interact with any other without bilateral relationship-building - the same architectural insight that enables the digital internet to function at scale.

3.2 Incomplete Contracts and Residual Control Rights

In complex and uncertain environments, it is infeasible to specify ex ante all relevant contingencies and actions in a contract (Hart & Moore, 1988; Hart, 1988; Rossi, 2019). Logistics networks operate under pervasive uncertainty, making it practically impossible to anticipate the full range of contingencies affecting contractual performance (Al-Najjar, 2016). Contracts governing inter-network interactions are therefore necessarily incomplete.

When contracts cannot specify all contingencies, efficiency depends on two design choices: which variables are made contractable, and how residual control rights are allocated over non-contractable decisions (Grossman & Hart, 1986; Hart & Moore, 2007; Aghion & Holden, 2011). Variables verifiable by third parties belong in contracts; non-verifiable decisions requiring local information and real-time adaptation belong under the discretion of informed actors. When actors retain control over how outcomes are achieved, they capture returns from investments in superior processes and technologies - preserving innovation incentives that centralized control would weaken (Grossman & Hart, 1986; Wang, 2019).

Applied to logistic networks, externally observable outcomes are natural candidates for contractual commitment. Examples are whether a container arrived within a committed time window or whether committed capacity was delivered. Internal operational decisions such as route selection, consolidation timing, asset deployment, and algorithm choice require local information and benefit from real-time adaptation. Assigning residual control over these decisions to individual networks is efficient on both counts.

3.3 Implications for Physical Internet Coordination

These insights reframe the PI coordination challenge. The question is not how decisions should be optimized centrally, but which outcomes must be reliably committed to enable coordination. Coordination then emerges from decentralized decision-making constrained by those commitments.

From this perspective, an effective PI coordination architecture should exhibit three properties:

1. Definition of a set of standardized, externally verifiable performance commitments that allow downstream actors to plan reliably without inspecting upstream operations.
2. Allocation of residual control over internal processes to individual networks, enabling local adaptation and optimization while preserving innovation incentives.
3. Support of interoperability among heterogeneous actors by using a common commitment language, even if the depth of commitments varies across participants.

This framing naturally accommodates heterogeneity. Different networks face different levels of uncertainty, have different capabilities, and pursue different competitive strategies. A coordination mechanism that forces uniform integration or identical commitments is at risks of excluding participants or inducing inefficient behavior. By contrast, a commitment-based approach allows networks to self-select the strength of their commitments while remaining interoperable with others.

3.4 Coordination Primitives

Two concepts operationalize this theoretical foundation.

A *performance commitment* is a binding, verifiable promise to deliver a specified outcome at a network interface, independent of the internal process by which it is achieved. Each commitment specifies *a set of variables* (outcomes observable at the network boundary without inspecting internal operations such as delivery windows, capacity guarantees, price bands, GHG thresholds, handoff quality states), a *threshold* (target value or acceptable range, deterministic or distributional), a *verification mechanism* (independent infrastructure recording committed and actual values), and an *accountability consequence for each variable* (what occurs upon deviation). The protocol specifies only that commitments take this form, not how networks achieve them.

Network data sovereignty is the exclusive authority of a logistics network over what operational information it discloses beyond what its commitments require. The commitment establishes the disclosure floor; data sovereignty governs everything above it such as routing paths, cost structures, pricing algorithms, and demand patterns not part of the commitment. This directly addresses the adoption barrier the PI literature consistently identifies: networks resist PI participation not because they fear losing operational control, but because they fear data exposure to competitors (Ballot et al., 2021; Pan et al., 2021) endangering their competitiveness. The graduated concession of data sovereignty across the three levels in Section 4 is the mechanism by which the framework accommodates heterogeneous participation while preserving competitive integrity.

3.5 Key Actors in the Physical Internet

To ground the framework in the institutional reality of Physical Internet operations, we identify three actor types whose interactions the coordination protocol is designed to govern. This typology is not

exhaustive; the PI ecosystem will give rise to diverse business innovations and service configurations. It reflects, however, the principal roles that animate the framework developed here.

A ***client*** is any actor that requests logistics services from Physical Internet networks. Clients include shippers seeking freight movement, brands and e-commerce players requiring storage and fulfillment, and any other party whose demand generates freight flows through the PI ecosystem. The client is the ultimate beneficiary of the service commitment and the party against whom end-to-end performance is measured.

A ***network*** is a logistics entity with its own corporate identity, decisional sovereignty, and responsibility for freight within its geographic and/or functional domain. Networks are the sovereign actors whose coordination the framework governs. At its limiting case, a network may consist of a single node such as a hub or fulfillment center, treated as a special case of the general network definition.

An ***orchestrator*** is an actor that takes contractual responsibility for composing end-to-end freight journeys on behalf of clients, coordinating across multiple networks without necessarily owning physical infrastructure. Orchestrators are not a central authority; they are a type of market actor competing for end-to-end freight journey mandates on behalf of clients. A client with sufficient capability may internalize orchestration; alternatively, it delegates orchestration to an external orchestrator. A network with clients requesting transportation or deployment services not available within itself may internalize orchestration for its direct clients or alternatively require the services of an orchestrator for the external activities. When the term orchestrator appears in this paper, it thus refers to an actor taking an orchestration role, whether this actor is dedicated to orchestration roles, a network, or the client itself.

4. The Graduated Commitment Framework

Before detailing the three levels of the graduated sovereignty framework, we establish the core principle that unifies them: networks commit to what they will deliver (externally observable outcomes), not to how they will deliver it (internal operational processes). This principle manifests in two structural features present at all coordination levels:

- Contractable Performance Metrics: Commitments specify targets and thresholds for outcomes verifiable at network interfaces without inspecting internal operations. These include delivery time windows, price bands, capacity guarantees, as well as greenhouse gas emission and quality thresholds. Critically, these outcomes are what downstream networks and end customers fundamentally care about - they need to know when shipments will arrive and what they will cost, not what route was used or what algorithm computed the schedule. Networks retain sovereignty over internal operational decisions, subject to compliance with applicable legal, regulatory, and certified ethical standards.

- Residual Operational Control: Networks retain complete decision authority over all operational variables not explicitly committed to in the protocol. These include routing algorithms, consolidation strategies, asset allocation, exception handling methods, and pricing formulas. Because these decisions require private information about costs, capabilities, and real-time operational states, and because flexibility to adapt these decisions creates substantial value, assigning residual control to individual networks is efficient.

This implements the incomplete contracting logic directly: contract on verifiable outcomes where coordination value is high, assign residual control to informed parties where adaptation value is high. The result is performance-based sovereignty where networks are sovereign over methods but accountable for results.

4.1 Level 1: Full Operational Sovereignty

Level 1 represents the mandatory baseline for Physical Internet participation, preserving maximum operational autonomy while enabling multi-network coordination through endpoint verification. While respecting Physical Internet concepts and principles, networks retain complete decision-making authority over all internal operations - routing, scheduling, consolidation, resource allocation, and exception handling - while committing to verifiable performance outcomes at network boundaries.

Information disclosed: Networks commit to three verifiable boundary observations: pickup timestamp (e.g., when network accepts container), delivery timestamp (e.g., when network delivers container), and quality state at delivery (e.g., temperature logs, damage assessment, seal integrity). These endpoint observations, independently verified through sensors and/or blockchain-type certification, constitute complete external visibility into network operations.

Level 1 provides sufficient coordination for basic PI functionality through endpoint verification alone. Multi-network coordination does not require internal visibility: it requires only verified handoff points where delivering network responsibility ends and receiving network responsibility begins. Critically, network commitments are not fixed promises but window-conditional declarations. Each network states, at contract time, its delivery commitment conditioned on the arrival time window in which it receives freight. The orchestrator draws on historical performance data to construct, for each network, the deviation distribution around its declared commitments, capturing how frequently commitments are met and by how much they are typically exceeded. These distributions are then aggregated across the chain to derive a single robust delivery window to quote to the client. The client receives one commitment, and the penalty structure is defined against it.

As historical endpoint observations accumulate, the orchestrator constructs for each network a performance distribution around its declared commitments. A network declaring a 12-hour commitment may exhibit realized performance of 12.8h in average with a standard deviation of 1.4h for such a service. This distributional characterization requires nothing beyond timestamps at handoff points, yet it enables the orchestrator to quantify chain-level uncertainty and set a client-facing delivery window calibrated to a chosen service level rather than a best-case assumption.

Consider Network A transporting freight from Atlanta to Memphis, where Network B continues to Los Angeles. Network A declares a 12-hour commitment. Network B declares a 36-hour commitment if freight arrives in its standard inbound window, and 42 hours if freight arrives during high-congestion periods. The orchestrator, drawing on historical performance data for both networks and accounting for the probability that A's delays push freight into B's congestion window, quotes the client a 52-hour delivery window rather than the nominal 48, calibrated to be robust to plausible chain disruptions. If Network A takes 14 hours instead of its committed 12, freight arrives in B's congestion window and B's commitment automatically adjusts to 42 hours. If B performs within that commitment, no penalty applies to B. Network A is penalized for its 2-hour overrun, which also triggered the less favorable downstream condition. If total delivery falls within the quoted 52 hours, the orchestrator has fulfilled its commitment to the client. Attribution at every handoff requires only recorded arrival timestamps and pre-declared window-conditional commitments. Using this logic, no network requires visibility into any other actor's

operations. The logic applies to arbitrarily complex routes: five networks require only four handoff points, each creating a fault isolation boundary. The orchestrator identifies which segment caused delays without inspecting any network's internal operations. The logic extends cleanly to arbitrary network-of-networks depth.

4.2 Level 2: Performance-Observable Sovereignty

Level 2 represents a voluntary sovereignty configuration wherein networks disclose intermediate progression data (such as waypoint timestamps, routing paths, and segment-level timing) while retaining complete decision-making autonomy over operational methods.

Network incentives:

Networks adopt Level 2 when transparency enables premium pricing exceeding infrastructure and exposure costs. The core value is credible quality differentiation - progression data provides verifiable evidence of performance that Level 1's aggregate statistics cannot match. Progression-specific commitments create revenue opportunities. Networks can offer guaranteed routing through specific facilities, preferred routing with rerouting flexibility, or flexible network-optimized paths, each priced differently to capture value from heterogeneous client needs. The same waypoint infrastructure supports multiple service tiers, enabling price discrimination across client segments. The trade-off exists: infrastructure costs and competitive exposure (topology revelation, disclosed performance patterns) versus premium pricing markup, market differentiation, and operational insights. Networks with strong performance serving quality-conscious markets may adopt Level 2; cost-focused networks serving price-sensitive segments may decide to adopt Level 1.

Client benefits:

Real-time waypoint visibility enables early exception detection (by the orchestrator). Observing delayed progression at an early phase allows intervention, activating backup plans, adjusting operations, notifying recipients hours before delivery failures occur. In a transport context, path-specific options can help match service to freight requirements. For example, a pharmaceutical shipper may select certified-route guarantees; shippers of standard goods may choose preferred routing balancing cost and certainty, while shippers of commodities may accept flexible paths minimizing price.

Orchestrator and ecosystem benefits:

Segment-level temporal performance data enables substantially improved route composition and performance prediction. While Level 1 provides network-level performance distributions (as introduced in Section 4.1), it cannot account for temporal variation or current conditions within a network. Level 2's intermediate progression data decomposes these aggregate distributions into segment-level distributions conditioned on observable factors revealing, for example, that Network A's Birmingham hub processes freight in 1.5h at the median and 2.0h at the 99th percentile on Tuesday mornings, but 2.5h and 3.6h respectively on Friday afternoons, or that peak season increases processing times on average 60% over baseline. This data richness enables durable route composition as explored in Muthukrishnan et al. (2025). Orchestrators can compose segment-level distributions accounting for temporal factors, constructing routes that remain reliable under varying conditions rather than optimizing only for average-case scenarios. This condition-aware composition reduces prediction error and improves matching quality compared to Level 1's condition-blind aggregate statistics. Critically, segment-level visibility enables proactive measures: an orchestrator observing degraded performance at

a specific hub can reroute flows before breaches of commitment materialize, rather than attributing delays only after the fact at the endpoint.

4.3 Level 3: Proactive State Sovereignty

Level 3 represents a qualitative shift in both the timing and the nature of disclosure. Where Level 2 networks disclose container-level progression data - revealing where a specific shipment is within its journey - Level 3 networks continuously broadcast network-level operational state: hub utilization, queue lengths, processing backlogs, and equipment availability. This is a categorically different class of information. Level 2 answers the question of where a given container is; Level 3 answers the question of how the network itself is performing at any given moment, independent of any specific shipment. The data sovereignty concession at this level is therefore twofold: networks relinquish control over when operational information exits the network, and they expose the internal health of their operations rather than merely the progress of flows passing through them.

Information disclosed (in addition to Level 2): Networks continuously broadcast operational state: current hub utilization percentages and trends, real-time queue lengths and processing backlogs, available capacity across facilities, current equipment availability and location, active processing times reflecting actual current performance, developing capacity constraints before they impact delivery commitments, and exception events immediately upon detection. Update frequencies are established contractually and may take one of three forms. In the most sovereignty-preserving configuration, frequency is negotiated on a per-commitment basis: each time a network agrees to Level 3 disclosure for a specific flow or batch of flows, the update frequency is scoped to that commitment and dissolves when it ends. Alternatively, two parties may establish a bilateral standing agreement governing all flows between them over a defined period, durable and multi-commitment in nature, requiring explicit renewal or termination. At the least sovereignty-preserving end, a network may declare an ecosystem-wide broadcast policy, making its Level 3 update frequency available as a general declaration to any counterpart that triggers Level 3 coordination.

Network Incentives:

Networks adopt Level 3 when the benefits of continuous operational state visibility to their client justify the sovereignty concession. A network with available capacity receives freight assignments faster when its client can see that capacity in real time, capturing revenue before it flows to competing networks. Penalty exposure is reduced through early constraint signaling: the earlier a developing constraint is visible, whether through the network's own proactive broadcast or through the client's monitoring of the operational state, the more time remains for rerouting before a commitment is breached. The penalty impact of a breach scales inversely with warning time - a constraint flagged days ahead carries far lower consequences than one surfacing at the moment of failure or not at all. Continuous visibility also reduces the transaction cost of participation by replacing repeated bilateral query-response cycles with standing operational transparency. Over time, a network's Level 3 transparency may become a competitive differentiator: clients routing high-value or time-sensitive freight may preferentially select networks whose real-time visibility reduces coordination risk, making transparency itself a revenue-generating asset.

Client Incentives:

Clients contracting with Level 3 partnerships receive superior service reliability through orchestrator optimization across real-time operational state. Booking requests receive immediate capacity

confirmation based on current partnership conditions. Delivery time estimates reflect actual current operational state - when routing encounters elevated hub utilization or processing backlogs, estimates automatically adjust based on observed conditions. Exception recovery improves as the orchestrator rapidly identifies alternative capacity and coordinates redistributed routing based on real-time availability.

Orchestrator and ecosystem benefits:

Continuous operational state broadcasts transform orchestrator capabilities from reactive coordination to proactive system management through standing analytical visibility. The orchestrator maintains time-series data spanning all client facilities - hub utilization patterns, queue length distributions, processing time variability, equipment availability trends, and exception frequency - enabling sophisticated aggregate analytics. Real-time system-wide dashboards reveal client-level patterns: systematic bottlenecks, temporal dynamics, and directional imbalances enabling optimization opportunities. This analytical capability also supports strategic decisions including where to invest in collective capacity expansion, which corridors require additional networks, and how seasonal patterns necessitate dynamic capacity planning.

4.4 Commitment Deviation Monitoring and Reliability Intelligence

The graduated commitment framework generates a natural byproduct: for every fulfilled commitment it oversees, the orchestrator observes a committed value and an actual value for every commitment variable across the networks it engaged.

At every network interface, the verification mechanism produces a structured observation: the committed value c and the realized value r for each commitment variable k at time t for network n on route segment s. The signed scalar deviation $\delta = r - c$ captures direction and magnitude, with the interpretation of sign depending on the variable type: for time, positive deviation indicates late delivery; for price, positive deviation indicates overcharge; for emissions, positive deviation indicates threshold exceedance. In each case, the direction of underperformance is variable-specific and defined at contract time. For non-scalar commitment variables such as quality states or categorical compliance indicators, deviation is recorded as a structured comparison against the committed specification rather than a signed scalar. The orchestrator accumulates these into a time-indexed record:

$$D(n, s, k, t) = \{(c, r, \delta) : \text{all transits through segment } s \text{ of network } n \text{ on commitment variable } k, \text{ as observed by this orchestrator, up to time } t\}$$

It is important to note that in a PI ecosystem with multiple competing orchestrators, no single orchestrator observes all flows through a network. Each reliability profile is therefore partial and conditioned on the flow patterns that orchestrator has routed, meaning the same network may be assessed differently across orchestrators. Cross-orchestrator aggregation of reliability records would raise both data sovereignty and competitive intelligence concerns that fall outside this framework's scope. We flag this as an important open question for PI governance design, as open performance monitoring is a key Physical Internet principle (Montreuil, 2011).

4.4.1 Deviation Taxonomy

Orchestrators and clients classify deviations along two dimensions:

Random vs. systematic. Random deviations are zero-mean fluctuations consistent with genuine operational uncertainty. Systematic deviations exhibit persistent directional bias, for example a network consistently delivering 45 minutes later than committed is misrepresenting its capability. Detecting systematic deviation requires modest time-series depth and triggers commitment renegotiation or penalty escalation.

Stationary vs. structural. Stationary patterns are stable over time. Structural deviations represent regime shifts, i.e. significant distributional changes indicating capability transitions from infrastructure investment, capacity expansion, or operational degradation. Detecting structural breaks enables the orchestrator to update reliability estimates rather than extrapolating stale distributions.

4.4.2 Level-Differentiated Observability

The three coordination levels produce qualitatively different deviation records. **At Level 1,** the orchestrator observes only endpoint deviations, i.e. total deviation across the network's entire segment. **At Level 2,** progression timestamps decompose an endpoint deviation into segment-level contributions. **At Level 3,** continuous operational state broadcasts provide leading indicators that precede deviation events. The orchestrator accumulates not just (committed, actual) pairs but (operational state at t, subsequent deviation at $t+\Delta$) pairs, enabling regression of deviation magnitude on observable precursors.

The deviation record closes two feedback loops. Over time, empirically observed reliability distributions - mean deviation, variance, tail behavior, and temporal conditioning factors - replace network reported capability estimates, grounding route composition in actual rather than claimed performance. For incentive alignment, networks with consistent adherence receive preferential allocation while those with persistent systematic deviations face reduced allocation or eventual exclusion, making the deviation record a protocol-native trust mechanism that replaces bilateral relationship-building.

5. Discussion and Conclusion

The performance commitment framework resolves the PI coordination paradox by correctly identifying its source. The barrier to PI adoption is that prevailing architectures demand full data transparency, which competitive networks are unwilling to provide. By contracting only on verifiable outcomes and formally protecting everything else as data sovereignty, the framework makes coordination institutionally feasible without sacrificing theoretical coherence. Incomplete contract theory provides the normative foundation: outcomes are verifiable at interfaces and belong in contracts; internal decisions require local information and belong under network control.

The graduated data sovereignty structure has an important ecosystem property: heterogeneous participation strengthens rather than fragments the network-of-networks. A PI ecosystem containing Level 1, Level 2, and Level 3 networks is more robust than one requiring uniform disclosure depth, because it accommodates networks at different capability and trust stages while maintaining interoperability through the common commitment language. This supports incremental adoption as networks can enter at Level 1 and deepen participation as governance trust develops. Also prevented is

the fragmentation that uniform disclosure requirements would produce by excluding unwilling participants entirely.

The framework is conceptual and theoretical, opening several research and innovation avenues. Mechanism design questions arise around commitment renegotiation protocols, penalty calibration, and orchestrator governance structures. The deviation monitoring architecture introduced in Section 4.4 wherein the orchestrator accumulates committed versus actual performance observations into time-indexed reliability records warrants formal development as a trust infrastructure for PI ecosystems. Empirical validation through simulation-based assessment and PI field pilot testing of graduated participation dynamics are natural next steps.

References

1. Aghion, P. & Bolton, P., 1992. An incomplete contracts approach to financial contracting. The Review of Economic Studies, 59(3), pp.473–494.
2. Aghion, P. & Holden, R., 2011. Incomplete contracts and the theory of the firm: What have we learned over the past 25 years? Journal of Economic Perspectives, 25(2), pp.181–197.
3. Al Najjar, N.I., 2016. Incomplete contracts and the governance of complex contractual relationships. American Economic Review, 106(2), pp.1–30.
4. Asian, S. & Nie, X., 2014. Coordination in supply chains with uncertain demand and disruption risks: Existence, analysis, and insights. IEEE Transactions on Systems, Man, and Cybernetics: Systems, 44(9), pp.1139–1153.
5. Ballot, E. et al., 2021. Physical Internet: First results and next challenges. J. of Business Logistics, 42(1).
6. Christensen, H.B. et al., 2016. Accounting information in financial contracting: The incomplete contract theory perspective. Chicago Booth Research Paper No. 16 02.
7. D'Amours S. et al., 1999. Networked Manufacturing: Impact of Information Sharing on Scheduling Performance, International Journal of Production Economics, vol. 58, p.63-79.
8. Dalmolen, S. et al., 2019. Maintaining control over sensitive data in the Physical Internet: Towards an open, service oriented, network-model for infrastructural data sovereignty, IPIC 2019, 16 p.
9. Dubey, V., 2023. A comprehensive theoretical framework for sustainable network contracts: Contracting dimensions and contract classification. Heliyon, 9(12), e22660.
10. Fazili, M et al., 2017. Physical Internet, conventional and hybrid logistic systems: a routing optimisation-based comparison using the Eastern Canada road network case study. International Journal of Production Research, 55(9), pp.2703-2730.
11. Grossman, S.J. & Hart, O.D., 1986. The costs and benefits of ownership: A theory of vertical and lateral integration. Journal of Political Economy, 94(4), 691–719.
12. Hakimi, D. et al., 2012. Simulating a physical internet enabled mobility web: the case of mass distribution in France. In 9th Internet. Conference on Modeling, Optimization & SIMulation-MOSIM'12, 10p.
13. Hart, O. and Moore, J., 1988. Incomplete contracts and renegotiation. Econometrica: Journal of the Econometric Society, pp.755-785.
14. Hart, O. & Moore, J., 2007. Incomplete contracts and ownership: Some new thoughts. American Economic Review, 97(2), pp.182–186.
15. Hart, O. & Moore, J., 1990. Property Rights and the Nature of the Firm. J. of Political Economy, 98(6), 1119-1158.
16. Hart, O., 1988. Incomplete contracts and the theory of the firm. Journal of Law, Economics, & Organization, 4(1), pp.119–139.
17. Hart, O., 1995. Firms, contracts, and financial structure. Clarendon press.
18. He, C. et al., 2024. Enhancing resilience and coordination in sustainable logistics service supply chains through integrated investment contracts. Sustainability, 16(1), 1234.
19. Jung, C. and Dörr, J., 2022. Data usage control. In Designing data spaces: The ecosystem approach to competitive advantage (pp. 129-146). Cham: Springer International Publishing.

20. Leung, E. et al. 2022. From traditional warehouses to Physical Internet hubs: A digital twin-based inbound synchronization framework for PI-order management. International Journal of Production Economics. https://doi.org/10.1016/j.ijpe.2021.108353.
21. Montreuil B. et al. (2000). A Strategic Framework for Networked Manufacturing, Computers for Industry, v42, 299-317, https://doi.org/10.1016/S0166-3615(99)00078-0
22. Montreuil, B., 2011. Toward a Physical Internet: meeting the global logistics sustainability grand challenge. Logistics Research, 3(2), pp.71-87, https://doi.org/10.1007/s12159-011-0045-x
23. Montreuil B. et al. (2016). Modular Design of Physical Internet Transport, Handling and Packaging Containers, Progress in Material Handling Research V13, Ed. J. Smith et al., MHI, Charlotte, NC, U.S.A. https://digitalcommons.georgiasouthern.edu/pmhr_2014/1/
24. Mouzas, S. & Naidu, N.T.S., 2024. Overcoming contract incompleteness: Evidence from long-term supply relationships. Production Planning & Control, 35(3), pp.1–14.
25. Pan, S. et al., 2021. Digital interoperability in logistics and supply chain management: state-of-the-art and research avenues towards Physical Internet. Computers in industry, 129, p.103435.
26. Rossi, M.A., 2019. Incomplete contracts. In Encyclopedia of Law and Economics. Dordrecht: Springer.
27. Sarraj, R. et al., 2014. Interconnected logistic networks and protocols: simulation-based efficiency assessment. International Journal of Production Research, 52(11), 3185-3208.
28. Sun, Y. et al., 2025. Research on logistics service supply chain coordination in the context of green innovation. Sustainability, 17(2), 5678.
29. Wang, S., 2019. Incomplete contracts with disparity, uncertainty, information and incentives. Theory and Decision, 87(2), pp.213–235.
30. Zheng, L., Beem, P. and Bae, K.H.G., 2019. Assessment of the physical internet enabled urban logistics using agent-based simulation. International Journal of Logistics Systems and Management, 33(4), pp.441-466.

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE
ENCGT
المدرسة الوطنية للتجارة والتسيير
ENCG TANGER

From Digital Supply Chains to Network Intelligence in the Physical Internet: An Empirical Consolidation of Capability Pathways

Nadia HMITTI[1] and Ahmed MAGHNI[2]
[1]ENCG Tangier, Abdelmalek Essaâdi University, Morocco
[2]ENCG Tangier, Abdelmalek Essaâdi University, Morocco
Corresponding author: nadia.hmitti@etu.uae.ac.ma

Abstract: *This paper builds on a structured systematic literature review and interpretative consolidation of 36 empirical studies examining the relationship between digital supply chain (DSC) and performance. The study reconstructs and harmonizes intermediate explanatory mechanisms across existing empirical models to identify five recurrent functional domains through which digitalization affects supply chain performance: enabling (digital infrastructure and technological foundations), operational (process and efficiency mechanisms), coordination (integration and synchronization mechanisms), information (visibility and data-sharing mechanisms), and adaptive (flexibility, resilience, and viability mechanisms).*

The findings show that digitalization rarely operates as a direct performance driver. Instead, its impact is predominantly mediated by these intermediate functional domains, with adaptive capabilities emerging as the most frequently observed pathway. The results further reveal that existing empirical models remain largely fragmented, with limited systemic integration across capability functions.

Building on the Physical Internet (PI) perspective, the study reconceptualizes these functional domains as distributed intelligence functions embedded within interconnected logistics networks. In this interpretation, coordination and information capabilities support interoperability, synchronization, and hyperconnectivity across network actors, while enabling capabilities provide the technological and infrastructural foundations required for modular and interconnected logistics systems. Adaptive capabilities further contribute to resilience and dynamic responsiveness within complex network environments.

This shift moves the analytical focus from firm-level capabilities to network-level intelligence, emphasizing coordinated information processing, recursive coordination mechanisms, and adaptive responses as key enablers of operational, resilience-related, and sustainability-related performance.

By bridging empirical consolidation and conceptual structuring, the paper proposes a functional intelligence-based model of digital supply chains, providing an empirically grounded foundation for understanding how digital supply chains may progressively evolve toward interoperable, hyperconnected, and Physical Internet-compatible logistics systems.

Keywords: *Physical Internet, Digital Supply Chain, Network Intelligence, Supply Chain Performance, Interoperability, Coordination and Information.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☒ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

1 Introduction

The increasing complexity, volatility, and sustainability pressures faced by contemporary supply chains have accelerated the adoption of digital technologies across logistics and industrial systems. Technologies associated with Industry 4.0, such as artificial intelligence, the Internet of Things, big data analytics, and digital platforms, are increasingly deployed to enhance operational efficiency, responsiveness, and visibility within supply chains. As a result, digital supply chain has emerged as a central topic in both academic research and managerial practice.

Despite this growing interest, the relationship between digital supply chain and performance remains fragmented in the literature. Existing studies adopt heterogeneous conceptualizations of digitalization, ranging from specific technologies to broader notions such as smart supply chains or digital transformation. Similarly, supply chain performance is measured through multiple dimensions, including operational efficiency, resilience, and sustainability. While numerous studies report positive effects of digitalization, empirical evidence suggests that these effects are not systematic nor purely direct.

In particular, a growing body of research indicates that the impact of digitalization on performance is predominantly mediated by intermediate capabilities such as supply chain integration, visibility, flexibility, agility, and resilience. These capabilities enable organizations to process information more effectively, coordinate activities, and respond to environmental uncertainty. At the same time, several studies highlight the contingent nature of digitalization-performance relationships, emphasizing the role of organizational, environmental, and strategic factors.

Taken together, these findings suggest that digitalization should not be understood as a standalone technological driver of performance, but rather as an enabling foundation whose value materializes through organizational and operational capabilities. However, existing empirical models remain largely focused on firm-level or supply chain-level analyses and tend to examine these capabilities in isolation, without capturing their potential integration within broader system-level configurations.

In parallel, the Physical Internet (PI) has emerged as a paradigm that envisions a transition toward open, interconnected, and interoperable logistics networks. By promoting the interconnection of logistics systems through standardized protocols, modular containers, and shared infrastructures, the PI aims to achieve significant improvements in efficiency, sustainability, and resilience. Central to this vision are principles such as interconnectivity, interoperability, coordination, and information sharing across distributed actors and systems.

While the PI provides a systemic and network-oriented perspective on logistics transformation, it remains largely conceptual, with limited empirical grounding regarding how such interconnected systems can emerge from existing digital supply chains. In particular, the link between empirically observed digitalization capability pathways and the intelligence mechanisms required for network-level coordination remains underexplored.

This study addresses this gap by building on a structured consolidation of empirical research on digital supply chain and performance. It aims to identify and organize intermediate functional capability domains through which digitalization contributes to performance and reframe these domains as distributed intelligence functions at the network level. Specifically, the study proposes a conceptual shift from firm-level capability pathways to distributed intelligence functions, providing a structured bridge between empirical digital supply chain research and the systemic logic of the Physical Internet.

By doing so, the study contributes to advancing the understanding of how digital supply chains may evolve toward interconnected and intelligent logistics networks, offering both theoretical insights and a foundation for future research on Physical Internet-compatible systems.

To address this objective, the study builds on a structured systematic literature review and interpretative synthesis of empirical studies examining the relationship between digital supply chain and performance.

2 Literature Background

The Physical Internet (PI) has emerged as a paradigm-shifting concept aiming to redesign logistics and supply chain systems through the interconnection of logistics networks. It is commonly defined as an open global logistics system founded on physical, digital, and operational interconnectivity, enabled through standardized encapsulation, interfaces, and protocols (Pan et al., 2021; Treiblmaier et al., 2020). Inspired by the Digital Internet, the PI envisions a transition from fragmented and independently managed supply chains toward a "network of logistics networks," where goods flow seamlessly across interconnected systems (Sarraj et al., 2014a; Sternberg & Norrman, 2017).

At its core, the PI aims to significantly improve efficiency, sustainability, and overall logistics performance by enabling shared infrastructures, coordinated flows, and collaborative resource utilization. This objective relies on a set of core principles consistently highlighted in the literature, including interoperability, interconnectivity, coordination, and information sharing.

Interoperability refers to the ability of heterogeneous systems, organizations, and logistics networks to operate jointly and exchange data seamlessly, while interconnectivity emphasizes the creation of globally connected logistics systems. Coordination mechanisms are central to aligning flows and decisions across decentralized actors, often supported by protocols and shared rules (Sarraj et al., 2014b). Information sharing and visibility enable real-time tracking, transparency, and data exchange across the network, forming the informational backbone of the system. Modularity and standardization, particularly through the use of standardized containers, allow goods to be handled independently of their content, facilitating seamless transfer across nodes. Finally, collaboration, both vertical and horizontal, is required to enable resource pooling and collective optimization across the network.

These principles are predominantly conceptualized at the supply chain and network levels, with several contributions explicitly framing the PI as a system-level transformation toward hyperconnected logistics infrastructures (Tran-Dang et al., 2020; Tran-Dang & Kim, 2021). This network-oriented perspective is further reinforced by architectural and technological developments within the PI literature. Existing studies highlight the importance of layered architectures, service-oriented systems, and protocol-based coordination mechanisms, as well as the integration of enabling technologies such as the Internet of Things, cyber-physical systems, artificial intelligence, and multi-agent systems.

In this context, the PI can be understood as a distributed and intelligent system, where decision-making, coordination, and information processing are embedded across multiple interconnected entities. Intelligence is not centralized but distributed across containers, hubs, and platforms, enabling real-time responsiveness, autonomous decision-making, and adaptive system behavior (Sallez et al., 2015; Tran-Dang et al., 2020).

Despite these advances, the PI literature remains largely conceptual, with limited empirical grounding regarding how such system-level capabilities emerge from existing supply chain practices. In particular, while the principles of interoperability, coordination, and visibility are

well established, their operationalization through concrete capability pathways remains underexplored.

This gap highlights the need to connect the systemic vision of the Physical Internet with empirical insights from digital supply chain research. In this regard, recent studies on digital supply chain provide evidence on how digital technologies contribute to performance through intermediate capabilities such as integration, visibility, flexibility, and resilience. However, these studies remain primarily focused on firm-level or supply chain-level analyses and do not explicitly address the transition toward network-level systems.

This study builds on this gap by linking empirically identified capability pathways with the systemic principles of the Physical Internet. By doing so, it contributes to bridging the divide between conceptual PI frameworks and empirical research on digital supply chains, offering a foundation for understanding how digital supply chain capabilities may evolve into distributed intelligence functions within interconnected logistics networks.

Building on these theoretical and conceptual foundations, the following section presents the methodological approach adopted to consolidate and analyze the empirical literature on digital supply chain and performance.

3 Methodology

This study builds on a structured systematic literature review and interpretative synthesis of empirical studies examining the relationship between digital supply chain (DSC) and performance. The objective was not to statistically aggregate effect sizes, but rather to identify, compare, and structurally consolidate recurring intermediate capability pathways through which digitalization contributes to supply chain performance.

The literature search was conducted using major academic databases, namely Scopus and Web of Science, and focused on empirical studies addressing digital supply chain, supply chain digitalization, smart supply chains, or Industry 4.0-related supply chain transformations in relation to performance outcomes. The selection process followed predefined inclusion criteria, prioritizing empirical studies explicitly modeling intermediate or mediating mechanisms between digitalization and performance.

A total of 36 empirical studies were retained for the consolidation process (e.g., Lee et al., 2023; Haddud & Khare, 2020; Hallikas et al., 2021; Yu et al., 2024; Cheng et al., 2022). The selected studies included heterogeneous empirical models involving operational, resilience-related, and sustainability-related performance dimensions. Given this heterogeneity, the study adopted an interpretative and comparative synthesis approach rather than a statistical meta-analytic synthesis.

The analysis involved iterative coding, cross-study comparison, and harmonization of recurrent mechanisms across the empirical literature. Through this process, the identified intermediate mechanisms were consolidated into five recurrent functional domains: enabling, operational, coordination, information, and adaptive capabilities.

The resulting framework was subsequently reinterpreted through the lens of the Physical Internet (PI) in order to examine how these capability pathways may contribute to network-level intelligence, interoperability, hyperconnectivity, and interconnected logistics systems.

4 Results and Discussion

[redacted]Functional Capability Domains

The empirical consolidation of the selected studies reveals that the relationship between digital supply chain and performance is structured around a set of intermediate functional capabilities. These capabilities operate as explanatory mechanisms through which digitalization translates into operational, resilience-related, and sustainability outcomes.

4.1.1 Functional hierarchy of capability domains

The analysis of frequency distribution across 30 empirical studies involving mediation effects highlights a differentiated functional hierarchy among capability domains.

Adaptive capabilities emerge as the dominant mechanism (n = 11). The literature predominantly mobilizes constructs related to resilience, flexibility, agility, and viability. Digitalization is therefore primarily conceptualized as a lever of adaptability, enabling supply chains to respond to uncertainty, disruptions, and environmental turbulence.

Operational capabilities represent the second most frequent domain (n = 9). These mechanisms are associated with process improvement and operational excellence, including lean practices, green supply chain initiatives, procurement optimization, and process efficiency. In this perspective, digitalization contributes to performance through the enhancement of operational practices and resource utilization.

Coordination and information capabilities appear at a structural level (n = 7 each). Although less frequent individually than adaptive capabilities, they form a significant functional layer related to inter-organizational integration, visibility, and information sharing. These capabilities support synchronization across supply chain actors and enable coordinated decision-making.

Enabling capabilities, associated with technological infrastructure and digital tools, remain marginal as intermediate mechanisms (n = 3). This finding indicates that technology itself is rarely modeled as a mediator. Instead, it functions as an enabling layer whose value materializes through higher-level capabilities.

Overall, this hierarchy suggests that digitalization is not directly linked to performance outcomes, but operates through distinct functional capability domains, with a strong emphasis on adaptive and operational mechanisms.

4.1.2 Structure of capability combinations

The analysis of co-occurrence patterns provides further insights into how these capability domains are combined within empirical models.

The most significant combination is observed between coordination and information capabilities (n = 3), indicating that integration and visibility are frequently modeled together as complementary mechanisms. This pairing reflects an underlying network logic, where information sharing and coordination jointly support performance outcomes.

In contrast, adaptive capabilities, despite being the most frequent, are predominantly modeled in isolation. Although present in 11 studies, they are combined with other domains only in a limited number of cases. This suggests that empirical models often focus on direct pathways such as digitalization → resilience/flexibility → performance, without integrating coordination or information mechanisms within the same framework.

A similar pattern is observed for operational capabilities, which are also largely treated as independent mechanisms. Operational improvements are typically modeled as standalone pathways linking digitalization to performance.

Enabling capabilities are not combined with other domains, reinforcing the idea that technological elements are not conceptualized as part of an integrated functional system but rather as independent or upstream variables.

4.1.3 Emerging structural configuration

The combined analysis of frequency and co-occurrence patterns reveals a structure of largely independent capability pathways.

First, there is no strong evidence of integrated multi-functional architectures within the empirical models. Instead, the literature is characterized by relatively independent mechanisms, each explaining a specific pathway from digitalization to performance.

Second, three dominant pathways can be identified:

- An adaptive pathway, where digitalization enhances resilience, flexibility, and adaptability, leading to improved performance.
- An operational pathway, where digitalization improves operational practices and process efficiency.
- A network-oriented pathway, based on the joint role of coordination and information capabilities, enabling integration and visibility across supply chain actors.

This configuration highlights that while digitalization activates multiple capability domains, these are rarely integrated within a unified framework. The empirical literature thus reflects a fragmented but structured set of pathways, where capabilities operate as distinct building blocks rather than components of a systemic architecture.

From a broader Physical Internet perspective, these pathways may also be interpreted as complementary network-level functions supporting interoperability, coordination, visibility, adaptability, and interconnected logistics operations among distributed supply chain actors.

Overall, the findings confirm that digital supply chain primarily affects performance through intermediate capabilities rather than through direct effects, thereby supporting a broader interpretation beyond firm-level analysis.

From Capability pathways to Network Intelligence in the Physical Internet

The results of the empirical consolidation show that digital supply chain affects performance primarily through intermediate functional capabilities rather than through direct technological effects. From a Physical Internet (PI) perspective, these findings can be reinterpreted at a broader analytical level.

Consistent with the literature reviewed earlier, the Physical Internet is conceptualized as an open global logistics system based on the interconnection of logistics networks through standardization, modular containers, interfaces, and protocols, with the objective of improving efficiency and sustainability. This vision is commonly associated with three core properties: interconnectivity, interoperability, and coordination across networks of logistics services. Rather than focusing on isolated firms or technologies, the PI literature emphasizes a network-level logic in which multiple actors, systems, and flows must operate in a synchronized and information-rich environment.

This interpretation is particularly consistent with the structural patterns identified in the empirical consolidation. From a Physical Internet perspective, these functional domains may be associated with key system-level properties such as interoperability, modularity, openness, hyperconnectivity, and recursive coordination across interconnected logistics networks. The first important convergence concerns the role of coordination and information capabilities. In the empirical studies, these two domains form the most structured co-occurrence pattern, indicating that integration and visibility are frequently modeled together. In the PI literature, the same logic appears as a central requirement of interoperable logistics networks. Digital interoperability, seamless information exchange, tracking, routing, and visibility are presented as necessary conditions for interconnecting and coordinating supply chain actors within a system of logistics networks. This suggests that the coordination-information pairing identified in the empirical literature can be interpreted as a network-level pathway aligned with PI principles.

A second convergence concerns the enabling role of digital technologies. In the empirical consolidation, enabling capabilities remain marginal as mediating mechanisms, indicating that technology itself is rarely modeled as the direct pathway to performance. Instead, digital infrastructure serves as an enabling layer whose value materializes through higher-order capabilities. This interpretation is strongly supported by the PI literature, where technologies such as IoT, cyber-physical systems, blockchain, smart interfaces, and service-oriented architectures are presented as foundational components of interconnection, visibility, and coordination rather than as performance outcomes in themselves, while supporting modular and interoperable logistics infrastructures. In other words, technology enables the system, but does not constitute the mediating capability by itself.

A third convergence concerns adaptive capabilities. The empirical literature shows that resilience, flexibility, agility, and viability constitute the most frequently observed pathway between digitalization and performance. Although these mechanisms are often modeled independently from coordination and information capabilities, they are highly relevant in a PI context. The PI literature repeatedly associates interconnected logistics systems with improved responsiveness, visibility, dynamic routing, and the capacity to manage complexity under uncertain conditions. This suggests that adaptive capabilities represent a major performance-oriented intelligence function within digital supply chains, even when empirical models do not yet integrate them into a broader systemic architecture.

Operational capabilities also remain important in this interpretation. The empirical studies frequently conceptualize digitalization through operational improvement mechanisms such as lean, green, procurement, and process excellence practices. In the PI literature, comparable concerns appear in relation to routing, consolidation, scheduling, hub coordination, and system-wide efficiency. However, as shown by the empirical findings, these operational pathways are usually modeled as autonomous mechanisms rather than as part of an integrated multi-functional system.

Taken together, these convergences support a conceptual shift from firm-level capability pathways to network-level intelligence functions. The proposed model conceptualizes digitalization as an enabling layer that activates distributed intelligence functions (operational, coordination, information, and adaptive), which in turn drive system-level performance outcomes, including operational performance, resilience, and sustainability. From this perspective, the five functional domains identified in the empirical consolidation constitute network-level intelligence dimensions of digital supply chains:

- **Enabling intelligence**, referring to the technological and infrastructural layer that supports interconnection across the network;

- **Operational intelligence**, capturing process optimization and execution efficiency within supply chain activities;
- **Coordination intelligence**, reflecting integration and synchronization mechanisms across actors and flows;
- **Information intelligence**, encompassing visibility, traceability, and data-sharing capabilities across the network;
- **Adaptive intelligence**, representing flexibility, resilience, viability, and dynamic response under changing conditions.

Among these dimensions, coordination and information jointly form the core of network intelligence, as they enable synchronization and visibility across interconnected supply chain actors.

This conceptualization does not imply that the empirical studies explicitly adopt a Physical Internet perspective. Rather, it suggests that the functional mechanisms observed in digital supply chain research provide an empirical basis for understanding how digital supply chains may evolve toward interoperable, hyperconnected, and intelligent logistics systems. The contribution of this paper is therefore not to claim that digital supply chains are already Physical Internet systems, but to show that the capability pathways identified in empirical DSC research can be conceptually reorganized as intelligence functions consistent with the PI vision of interconnected logistics networks.

This also helps explain an important limitation of the current empirical literature. While the identified functions are individually well represented, they are rarely integrated within the same empirical model. The findings therefore suggest not the presence of a mature multi-functional architecture, but rather an independent set of pathways that together point toward a broader system logic. In this sense, the transition from digital supply chains to Physical Internet-compatible systems appears less as a purely technological shift than as a process of functional integration across distributed intelligence domains. The transition from largely independent capability pathways to network-level intelligence is illustrated in Figure 1.

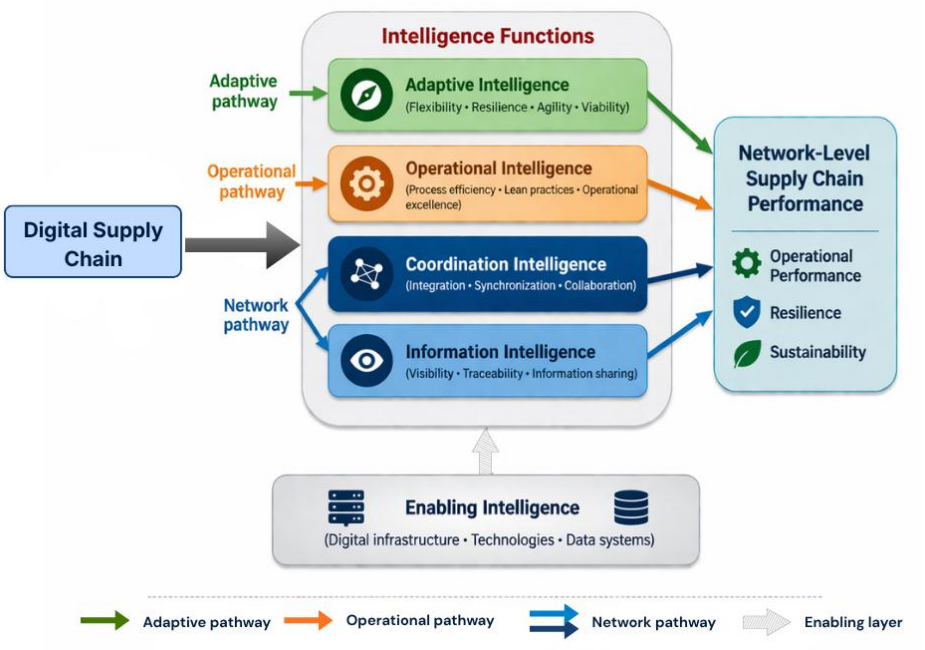


Figure 1: From digitalization to network intelligence: functional capability pathways toward Physical Internet-compatible supply chains.

The figure illustrates three dominant empirical pathways through which digitalization affects supply chain performance: adaptive, operational, and network (coordination-information) pathways. Enabling technologies act as an underlying infrastructure rather than a direct mediating mechanism. The model highlights the fragmented yet structured nature of existing empirical studies and their potential integration into network-level intelligence functions aligned with Physical Internet principles, including interoperability, coordination, and interconnected logistics operations.

5 Contributions

This study makes four main contributions to the literature on digital supply chains and the Physical Internet.

First, it consolidates fragmented empirical evidence by identifying a limited set of recurrent intermediate capability pathways through which digitalization influences supply chain performance. The results demonstrate that performance effects are predominantly mediated, rather than direct, through adaptive, operational, coordination, and information capabilities.

Second, it reveals the underlying structure of these pathways by showing that they are organized as distinct and weakly interconnected configurations, namely adaptive, operational, and coordination-information pathways, thereby exposing the absence of integrated multi-functional architectures in existing empirical models.

Third, it reconceptualizes these capability pathways as distributed intelligence functions, positioning digitalization as an enabling layer that activates operational, coordination, information, and adaptive intelligence at the system level.

Fourth, it extends this reconceptualization to the Physical Internet by demonstrating that coordination and information functions form the core of network intelligence, providing an empirically grounded explanation of how digital supply chains can evolve toward interoperable, hyperconnected, and Physical Internet-compatible logistics systems.

From a managerial perspective, the findings suggest that organizations should move beyond isolated digital investments and instead develop coordination and information capabilities as central drivers of network-level performance.

Finally, the study identifies a critical gap, the lack of integrated multi-functional models, and calls for future research on the integration of distributed intelligence functions to enable system-level optimization in Physical Internet environments.

6 Conclusion

This study provides an empirically grounded synthesis of how digital supply chain contributes to supply chain performance through intermediate functional capabilities. The findings show that digitalization primarily operates through mediated pathways, where adaptive, operational, coordination, and information capabilities act as key explanatory mechanisms, while technological elements function as an enabling infrastructural layer.

The results further reveal that these mechanisms are predominantly structured as largely independent pathways, with limited integration across functional domains. Building on this observation, the study proposes a shift from firm-level capability pathways to network-level intelligence functions, offering a structured interpretation of digital supply chains as distributed systems of operational, coordination, information, adaptive, and enabling intelligence.

This perspective aligns with the Physical Internet paradigm, highlighting that the transition toward interconnected, interoperable, and increasingly hyperconnected logistics networks is not solely technological, but depends on the integration of multiple functional dimensions across actors and systems.

Overall, the study provides a coherent bridge between empirical research on digital supply chains and the systemic vision of the Physical Internet, offering a foundation for future research on integrated logistics systems within Physical Internet environments.

7 References

- Cheng T.C.E., S.S. Kamble, A. Belhadi, N.O. Ndubisi, K.-H. Lai, M.G. Kharat (2022): Linkages between big data analytics, circular economy, sustainable supply chain flexibility, and sustainable performance in manufacturing firms. International Journal of Production Research, v60, no22, 6908-6922. https://doi.org/10.1080/00207543.2021.1906971
- Haddud A., A. Khare (2020): Digitalizing supply chains: Potential benefits and impact on lean operations. International Journal of Lean Six Sigma, v11, no4, 731-765. https://doi.org/10.1108/IJLSS-03-2019-0026
- Hallikas J., M. Immonen, S. Brax (2021): Digitalizing procurement: The impact of data analytics on supply chain performance. Supply Chain Management: An International Journal, v26, no5, 629-646. https://doi.org/10.1108/SCM-05-2020-0201
- Lee K.L., S.Y. Wong, H.M. Alzoubi, B. Al Kurdi, M.T. Alshurideh, M. El Khatib (2023): Adopting smart supply chain and smart technologies to improve operational performance in manufacturing industry. International Journal of Engineering Business Management, v15, 18479790231200614. https://doi.org/10.1177/18479790231200614
- Pan S., D. Trentesaux, D. McFarlane, B. Montreuil, E. Ballot, G.Q. Huang (2021): Digital interoperability in logistics and supply chain management: state-of-the-art and research avenues towards the physical internet. Computers in Industry, v128, 103435. https://doi.org/10.1016/j.compind.2021.103435
- Sallez Y., B. Montreuil, E. Ballot (2015): On the activeness of physical internet containers. In Service orientation in holonic and multi-agent manufacturing, Studies in Computational Intelligence, v594, 259-269. https://doi.org/10.1007/978-3-319-15159-5_24
- Sarraj R., E. Ballot, S. Pan, B. Montreuil (2014): Analogies between internet network and logistics service networks: challenges involved in the interconnection. Journal of Intelligent Manufacturing, v25, no6, 1207-1219. https://doi.org/10.1007/s10845-012-0697-7
- Sarraj R., E. Ballot, S. Pan, D. Hakimi, B. Montreuil (2014): Interconnected logistic networks and protocols: simulation-based efficiency assessment. International Journal of Production Research, v52, no11, 3185-3208. https://doi.org/10.1080/00207543.2013.865853
- Sternberg H., A. Norrman (2017): The physical internet - review, analysis and future research agenda. International Journal of Physical Distribution & Logistics Management, v47, no8, 736-762. https://doi.org/10.1108/IJPDLM-12-2016-0353
- Tran-Dang H., N. Krommenacker, P. Charpentier, D.-S. Kim (2020): Towards the internet of things for physical internet: perspectives and challenges. IEEE Internet of Things Journal, v7, no6, 4711-4736. https://doi.org/10.1109/JIOT.2020.2971736
- Tran-Dang H., D.-S. Kim (2021): The physical internet in the era of digital transformation: perspectives and open issues. IEEE Access, v9, 164613-164631. https://doi.org/10.1109/ACCESS.2021.3131562
- Treiblmaier H., K. Mirkovski, P.B. Lowry, Z.G. Zacharia (2020): The physical internet as a new supply chain paradigm: a systematic literature review and a comprehensive framework. The International Journal of Logistics Management, v31, no2, 239-287. https://doi.org/10.1108/IJLM-11-2018-0284

- Yu W., R. Chavez, Q. Liu, T. Cadden (2024): Examining the effects of digital supply chain practices on supply chain viability and operational performance: A practice-based view. IEEE Transactions on Engineering Management, v71, 10413-10426. https://doi.org/10.1109/TEM.2023.3294670

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

PI-Container flow optimization in a Multi-Echelon Physical Internet network *

Hadjer LOUZIM[1,2] , Abdelghani BEKRAR[2], Fouad MALIKI[3,4], Tarik CHARGUI[2] , Atika RIVENQ[1,5] and Marwane AYAIDA[1]

1. IEMN, UMR CNRS 8520, Université Polytechnique Hauts-de-France, Valenciennes, France
2. LAMIH UMR CNRS 8201, Université Polytechnique Hauts-de-France, Valenciennes, France
3. Ecole Supérieure en Sciences Appliquées de Tlemcen, ESSA-Tlemcen, BP 165 RP Bel Horizon, Tlemcen, 13000, Algeria.
4. Manufacturing Engineering Laboratory of Tlemcen (MELT), University of Tlemcen, Tlemcen, Algeria
5. TicLab UIR (Université Internationale de Rabat)

Abstract: *Efficient coordination of inventory and transportation on decisions is essential for improving the performance of interconnected logistics systems, especially in Physical-Internet-based supply chains. This paper propose a mixed-integer linear programming (MILP) model for a multi-period inventory routing problem (IRP) in a Physical-Internet network integrating both customer delivery and outsourced replenishment flows with vertical and horizontal supply chain alignment. The model is based on three objectives: minimization of transportation costs, inventory holding costs and vehicle space utilization costs across a network of π-hubs, retailers, and suppliers over a multi-period planning horizon. A hierarchical (lexicographic) objective is employed: giving different priorities for the objectives to illustrate the trade-off between the objectives. The proposed model is implemented and solved using Gurobi solver. Computational experiments on a small-scale instance demonstrate the model's ability to coordinate replenishment and distribution decisions while respecting capacity, volume, and flow conservation constraints.*

Keywords: *Physical Internet, Encapsulation, PI-containers, VRP, Multi-objective optimization, PI-hubs, Multi-echelon networks, Gurobi.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan: ☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),☐ Transportation Equipment, ☒ PI Networks, ☐ System of Logistics Networks, ☒ Vertical Supply Consolidation, ☒ Horizontal Supply Chain Alignment, ☐ Logistics/Commercial Data Platform, ☐ Access and Adoption, ☐ Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

1 Introduction

In today's rapidly changing world of global commerce and logistics, the quest for efficiency, sustainability, and resilience stands as a top priority. Traditional supply chain models, while effective in their time, are increasingly facing challenges. These challenges manifest in various aspects of supply chain operations: in an economic, environmental and societal aspect [1]. Recognizing the need for transformation and change, the concept of the Physical Internet emerges as a promising solution to overcome the limitations of classical supply chains. The Physical Internet (π) concept has been recently introduced as a response to the Global Logistics Sustainability Grand Challenge [1]. It is defined as an open global

logistics system founded on physical, digital and operational interconnectivity through encapsulation, interfaces and protocols, aiming to change the way physical objects are moved, handled, stored and transported based on the structure and principles of the Digital Internet [2]. The main objective of PI is to move from a closed, independent logistics network into an open, dependent logistics network. Inventory routing problem is a complex challenge in supply chain management where suppliers manage inventory levels for geographically distributed customers through intermediary hubs. Many articles studied this type of problem, with different numbers of echelons, and since our problem is a two-echelon network, we were more interested in this type of problem. One study of [3] proposes a two-phase metaheuristic approach for solving the two-echelon inventory-routing problem, where the objective is to optimize both routing and inventory costs over a fixed planning horizon between suppliers, satellites and customers, while considering the inventory levels of satellites and customers. The proposed approach combines tabu search and mathematical programming to efficiently handle small to large instances, and provides valuable insights into managing both inventory demands and routing cost reduction in a supply chain with flexible satellite assignment. Another study [4] proposed a network composed of a supplier that delivers products through distribution centers and end customers, with the main objective of minimizing total costs by combining inventory, transportation, and order costs, while respecting capacity constraint. To solve this problem, the article proposes a mixed-integer linear programming formulation and develops an exact branch-and-cut algorithm. Another contribution [5] introduces the Two-Echelon Multi-Depot Inventory-Routing Problem with pickup and deliveries, involving both inbound and outbound flows in a multi-depot environment. The authors proposed a mathematical model that incorporates two inventory policies. The objective function in the proposed model aims to minimize the total system cost by combining transportation, storage, and routing costs. To solve the problem, they propose a hybrid algorithm that combines branch-and-cut with an adaptive large neighborhood search. Another study takes an exact optimization perspective [10]. The authors present a MILP model integrating at the same time warehouse, inventory, and transportation planning in a multi-echelon supply chain, minimizing total operational costs while ensuring the respect of deliveries time and optimal inventory levels. Key decision variables include warehouse capacities, inventory replenishment policies, and transportation routes. Computational experiments on realistic supply chain data demonstrate significant cost savings compared to conventional approaches. In the Physical Internet context, one study [6] explores how the Physical Internet can improve the management of production, inventory, and distribution in global logistics networks. Unlike traditional supply chains, PI uses interconnected hubs and standardized containers to enhance efficiency. The authors present a mixed-integer linear programming model solved using CPLEX, with the main objective of minimizing total cost in one objective function that integrates production, inventory, and distribution decisions, which are often studied separately in the literature. The results show the superiority of PI in terms of cost reduction while providing similar or better service levels compared to traditional supply chain configurations. The same authors [7] proposed in another article a multi-objective mixed-integer linear programming model for the same problem that focuses on minimizing at the same time total costs, greenhouse gas emissions, and accident risks. The study compares the PI-enabled system with traditional models, demonstrating that PI offers substantial environmental, economic, and social benefits. For the resolution of this problem, the authors used the augmented epsilon-constraint method, which helps generate a set of Pareto-optimal solutions by optimizing the primary objective while imposing constraints on the other objectives. Another work introduces an advanced inventory replenishment optimization approach in the Physical Internet framework,

proposing a hybrid Simulated Annealing-Genetic Algorithm to optimize the balance between exploration and exploitation. The study presents an enriched mathematical model integrating dynamic demand and multi-objective optimization. This paper addresses transportation in a Physical Internet network with the consideration of inventory control. The problem under study is formulated as a two-echelon inventory routing problem with the introduction of the concept of encapsulation of PI-containers. Section 2 presents the problem description, Section 3 addresses our problem and its mathematical formulation, Section 4 discusses the implementation of the model in the solver and the obtained results.

2 Problem Description and Notation

2.1 General description

This model is a multi-period, two-echelon inventory routing problem designed to optimize logistics by integrating routing and inventory decisions within a network of suppliers, hubs, and retailers. At its core, the model manages inventory levels at each π-hub over time, while balancing capacities and demands through replenishment strategies such as just-in-time or maximum-level policies. The aim is to avoid both stock outs and excessive holding costs. In this proposed model, suppliers deliver products to hubs then hubs distribute them to retailers through optimized routes, with the objective of minimizing transportation costs, holding costs, and improving the filling of PI-containers based on the concept of encapsulation. Exchange between π-hubs is allowed as another strategy for replenishment (resource sharing). PI-boxes can be considered as P-Containers (packaging containers) and π-containers can be considered as T-containers. Demands are supposed expressed in term of π-boxes .The work of [9] that introduced a multi-phase model, inspired the idea of PI-boxes, which represents another aspect of Physical Internet principles. In their work, demands were expressed in terms of π-boxes of different sizes, represented by fractional volumes, which are transported using a set of π -containers with a maximum capacity, also expressed as term of volume. This work extends our previous work [8]. The main objective of the model is to simultaneously minimize transportation costs, holding costs, and the empty space inside the PI-containers. In addition, the proposed model considers the number of utilized PI-containers and their filling rates as decision variables, in order to compare how the system reacts under different scenarios. More specifically, the distribution network consists of a set of suppliers, denoted by F, where f = 1… F, responsible for replenishing a set of hubs, denoted by H, where h = 1, …, H using a set of PI-containers OC where oc=1,…,OC. Each hub serves a designated group of retailers, denoted by R, where i = 1, …,R, with different demands expressed in terms of PI-boxes, denoted by P, where p = 1, …, P. These PI-boxes are transported using a set of PI-containers, denoted by C , where c = 1,… C, over a set of time periods, denoted by T, where t = 1,… T. This network structure allows the model to capture the flow of goods from suppliers to retailers through hubs, while managing both inventory and routing decisions to satisfy retailers demand.

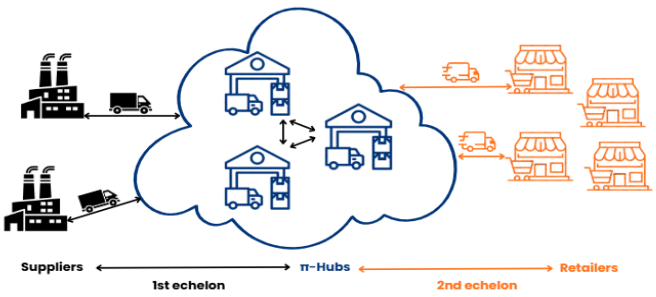


Figure 1: Proposed Multi-Echelon π-Network

2.2 Assumptions

The proposed model addresses the inventory multi-period routing problem within the Physical Internet framework. Before presenting the model in detail, it is important to mention the main assumptions that guide our approach.

- First, the model operates within a two-echelon distribution system. In this structure, products are first transported from suppliers to hubs. The hubs then replenish their inventory before delivering products to retailers.
- To simplify the model, only one supplier is considered, and this supplier is assumed to have infinite capacity.
- Retailers' demands are expressed in terms of PI-boxes, which are transported using PI-containers.
- Following the work of [9] both PI-containers and PI-boxes are defined using fractional volumes. All PI-containers have the same fractional size of 2, while PI-boxes are available in four different fractional sizes: 0.25, 0.5, 0.75, and 1.
- A penalty cost is associated with the remaining free space inside PI-containers. In addition, a holding cost is associated with the remaining stock of each type of PI-box stored at each hub.
- The model is developed over multiple time periods.
- A retailer's demand may be satisfied using different PI-containers, and each retailer may be visited multiple times within a route.
- Each hub has an inventory level that must be respected throughout the planning horizon.
- Hubs are also allowed to share different resources, including trucks, drivers, and stock levels.
- Finally, the starting and ending hub of a truck may be the same or different. At the end of its route, the truck returns to the closest hub included in the route.

2.3 Sets and Parameters

Table 1: Sets and Parameters

Parameter	Description
H, I, P, C, T, F, OC	Number of π-hubs, retailers, π-Boxes, Second echelon π-containers, time periods, Facilities, first echelon π-containers respectively
D[h][i]	Distance from hub h to retailer i
Dist[i][j]	Distance from retailer i to retailer j
D2[f][h]	Distance from supplier f to hub h
Dist2[l][k]	Distance between hub l and hub k
Dem[i][p][t]	Demand of product p at retailer i in period t
CAP[h][p]	Storage capacity for π-Boxes p at hub h
Vol[c]	Volume of Second echelon π-container c

VolOC[oc]	Volume of first echelon π-container oc
V[p]	Volume per unit of π-Box p
HC[p]	Holding cost per unit of π-Box p per time period
TC	Transportation cost per unit distance per trip
Cout	Cost per unit of unused truck volume capacity
IL0	Initial inventory level at all hubs

2.4 Decision Variables

Table 2: Decision variables

Variable	Description
B[h,i,p,c,t] ∈ Z+	Quantity of π-Box p delivered from π-hub h to retailer i by Second echelon π-container c in time period t
Q[f,h,p,oc,t] ∈ Z+	Quantity of π-Box p shipped from supplier f to hub h by first echelon π-container oc in time period t
QN[k,l,f,p,oc,t] ∈ Z+	Quantity of π-Box p transferred between hub k and l via π-container oc from supplier f in period t
IL[h,p,t] ∈ Z+	Inventory level of π-Box p at π-hub h at end of period t
Y[h,i,c,t] ∈ {0,1}	1 if π-container c travels directly from hub h to retailer i in period t
X[i,j,h,c,t] ∈ {0,1}	1 if π-container c travels from retailer i to retailer j (via hub h route) in period t
Z[i,h,c,t] ∈ {0,1}	1 if π-container c returns from retailer i to hub h in period t
L[c,t] ∈ {0,1}	1 if second echelon truck c is used in period t
L2[oc,t] ∈ {0,1}	1 if π-container oc is used in period t
M[f,h,oc,t] ∈ {0,1}	1 if π-container oc starts from supplier f to π-hub h in period t
N[l,k,f,oc,t] ∈ {0,1}	1 if π-container oc travels between hubs l and k from supplier f in period t
O[h,f,oc,t] ∈ {0,1}	1 if OC vehicle oc goes back from hub h to supplier f in period t
VolOCP[c,t] ≥ 0	Volume occupied in second echelon π-container c in period t
VolOCP2[oc,t] ≥ 0	Volume occupied in π-container oc in period t
TR[c,t] ≥ 0	Fill rate (%) of π-container c in period t
TR2[oc,t] ≥ 0	Fill rate (%) of π-container oc in period t
U[i,t], K[h,t] ∈ Z+	Auxiliary subtour-elimination variables for retailers and hubs respectively

3 Mathematical Formulation

The model is formulated as a Mixed-Integer linear Program (MILP) with a hierarchical (lexicographic) objective and was inspired from the work of [9] and our previous work [8] where we extended our published model into a multi-echelon network with other assumptions already exposed. The primary objective minimizes the cost of unused truck capacity (empty space cost), and the secondary objective minimizes inventory holding costs.

3.1 Objective Functions

$F_1 = Cout \cdot \Sigma c,t\ [(Vol_c \cdot L_\{c,t\} - VolOCP_\{c,t\}) / 0.25]$ (1)

+ Cout · Σ oc,t [(VolOC_{oc} · L2_{oc,t} - VolOCP2_{oc,t}) / 0.25]

$F_2 = \Sigma_\{h,p,t\}\ HC_p \cdot IL_{h,p,t}$ (2)

F_3 = TC · Σ_{h,i,c,t} $D_{hi} \cdot Y_{hict}$ + TC · Σ_{i,h,c,t} $D_{h,i} \cdot Z_{i,h,c,t}$ (3)

+ TC · Σ_{i,j,h,c,t} $Dist_{i,j} \cdot X_{i,j,h,c,t}$ + TC · Σ_{f,h,oc,t} $D2_{f,h} \cdot M_{f,h,oc,t}$

+ TC · Σ_{h,f,oc,t} $D2_{f,h} \cdot O_{h,f,oc,t}$ + TC · Σ_{l,k,f,oc,t} $Dist2_{l,k} \cdot N_{l,k,f,oc,t}$

3.2 Constraints

$\Sigma_\{h \in H, c \in C\}\ B_{h,i,p,c,t} = Dem_{i,p,t}$ ∀i ∈ I, p ∈ P, t ∈ T (4)

$IL_{h,p,t} = IL_{h,p,t-1} + \Sigma_\{f,oc\}\ Q_{f,h,p,oc,t} + \Sigma_\{k,f,oc\}\ QN_{k,h,f,oc,p,} - \Sigma_\{k,f,oc\}\ QN_{h,k,f,oc,p,t} - \Sigma_\{i,c\}\ B_{h,i,p,c,t}$
∀h ∈ H, p ∈ P, t ≥ 1 (5)

$IL_{h,p,0} = IL_0 + \Sigma_\{f,oc\}\ Q_{f,h,p,oc,0} + \Sigma_\{k,f,oc\}\ QN_{k,h,f,p,oc,0} - \Sigma_\{k,f,oc\}\ QN_{h,k,f,p,oc,0} - \Sigma_\{i,c\}\ B_{h,i,p,c,0}$
∀h ∈ H, p ∈ P (6)

$IL_{h,p,t} \le CAP_{h,p}$ ∀h ∈ H, p ∈ P, t ∈ T (7)

$\Sigma_\{h,p,i\}\ V_p \cdot B_{h,i,p,c,t} \le Vol_c \cdot L_{c,t}$ ∀c ∈ C, t ∈ T (8)

$\Sigma_\{f,h,p\}\ V_p \cdot Q_{f,h,p,oc,t} \le VolOC_{oc} \cdot L2_{oc,t}$ ∀oc ∈ OC, t ∈ T (9)

$\Sigma_\{c,i,t\}\ B_{h,i,p,c,t} \le CAP_{h,p}$ ∀h ∈ H, p ∈ P (10)

$I \cdot \Sigma_\{i \in I\}\ Y_{h,i,c,t} \ge \Sigma_\{i,j \in I\}\ X_{i,j,h,c,t}$ ∀h ∈ H, c ∈ C, t ∈ T (11)

$I \cdot \Sigma_\{i,h\}\ Z_{i,h,c,t} \ge \Sigma_\{i,j,h\}\ X_{i,j,h,c,t}$ ∀c ∈ C, t ∈ T (12)

$\Sigma_h\ Y_{h,i,c,t} + \Sigma_\{j,h\}\ X_{j,i,h,c,t} = \Sigma_h\ Z_{i,h,c,t} + \Sigma_\{j,h\}\ X_{i,j,h,c,t}$ ∀i ∈ I, c ∈ C, t ∈ T (13)

$H \cdot \Sigma_\{h \in H\}\ M_{f,h,oc,t} \ge \Sigma_\{l,k \in H\}\ N_{l,k,f,oc,t}$ ∀f ∈ F, oc ∈ OC, t ∈ T (14)

$H \cdot \Sigma_\{h,f\}\ O_{h,f,oc,t} \ge \Sigma_\{i,j,h,c\}\ X_{i,j,h,c,t}$ ∀oc ∈ OC, t ∈ T (15)

$\Sigma_f\ Y_{h,oc,t} + \Sigma_\{k,f\}\ N_{h,k,f,oc,t} + \Sigma_f\ O_{h,f,oc,t} =$ Σ_i $Z_{h,oc,t}$ + Σ_{k,f} $N_{k,h,f,oc,t}$ + Σ_f $M_{f,h,oc,t}$
∀h ∈ H, oc ∈ OC, t ∈ T (16)

$X_{i,i,h,c,t} = 0 \quad \forall i, h, c, t$ (17)

$N_{l,l,f,oc,t} = 0 \quad \forall l, f, oc, t$ (18)

$U_{i,t} - U_{j,t} + I \cdot X_{i,j,h,c,t} \le I - 1$ ∀i ≠ j, h, c, t (19)

$K_{l,t} - K_{k,t} + H \cdot N_{l,k,f,oc,t} \le H - 1$ ∀l ≠ k, f, oc, t (20)

$Y_{h,i,c,t} \le L_{c,t}$ ∀ h ∈ H, i ∈ I, c ∈ C, t ∈ T (21a)

$Z_{i,h,c,t} \le L_{c,t}$ ∀ h ∈ H, i ∈ I, c ∈ C, t ∈ T (21b)

$X_{i,j,h,c,t} \le L_{c,t}$ ∀ i ∈ I, j ∈ I, h ∈ H, c ∈ C, t ∈ T (21c)

$M_{f,h,oc,t} \le L2_{oc,t}$ ∀ f ∈ F, h ∈ H, oc ∈ OC, t ∈ T (22a)

$O_{h,f,oc,t} \le L2_{oc,t}$ ∀ f ∈ F, h ∈ H, oc ∈ OC, t ∈ T (22b)

$N_{l,k,f,oc,t} \le L2_{oc,t}$ ∀ l ∈ H, k ∈ H, f ∈ F, oc ∈ OC, t ∈ T (22c)

$\Sigma_p\ B_{h,i,p,c,t} \le 1000 \cdot (Y_{h,i,c,t} + \Sigma_j\ X_{j,i,h,c,t})$ ∀ i ∈ I, c ∈ C, h ∈ H, t ∈ T (23)

$\Sigma_\{p,f\}\ Q_{f,h,p,oc,t} \le 1000 \cdot (\Sigma_f\ M_{f,h,oc,t} + \Sigma_\{k,f\}\ N_{k,h,f,oc,t})$ ∀ h ∈ H, oc ∈ OC, t ∈ T (24)

$\Sigma_p\ Q_{f,h,p,oc,t} \le 1000 \cdot M_{f,h,oc,t}$ ∀ f ∈ F, h ∈ H, oc ∈ OC, t ∈ T, t ≥ 1 (25)

$\Sigma_p\ QN_{l,k,f,p,oc,t} \le 1000 \cdot N_{l,k,f,oc,t}$ ∀ l ∈ H, k ∈ H, f ∈ F, oc ∈ OC, t ∈ T (26)

$VolOCP_{c,t} = \Sigma_\{h,p,i\}\ V_p \cdot B_{h,i,p,c,t}$ ∀ c ∈ C, t ∈ T (27)

$$VolOCP2_{oc,t} = \Sigma_{f,h,p} V_p \cdot Q_{f,h,p,oc,t} \quad \forall\, oc \in OC, t \in T \quad (28)$$

$$TR_{c,t} \cdot Vol_c = VolOCP_{c,t} \cdot 100 \quad \forall\, c \in C, t \in T \quad (29)$$

$$TR2_{oc,t} \cdot VolOC_{oc} = VolOCP2_{oc,t} \cdot 100 \quad \forall\, oc \in OC, t \in T \quad (30)$$

$$0 \leq VolOCP_{c,t} \leq Vol_c \quad \forall\, c \in C, t \in T \quad (31)$$

$$0 \leq VolOCP2_{oc,t} \leq VolOC_{oc} \quad \forall\, oc \in OC, t \in T \quad (32)$$

$$B, Q, QN, IL, U, K \in \mathbb{Z}_+ \quad (33)$$

$$Y, X, Z, L, L2, M, N, O \in \{0, 1\} \quad (34)$$

$$TR, TR2, VolOCP, VolOCP2 \geq 0 \quad (35)$$

Our proposed model is based on three objective functions. The first objective function, given in Eq. (1), minimizes the cost of unused space inside the π-containers in both echelons of the network. The second objective function given in Eq. (2), minimizes the inventory holding cost at different hubs over the planning horizon for all π-boxes types. Eq. (3) minimizes the transportation cost generated by all movements in the first and second echelons in the network.

The constraints ensure that the system operates consistently and efficiently while respecting all the proposed hypothesis. Eq. (4) enforces demand satisfaction at the retailer level for each π - box type and time period. Eqs. (5) and (6) model the inventory evolution at the hubs over time, considering replenishment, inter-hub transfers, and outbound deliveries (Eq. (6) concerns the first period). Eqs. (7)- (10) impose storage and π -container capacity limitations. Eqs. (11)- (16) maintain routing feasibility and flow conservation in both echelons. Eqs. (17)- (20) prevent invalid routing structures: self-loops and subtours. Eqs. (21)- (26) ensure consistency between activated routes and the associated delivery or replenishment quantities. Finally, Eqs. (27)- (32) define container occupancy and fill rates, and Eqs. (33)- (34) specify the admissible variable domains.

4 Computational Experiments

4.1 Data description

The model is tested on a small-scale instance with the following parameters: H = 3 π-hub , I = 4 retailers, P = 4 π-Boxes, C = 4 second echelon π-Containers, T = 2 periods, F = 2 suppliers, and OC = 4 first echelon π-Containers . π-Boxes volumes are V = [0.25, 0.50, 0.75, 1.00] and holding costs are HC = [2, 2, 2, 2] per unit per period. All second echelon and OC truck capacities are set to Vol = VolOC = 2 volume units. Transportation cost TC = 2.0 per unit distance, and empty space cost Cout = 5.0 per 0.25 volume unit. The model is solved using Gurobi with a lexicographic objective. More details are mentioned in *Table 3*.

4.2 Numerical results

The model was at first tested by minimizing each objective separately, namely transport cost, free-space cost and holding cost. This mono-objective analysis was then extended to several lexicographic bi-objective combinations (using Gurobi solver version 12.0) so we can observe how the system reacts when two performance criteria are considered simultaneously and to provide more informative support for managerial decision-making. This analysis is particularly relevant in a Physical Internet context, where improving one performance measure often affects the others. From a computational point of view, the transport-cost objective was taking the most important time for resolution, therefore, a time limit of 3600 seconds was imposed when the concerned objective was involved. By contrast, the free-space and holding-cost objectives were

solved much faster and did not require such a limit. Tests on the optimization of the three objectives were done but not presented due to long resolution times (a limit of three hours was fixed although no optimal solution was found). NC indicates that the corresponding objective was not included as an optimization criterion in that scenario and its value was therefore not tracked. It does not imply that the objective equals zero or is infeasible. More details are mentioned in *Table 4*.

Table 3: Data description

Scenario Family	Scenario
Distance	Random values in the range [1-20] km
Retailers demand	[1-3] PI-boxes for each type in all periods
Inventory levels	[1-3] PI-boxes for each type in all periods
Hub capacity	H1: 3 for all types of PI-boxes and H2, H3 : 10 for all types of PI-boxes
PI-box size	{0.25, 0.5, 0.75, 1}
Holding cost	2/ PI-Box
Transportation cost	2U/ Km
Space penality cost	5/ Unit of volume left in the container
Periods	2
IL0	0

Table 4: Numerical results

Scenario Family	**Transport Cost F3**	**Free space Cost F1**	**Holding Cost F2**	**FE trucks at t=0**	**FE trucks at t=1**	**SE trucks at t=0**	**SE trucks at t=1**
Mono-objective							
Transport cost	747	NC	NC	4	3	4	3
Volume (Space)	NC	10	NC	3	3	3	3
Holding cost	NC	NC	0	4	4	4	4
Multi-objective (lexicographic)							
F1 and F2 (Space priority)	NC	4	10	3	3	3	3
F1 and F2 (Holding cost priority)	NC	20	0	3	3	3	3
F1 and F3 (Holding cost priority)	754	NC	0	4	4	4	4
F3 and F1 (Free space priority)	793	10	NC	4	2	3	3
F3 and F2 (Transport cost priority)	747	NC	12	4	2	4	4
F3 and F1 (transportation cost priority)	747	180	NC	4	2	4	4

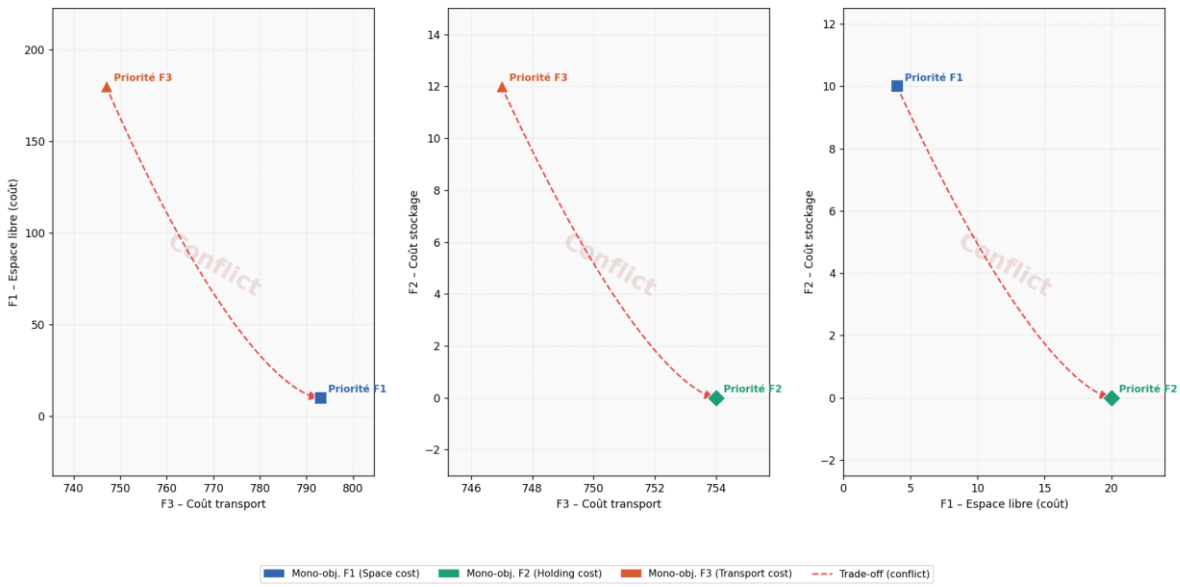


Figure 2: Conflict between objective functions

When the model focuses on free-space minimization, it uses fewer containers and improves their filling rate, which is consistent with the objective itself: reducing unused space naturally encourages better consolidation of flows and more efficient use of available vehicle capacity. By contrast, when holding cost is minimized, the model avoids keeping stocks at the hubs and bring to deliver, which completely eliminates storage cost. As a result, more vehicles are activated, since the system prefers transporting products directly when needed rather than storing them in advance. The lexicographic bi-objective results confirm the same logic. When free space is given priority over holding cost, the model seeks to load the containers more efficiently and therefore maintains a compact fleet configuration with lower unused capacity, even if some inventory must be kept at the hubs and being penalized on. On the other hand, when holding cost is prioritized, the model or eliminates stock, but this comes with poorer consolidation and higher free-space cost, since shipments are organized more to avoid storage than to maximize container utilization. A similar behavior appears when transport cost is combined with the other objectives: keeping transportation cost at its best level may require accepting either a certain amount of holding cost or a significant loss in space utilization, depending on the secondary priority. Overall, these results show that the three objectives are strongly conflicting. Therefore, the choice of the best solution depends on the operational strategy of the decision-maker: prioritizing free space leads to better container filling and fewer vehicles, while prioritizing holding cost supports a just-in-time logic with little or no stock, but generally requires more transport resources.

5 Conclusion

This paper proposed a MILP model for a multi-period inventory-routing problem in a multi-echelon Physical Internet network integrating both first-echelon replenishment flows and second-echelon customer deliveries. The proposed formulation integrates three objective to optimize: transportation cost, holding cost, and unused container space, while respecting inventory balance, hub and containers capacity limitations, and routing feasibility constraints in both echelons. The computational experiments were conducted on a small-scale proof-of-concept instance (3 hubs, 4 retailers, 2 periods), and the results should be interpreted accordingly. They demonstrate that the model produces feasible and coherent solutions and,

more importantly, that the three objectives are strongly conflicting. These findings do not constitute a validation of the model's scalability or practical performance on realistic Physical Internet networks, which remains an open direction for future work. The results highlighted that improving container filling generally leads to a more compact use of vehicles, whereas reducing holding cost supports a just-in-time strategy but may require more transport resources. Similarly, minimizing transport cost alone does not necessarily provide the best compromise in terms of inventory and space utilization. These findings confirm the relevance of a multi-objective perspective for decision-making in Physical Internet systems, where several operational goals must be balanced simultaneously. As future work, the model can be extended to larger-scale instances and tested using different meta-heuristics stochastic demand settings, and additional sustainability-related objectives, such as emissions or service quality, in order to better reflect real-world Physical Internet logistics networks.

References

[1] Montreuil B. (2011): Towards a Physical Internet: Meeting the Global Logistics Sustainability Grand Challenge. *Logistics Research*, v3, no2-3, 71-87.

[2] Montreuil B., R.D. Meller, E. Ballot (2013): Physical Internet Foundations. *Springer*, Berlin, pp. 151-166.

[3] Charaf S., D. Taş, S.D.P. Flapper, T. Van Woensel (2024): A Matheuristic for the Two-Echelon Inventory-Routing Problem. *Computers & Operations Research*, 171, 106778.

[4] Farias K., K. Hadj-Hamou, C. Yugma (2021): Model and exact solution for a two-echelon inventory routing problem. *International Journal of Production Research*, 59(10), 3109-3132.

[5] Guimaraes T.A., L.C. Coelho, C.M. Schenekemberg, C.T. Scarpin (2019): The two-echelon multi-depot inventory-routing problem. *Computers & Operations Research*, 101, 220-233.

[6] Ji S., X. Peng, R. Luo (2019): An integrated model for the production-inventory-distribution problem in the Physical Internet. *International Journal of Production Research*, 57(4), 1000-1017.

[7] Peng X., S. Ji, T. Ji (2020): Promoting sustainability of the integrated production-inventory-distribution system through the Physical Internet. *International Journal of Production Research*, 58(22), 6985-7004.

[8] LOUZIM, H., MALIKI, F., CHARGUI, T., & BEKRAR, A. (2024, November). PI-Container Flow Optimization in a Physical Internet Network. In 2024 International Conference of the African Federation of Operational Research Societies (AFROS) (pp. 1-5). IEEE.

[9] Fazili, M., Venkatadri, U., Cyrus, P., & Tajbakhsh, M. (2017). Physical Internet, conventional and hybrid logistic systems: a routing optimisation-based comparison using the Eastern Canada road network case study. International Journal of Production Research, 55(9), 2703-2730.

[10] Syofra, A. H., Tulus, T., Mawengkang, H., & Zarlis, M. (2025). An Optimization Model for Integrated Warehouse-Inventory-Transportation of a Multi-Echelon Supply Chain. *Engineering, Technology & Applied Science Research*, *15*(3), 23956-23964.

[11] Yan, W., Li, N., & Zhang, X. (2023). Enhancing supply chain management in the physical internet: a hybrid SAGA approach. *Scientific Reports*, *13*(1), 21470.

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE

A Conceptual Physical Internet Framework for EV Battery Reverse Logistics

Martyna Zielińska[1] and Marta Cudziło[2]
Łukasiewicz Research Network – Poznań Institute of Technology, Poznań, Poland[1,2,]
Corresponding author: martyna.zielinska@pit.lukasiewicz.gov.pl

Abstract: *The accelerating transition to electromobility generates a surge in end-of-life lithium-ion batteries, exposing the limitations of fragmented, semi-closed European reverse logistics networks. Current systems are structurally misaligned with exponential volume growth and regulatory requirements. Addressing this gap, this paper develops a conceptual framework for redesigning European battery reverse logistics through Physical Internet (PI) principles. Rather than simple transport optimization, the study proposes a multi-layered ecosystem: physical modularization via safety-compliant containers, network-level federated interconnectivity, and digital orchestration via the Digital Product Passport. This transformation shifts reverse networks from linear chains to open, interoperable, and digitally governed ecosystems. By embedding modular standards and distributed governance into network design, this framework reframes reverse logistics as a structurally central component of circular industrial systems. The proposed architecture provides a necessary evolution to ensure the scalability, resilience, and sustainability of Europe's EV battery recovery systems, facilitating a truly integrated circular value chain.*

Keywords: *Physical Internet, Reverse Logistics, EV Battery, Digital Battery Passport, Modular Containers, Open Mesh Network*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☒ *Transportation Equipment,* ☒ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☒ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☒ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☐ *In-Person presentation*

1 Introduction

The transition toward electric mobility is a cornerstone of the European Union's strategy to achieve climate neutrality by 2050 (Soavi et al., 2026). Lithium-ion batteries (LIBs) serve as the essential technological drivers of this shift, yet the rapid expansion of the electric vehicle (EV) market has placed significant pressure on the European supply chain. The EU remains highly dependent on imported raw materials for battery production, with key extraction processes concentrated in geopolitically sensitive regions, exposing the bloc to supply risks and price volatility (Gianvincenzi et al., 2025). To mitigate these vulnerabilities, the European Union has implemented Regulation 2023/1542, which establishes rigorous standards for battery lifecycle sustainability. A central pillar of this framework is the mandatory implementation of

the Digital Battery Passport (DBP) by February 2027, aimed at enhancing transparency, traceability, and material recovery across the value chain (Soavi et al., 2026).

Despite these regulatory developments, current infrastructure is largely unprepared to manage the anticipated 1.7 million tonnes of battery waste forecasted for 2030 (Altuntas Vural et al., 2025). Traditional reverse logistics (RL) models remain fragmented, characterized by isolated logistics silos, where original equipment manufacturers (OEMs) manage returns through proprietary and disconnected networks (Rouhani et al., 2025). These existing systems face multiple inefficiencies. Logistics and collection processes often account for a significant portion of total costs, estimated at approximately 41% of the total EVB recycling system expenses (Gonzales-Calienes et al., 2022). Furthermore, the lack of standardized handling and real-time data sharing between stakeholders limits the viability of second-life applications, such as stationary energy storage, often resulting in batteries being recycled prematurely (Altuntas Vural et al., 2025).

While significant research has addressed the optimization of forward logistics within the Physical Internet (PI) paradigm (Montreuil, 2011; Sallez et al., 2016), comparatively limited attention has been devoted to structurally redesigning reverse logistics networks for complex, hazardous, and circular industrial products. The application of PI principles to reverse battery logistics remains underexplored in both theoretical and systemic dimensions. Existing studies often focus on transport optimization, cost minimization, or facility location modeling, without questioning the underlying architectural logic of reverse networks (Rouhani et al., 2025).

This paper addresses this gap by developing a conceptual framework for redesigning the European reverse logistics network for EV batteries according to the foundational principles of the Physical Internet. The research is guided by the following questions:

1. How can the structural principles of the Physical Internet be systematically applied to reverse logistics networks for hazardous and high-value products such as EV batteries?
2. What architectural transformations are required to transition from fragmented, linear reverse chains to an open, modular, and scalable European network?
3. How does the integration of digital traceability frameworks (e.g., Digital Product Passport) reshape governance and interoperability within a PI-based reverse logistics system?

To answer these questions, the paper develops a multi-layered conceptual architecture grounded in PI theory and circular economy principles. The proposed framework integrates three interdependent layers: physical encapsulation and modularization, network topology and open interconnectivity, and digital orchestration and governance. By decoupling battery design complexity from the physical infrastructure and enabling interoperable data exchange, this framework offers a scalable blueprint for circular autonomy in Europe, aligning industrial operations with long-term climate and strategic resilience goals (Gianvincenzi et al., 2025). The following sections detail the three-layered architecture: Physical, Network, and Digital, designed to transform hazardous battery flows into a resilient, regulation-aligned circular ecosystem.

2 Literature Review

2.1 Current Paradigms in Reverse Logistics for EV Batteries

Current literature on the recovery of end-of-life (EoL) electric vehicle batteries (EVB) identifies several circular economy (CE) strategies, including repair, refurbishment, remanufacturing, and repurposing for secondary applications (Albertsen et al., 2021; Rahic and Vu, 2025). While the environmental benefits of extending battery service life are well-documented, the implementation remains challenging due to technical uncertainties and fragmented supply chains (Martinez-Laserna et al., 2018). For instance, Albertsen et al. (2021) observe that European original equipment manufacturers (OEMs) primarily focus on repurposing and repair, yet these activities are often performed within closed, proprietary networks.

Existing research in reverse logistics (RL) for batteries has largely concentrated on facility location modeling and route optimization to minimize costs (Gonzales-Calienes et al., 2022; Rahic and Vu, 2025). Gonzales-Calienes et al. (2022) utilize geographic information systems (GIS) to identify optimal recycling clusters, highlighting that transportation costs remain a dominant factor in the economic viability of EoL management. However, these models typically assume a centralized or semi-closed structure, which may lack the flexibility required to handle the projected exponential growth in battery returns (Soavi et al., 2026).

2.2 The Physical Internet (PI) Paradigm in Logistics

The Physical Internet (PI) represents a breakthrough innovation aimed at improving the efficiency and sustainability of logistics systems by an order of magnitude (Montreuil, 2011). It proposes an open, global system based on physical, digital, and operational interconnectivity (Ballot et al., 2014). Montreuil (2011) identified 13 symptoms of unsustainability in current logistics, such as empty travel and poor space utilization, which PI seeks to address through universal standards and shared resources.

A central component of the PI is the use of standardized, smart, and modular units known as π-containers (Sallez et al., 2016; Montreuil et al., 2013). These containers enable seamless handling and storage across different transport modes and operators. While the PI was initially conceptualized for forward logistics, its application to reverse industrial ecosystems is gaining attention. Treiblmaier (2019) argues that the integration of PI can achieve triple bottom line sustainability by optimizing the movement of physical components through a measurable and standardized supply chain framework.

2.3 Digital Traceability and the Digital Product Passport (DPP)

The transition toward a circular battery economy is increasingly tied to digital innovation. Regulation (EU) 2023/1542 introduces the Digital Battery Passport (BP) as a mandatory tool for ensuring transparency and traceability throughout the battery lifecycle (Soavi et al., 2026). The BP is designed to link each battery to its specific composition, performance data, and EoL instructions, facilitating more efficient sorting and dismantling (Gianvincenzi et al., 2025).

In the context of the Physical Internet, digital traceability serves as the „identity bone" of the network. Emerging technologies such as blockchain and Web3 are being explored to ensure data immutability and secure information exchange among decentralized actors (Betti et al., 2019; Duong et al., 2026). Duong et al. (2026) propose that Web3 and metaverse-enabled decentralized identifiers can reconfigure EoL diagnostics and waste management, turning retired batteries into active assets within a circular energy grid.

2.4 Research Gap: From Optimization to Architectural Redesign

Despite the growing body of literature on both PI and battery recovery, a significant research gap persists. Most existing studies on battery reverse logistics treat the network as a static, linear chain where optimization is limited to cost or distance (Gonzales-Calienes et al., 2022; Rahic and Vu, 2025). Conversely, general PI literature often overlooks the specific safety and regulatory constraints inherent in hazardous reverse flows of complex industrial products like EV batteries (Betti et al., 2019).

There is a lack of integrated studies that combine the architectural logic of the Physical Internet, specifically mesh topologies and open interconnectivity, with the mandatory safety and traceability requirements of the EU Battery Regulation. This paper addresses this gap by proposing a systemic redesign of the European battery reverse logistics network, moving beyond simple route optimization toward a modular, interoperable, and digitally orchestrated ecosystem.

3 The Conceptual PI-Based Reverse Logistics Framework

3.1 Framework Overview and Multi-Layered Architecture

The proposed framework transitions from traditional, linear, and semi-closed reverse supply chains toward an open global logistics system (Ballot et al., 2014). Grounded in the foundational principles of the Physical Internet (PI), the framework conceptualizes the European reverse battery network as a modular ecosystem rather than a series of dedicated bilateral arrangements. As noted by Altuntas Vural et al. (2025), circular models for second-life batteries require supply chains capable of handling heterogeneity and high uncertainty. This study addresses these requirements through a three-layer architecture: (1) Physical Layer, focusing on encapsulation; (2) Network Layer, defining the federated mesh topology; and (3) Digital Layer, providing orchestration and governance. The interconnected nature of these functional layers is synthesized in Figure 1: The Conceptual PI-Based Reverse Logistics Framework of EV Batteries, which illustrates the systemic transition from traditional, bilateral arrangements to an integrated, modular ecosystem designed for circular autonomy.

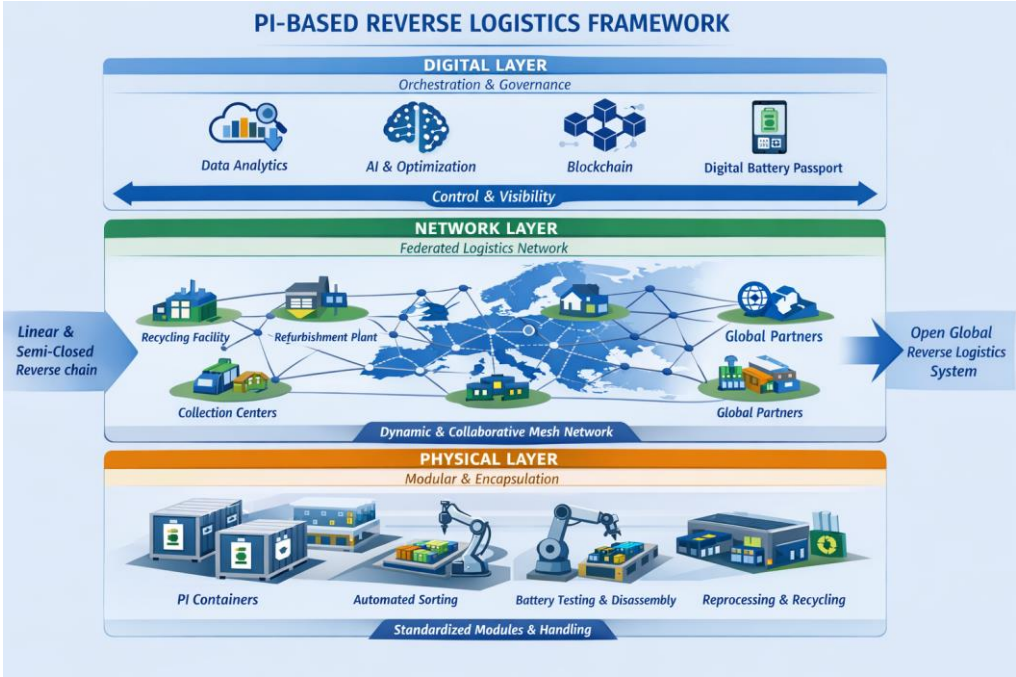


Figure 1: The Conceptual PI-Based Reverse Logistics Framework of EV Batteries

3.2 Layer 1: Physical Encapsulation and Modularisation (π-containers)

At the physical layer, the framework applies the concept of standard encapsulation (Montreuil, 2011). In a PI-based system, EV battery packs and modules are not transported as loose cargo but are enclosed in standardized, smart, and reusable π-containers (Sallez et al., 2016). This encapsulation decouples the technical specificities of the battery (e.g., specific brand dimensions or chemistries) from the handling and transport infrastructure, enabling universal interconnectivity (Sallez et al., 2016).

3.2.1 Standardising Safety and Handling for Hazardous Flows

Unlike standard freight, EV batteries are classified as hazardous materials (HazMat) due to risks of thermal runaway and toxic leakage (Betti et al., 2019). Consequently, π-containers in this framework serve as „safety envelopes" that must comply with strict regulatory certifications (Soavi et al., 2026). These containers are designed with integrated sensors for real-time monitoring of temperature, voltage, and gas levels, as suggested by the functional requirements for intelligent containers (Sallez et al., 2016). By modularizing these units into standard T-containers (transport) and H-containers (handling), the system facilitates the consolidation of diverse battery types into a single efficient flow, reducing the economic and safety risks associated with fragmented transport (Montreuil et al., 2013).

3.3 Layer 2: Network Topology and Federated Interconnectivity (π-nodes)

The network layer moves away from fixed facility locations toward a federated mesh of π-nodes. Current models, such as those analyzed by Gonzales-Calienes et al. (2022), focus on optimizing static routes between collection sites and specific dismantling hubs. The PI paradigm, however, utilizes an „open mobility web" where nodes, ranging from collection centers to certified recyclers, operate as shared infrastructure (Montreuil, 2011; Shaikh et al., 2024).

3.3.1 Decentralised Mesh Topology vs. Linear Chain Models

In this mesh topology, routing decisions are dynamic and informed by the overall state of the network. Following the protocols for inter-hub transportation (Shaikh et al., 2024), flows are reallocated based on the nearest available π-node with the required processing capacity or the optimal environmental performance. This structure mitigates the risks identified by Gianvincenzi et al. (2025) regarding the current lack of sufficient recycling capacity in the EU, as it allows for cross-border asset sharing and reduces the reliance on a single, centralized facility that may become a bottleneck or a safety risk.

3.3.2 Scalability Across Cross-Border Networks

The mesh topology is essential for scaling the network across European borders. As highlighted in the PI roadmap (ALICE, 2015), universal connectivity requires standardized protocols for the „Logistics Web." By conceptualizing reverse facilities as open nodes, the system can dynamically scale as battery return volumes grow, allowing new actors (e.g., third-party dismantlers or specialized second-life OEMs) to join the network seamlessly (Altuntas Vural et al., 2025; Chirumalla et al., 2024).

3.4 Layer 3: Digital Orchestration and Governance

The digital layer acts as the „operating system" of the reverse logistics network, ensuring that physical flows remain compliant and efficient. This layer leverages the Digital Product Passport (DPP) as the primary data backbone (Soavi et al., 2026).

3.4.1 The Digital Product Passport as an Identity Backbone

The DBP provides the mandatory and dynamic data attributes, such as State-of-Health (SoH), chemical composition, and usage history, necessary for automated decision-making (Gianvincenzi et al., 2025; Bina, 2025). In our framework, the DBP is integrated with a Decision Support System (DSS) to determine the optimal EoL pathway (e.g., remanufacturing vs. recycling) for each battery entering a π-node (Bina, 2025; Ma et al., 2024). This integration automates the „Product Disposition" phase, which is traditionally a manual and high-cost bottleneck in reverse logistics (Altuntas Vural et al., 2025).

3.4.2 Balancing Openness with Regulatory Control

A significant challenge in applying PI to hazardous flows is the tension between an open network and strict regulatory control. The orchestration layer introduces a governance dimension where decentralized technologies, such as Web3 and smart contracts, enforce compliance and safety protocols (Duong et al., 2026; Betti et al., 2019). For example, a π-container carrying a damaged battery can only be routed to a node certified for HazMat handling, with the smart contract automatically verifying the node's credentials and environmental permits before authorizing the transport (Duong et al., 2026). This ensures that the system remains open and flexible while guaranteeing the high levels of safety and liability required by the EU Battery Regulation.

4 Discussion and Evaluation

4.1 Structural Transformations: From Linear Chains to Open Ecosystems

The transition from current semi-closed reverse logistics (RL) to a Physical Internet (PI) based framework necessitates a fundamental shift in the architectural logic of battery recovery. Traditional RL models for EV batteries are often characterized by „logistics silos", where an OEM manages its own returns through dedicated contracts (Albertsen et al., 2021). Our framework replaces these rigid structures with a modular and federated mesh. As noted by Montreuil (2011), the inefficiency of current systems stems from poor asset utilization. In the context of batteries, this is exacerbated by the geographical dispersion of specialized recycling facilities (Gianvincenzi et al., 2025). By treating every certified dismantling center as a π-node, the network achieves a „pooling effect" that reduces empty backhauls and optimizes the use of high-cost, safety-certified infrastructure.

4.2 Safety and Hazardous Material (HazMat) Constraints

Reviewers of this conceptual framework correctly identified that EV batteries are structurally different from standard PI cargo due to their hazardous nature (Betti et al., 2019). To address this, the framework's feasibility rests on the „Smart π-container" concept. Unlike generic containers, battery-specific π-units must integrate fire-suppression materials and real-time gas monitoring sensors (Sallez et al., 2016). We argue that the PI paradigm actually enhances HazMat safety: in a traditional linear chain, a damaged battery might travel long distances to reach a specific OEM partner. In a PI-mesh, the digital orchestration layer can instantly reroute a critical battery to the nearest certified π-node capable of handling thermal runaway risks, thereby minimizing the „risk-exposure-time" on public roads.

4.3 Digital Traceability as an Enabler of Interoperability

The integration of the Digital Battery Passport (DBP) is the primary driver of interoperability within the framework. Under Regulation (EU) 2023/1542, the DBP provides the „State of Health" (SoH) and „State of Power" (SoP) data required for automated sorting (Soavi et al., 2026). For the PI-mesh to function, this data must be accessible across the network without compromising proprietary OEM secrets. We propose a decentralized governance model where the digital layer uses zero-knowledge proofs or blockchain-based access control (Duong et al., 2026). This allows a π-node to verify that a battery is safe for transport and identifies its optimal EoL path (repurposing vs. recycling) without requiring full access to the OEM's internal chemistry database.

4.4 Evaluation Metrics and Falsifiable Hypotheses

Translating this conceptual architecture into empirical evidence requires robust quantitative assessment. To bridge the gap between theoretical framework and practical application, we propose a set of evaluation metrics and hypotheses grounded in the foundational principles of the Physical Internet (PI).

4.4.1 Evaluation Metrics

To quantify system performance, we define three core metrics:

Interoperability index (I_{pi}): Defined as the ratio of successful battery handovers between non-affiliated network actors to the total number of attempted transfers. This metric is grounded in the PI foundation of Universal Interconnectivity, which demands that physical and informational entities flow seamlessly through heterogeneous systems, regardless of the specific actor or infrastructure involved.

Network resilience score (R_{net}): This metric evaluates system stability by measuring the capability of the network to maintain consistent end-of-life processing volumes during a hypothetical 20% failure rate among primary regional nodes. It draws upon the PI principle of the Mobility Web, which theorizes that redundant paths within a federated mesh topology provide superior robustness compared to rigid, linear reverse supply chains.

Asset utilization rate (U_{π}): This metric quantifies operational efficiency by tracking the proportion of time that π-containers and transport vehicles operate at their certified safety capacity. It directly addresses the PI core objective of eliminating "shipping air" through standardized modular encapsulation, thereby optimizing cargo load factors and reducing systemic waste.

4.4.2 Falsifiable Hypotheses

Guided by these principles, we posit two hypotheses to test the comparative advantages of the PI paradigm over traditional reverse logistics structures:

Hypothesis (H_1): A PI-based mesh network will demonstrate at least 30% higher resilience to cross-border logistical disruptions compared to traditional closed-loop models. This hypothesis is predicated on the theoretical assertion that hyperconnected, multi-path routing enables the system to absorb localized shocks far more effectively than the single-path dependency inherent in traditional hub-and-spoke models.

Hypothesis (H_2):The standardization of hazardous material flows via π-containers will reduce the average lead time from collection to dismantling by at least 25%. This expectation is derived from the PI tenet of Standard Coordination Protocols, which aim to minimize the administrative and regulatory friction currently exacerbated by non-standardized, siloed hazardous material workflows.

Together, these metrics and hypotheses provide a falsifiable foundation for subsequent simulation-based validation, enabling a rigorous assessment of the proposed reverse logistics ecosystem.

4.5 Broader Applicability to Circular Industrial Flows

While this framework is tailored to EV batteries, the underlying logic of „encapsulated hazardous modularity" is applicable to other complex circular flows, such as industrial chemicals or high-value electronics (Chirumalla et al., 2024). The synergy between the Physical Internet and the Digital Product Passport creates a scalable blueprint for the European „Circular Autonomy" strategy, providing a pathway to meet the 2030 recovery targets for critical raw materials (Gianvincenzi et al., 2025).

5 Conclusion and Recommendations

5.1 Summary of Research Contributions

This research has introduced a systemic conceptual framework for the redesign of the European reverse logistics network for electric vehicle batteries (EVB), utilizing the principles of the Physical Internet (PI). The study moves beyond traditional route optimization (Gonzales-Calienes et al., 2022) toward a three-layered architectural redesign focused on modularity, federated mesh topologies, and digital orchestration. By applying the logic of standard encapsulation to hazardous materials, the framework provides a pathway for decoupling battery technical heterogeneity from handling infrastructure (Sallez et al., 2016). Furthermore, the integration of the Digital Battery Passport (DBP) as the „identity backbone" enables automated, data-driven decision-making within the network (Soavi et al., 2026). The findings suggest that the transition from closed, proprietary chains to an open, PI-enabled ecosystem can significantly enhance network resilience and asset utilization while ensuring strict compliance with the safety and traceability requirements of Regulation (EU) 2023/1542.

5.2 Managerial and Policy Recommendations

To facilitate the systemic transition toward a Physical Internet (PI)-based reverse logistics network, industry stakeholders and policymakers must adopt a synchronized, multi-dimensional strategic approach.

Establishing standardized HazMat π-containers constitutes a critical prerequisite for achieving system-wide interoperability. Consequently, industry consortia should prioritize the design and adoption of certified, sensor-integrated modular units specifically tailored for hazardous battery flows. Such standardization is essential to enable universal handling across diverse transport modes and operational environments (Sallez et al., 2016; Montreuil et al., 2013).

In parallel, a strategic transition is required to move away from rigid, hierarchical buyer-supplier relations toward federated governance models. This shift is necessary to manage the inherent complexities of cross-border end-of-life (EoL) battery diagnostics (Duong et al., 2026). The implementation of decentralized technologies, such as blockchain, ensures data security and immutability, providing the requisite trust for sharing sensitive diagnostics across borders (Betti et al., 2019; Duong et al., 2026).

Furthermore, organizations should leverage the data embedded within Digital Battery Passports (DBP) to implement automated Decision Support Systems (DSS) for product disposition (Bina,

2025; Gianvincenzi et al., 2025). This integration serves to mitigate administrative bottlenecks and reduce the prohibitive manual testing costs that currently constrain the economic viability of second-life and recycling operations (Altuntas Vural et al., 2025; Bina, 2025).

Ultimately, attaining European circular autonomy necessitates proactive policy intervention focused on incentivizing the creation of open π-nodes. Expanding regional recycling capacity through these open nodes is crucial to reduce structural dependencies on high-risk third countries for critical raw materials, thereby aligning operational practices with the European Union's strategic objectives for supply chain resilience (Gianvincenzi et al., 2025).

5.3 Future Research Directions

While the proposed framework establishes a structural foundation for the systemic redesign of European battery reverse logistics, several avenues for further scholarly inquiry remain essential to bridge the gap between conceptual design and operational reality.

A primary avenue for further inquiry involves empirical validation through computational simulation. Future studies should employ agent-based modeling (ABM) to rigorously test the proposed hypotheses ($\boldsymbol{H_1}$ & $\boldsymbol{H_2}$) across diverse scenarios of demand volatility and logistical disruptions. Such simulations are critical to understanding how the federated mesh network behaves under stress and to refine the PI protocol parameters (Shaikh et al., 2024; Sallez et al., 2016).

Beyond empirical testing, the economic viability of multi-stakeholder PI-mesh networks warrants rigorous examination. There is a pressing need for a systematic analysis of cost-sharing and revenue-allocation mechanisms. Developing equitable incentive structures is necessary to facilitate participation among competing Original Equipment Manufacturers (OEMs), who must balance the benefits of open logistics with the inherent risks of co-opetition (Altuntas Vural et al., 2025; Chirumalla et al., 2024).

Additionally, comparative Life Cycle Assessments (LCA) are necessary to provide quantitative rigor. Further studies should quantify the net environmental savings of PI-based „proximity-routing" for hazardous battery flows compared to traditional, centralized recovery models (Gianvincenzi et al., 2025; Soavi et al., 2026).

Ultimately, the synergy between PI structural principles and digital traceability offers a scalable blueprint for achieving a resilient and sustainable European battery ecosystem by 2030, successfully aligning industrial operations with high-level circular economy imperatives (Gianvincenzi et al., 2025; Montreuil, 2011).

Acknowledgement

The research leading to these results has received funding from European Union's Horizon Europe Programme under the BatteReverse project, grant agreement n°101103706. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Commission. Neither the European Union nor the granting authority can be held responsible for them.

References

- Albertsen, L., Richter, J. L., Peck, P., Dalhammar, C., & Plepys, A. (2021). Circular business models for electric vehicle lithium-ion batteries: An analysis of current practices of vehicle manufacturers and policies in the EU. Resources, Conservation & Recycling, 172, 105658.

- Altuntas Vural, C., van Loon, P., Halldórsson, Á., Fransson, J., & Josefsson, F. (2025). Life after use: circular supply chains for second-life of electric vehicle batteries. Production Planning & Control, 36(9), 1229–1246.
- Bina, L. (2025). The Role of the Digital Product Passport in Unlocking Second-Life Opportunities for Batteries. [Master's thesis, Politecnico di Milano].
- Chirumalla, K., Kulkov, I., Parida, V., Dahlquist, E., Johansson, G., & Stefan, I. (2024). Enabling battery circularity: Unlocking circular business model archetypes and collaboration forms in the electric vehicle battery ecosystem. Technological Forecasting & Social Change, 199, 123044.
- Duong, Q. H., Arranz, C. F. A., Xu, M., Zhou, L., & Sun, W. (2026). Towards Web3- and metaverse-enabled decentralisation of electric vehicle battery closed-loop supply chains: insights from advanced text mining techniques. Transportation Research Part E, 206, 104583.
- Gianvincenzi, M., Marconi, M., Mosconi, E. M., & Tola, F. (2025). Toward a resilient European battery ecosystem by 2030: Strategic pathways to meet EU regulatory targets. Journal of Energy Storage, 132, 117851.
- Gonzales-Calienes, J., et al. (2022). Development of a Reverse Logistics Modeling for End-of-Life Lithium-Ion Batteries and Its Impact on Recycling Viability - A Case Study to Support End-of-Life Electric Vehicle Battery Strategy in Canada. Sustainability, 14, 15321.
- Guo, D., & Mantravadi, S. (2024). Improving Hyperconnected Logistics with Blockchains and Smart Contracts. [Article/Source Title as per library metadata].
- Larkomaa, A. (2024). DIGITAL BATTERY PASSPORT Data collecting and data value creation. [Master's thesis].
- Montreuil, B. (2011). Toward a Physical Internet: meeting the global logistics sustainability grand challenge. Logistics Research, 3(2-3), 71-87.
- Montreuil, B., Rougès, J. F., Cimon, Y., & Poulin, D. (2012). The Physical Internet and Business Model Innovation. Technology Innovation Management Review, 2(6), 32-37.
- Nakechbandi, M., Colin, J.-Y., & Mathieu, H. (n.d.). Management of mobile resources in Physical Internet logistic models.
- Pan, S., Ballot, E., George, Q. H., & Montreuil, B. (2017). Physical Internet and interconnected logistics services: research and applications. International Journal of Production Research, 55(9), 2603-2609.
- Rahic, M., & Vu, F. (n.d.). Exploring EV Battery Secondary life Business models and Reverse logistic perspectives. [Master's thesis, Mälardalen University].
- Rouhani, S., Wardley, L. J., & Amin, S. H. (2025). A comprehensive survey into reverse logistics and Closed-Loop Supply Chain aspects to provide analyses and insights for implementation. Journal of Cleaner Production, 490, 144743.
- Sallez, Y., Pan, S., Montreuil, B., Berger, T., & Ballot, E. (2016). On the activeness of intelligent Physical Internet containers. Computers in Industry, 76, 1-9.
- Shaikh, S. J., Montreuil, B., Hodjat-Shamami, M., & Gupta, A. (2018). Introducing Services and Protocols for Inter-Hub Transportation in the Physical Internet. [Technical Report].
- Shekarian, E. (2020). A review of factors affecting closed-loop supply chain models. Journal of Cleaner Production, 253, 119823.
- Soavi, F., Gregucci, A., Liverani, A., Mojtahedi, S., Petri, E., Mascetti, F., Capodarca, F., & Bel Hadj Jrad, E. (2026). The European Battery Regulation and Digital Battery Passport: Prospects and Challenges.
- Treiblmaier, H., Mirkovski, K., Lowry, P. B., & Zacharia, Z. G. (2020). The physical internet as a new supply chain paradigm: a systematic literature review and a comprehensive framework. The International Journal of Logistics Management, 31(2), 239-287.

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE
Physical Internet Center

Enabling scalable and resilient Supply Chains through Hyperconnected Mobile Supply Chains

Julien MAURICE[1], Benoit MONTREUIL[1,2,3], Walid KLIBI[4], Olivier LABARTHE[4] ,Eunhye SONG[1,2] and Todd Ullom[5]
1. Physical Internet Center, Supply Chain & Logistics Institute, Georgia Tech, Atlanta, USA
2. H. Milton Stewart School of Industrial and Systems Eng., Georgia Tech, Atlanta, USA
3. Coca-Cola Chair in Material Handling & Distribution, Georgia Tech, Atlanta, USA
4. CESIT, Kedge Business School, Bordeaux, France
5. MiTek.Inc, Chesterfield, USA
Corresponding author: julien.maurice@gatech.edu

Abstract: *The need to serve geographically dispersed and uncertain demand sustainably has increased interest in containerized and mobile production systems. Unlike traditional centralized manufacturing, containerized production encapsulates modular production lines within transportable units, enabling rapid deployment and relocation. While industrial initiatives have demonstrated benefits such as reduced capital cost, energy consumption, and footprint, most existing research focuses on static or single-tier systems and does not capture the complexity of multi-tier mobile production networks. This paper studies the design of a four-tier supply chain network composed of suppliers, consolidation facilities, production facilities, and distributed demand points. Consolidation and production tiers rely on modular mobile units that can be purchased, deployed, relocated, stored, or sold over a multi-period planning horizon. Strategic decisions include facility selection, activation over time, and capacity sizing, while tactical decisions include module relocation and material flow allocation. We adopt a simulation–optimization framework to evaluate network designs under demand uncertainty. To explore the high-dimensional discrete design space, we adapt the Adaptive Hyperbox Algorithm. Experiments on realistic instances show significant cost reductions and improved service robustness compared with baseline search methods.*

Keywords: *Physical Internet, Hyperconnected Mobile Supply Chains, Simulation-Optimization, Multi-Tier Networks, Modular Capacity, Distributed Construction.*

Physical Internet (PI) Roadmap Fitness*: ☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories) ☒ PI Networks, ☒ System of Logistics Networks*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

1 Introduction

In a world increasingly affected by environmental, political, and societal instabilities, the need to serve geographically dispersed and uncertain demand in an economically and environmentally sustainable manner has led to the emergence of containerized and mobile production systems (Ivanov, 2022). Unlike traditional centralized manufacturing networks, mobile production systems leverage the concept of containerized production units that encapsulate entire production lines within modular, transportable containers, enabling rapid deployment and relocation across the supply network.

A prominent demonstration of this concept was the *F3 Factory Project (2009–2014)* (Figure 1), a large European initiative to demonstrate the economic and environmental benefits of the mobile production system in various industries. The project reported several quantitative benefits, including reductions in capital and operational expenditure, energy consumption, and

physical footprint, alongside improvements in production yield and capacity. These results highlight the transformative potential of containerized systems in enhancing modularity, scalability, and energy efficiency within production networks.

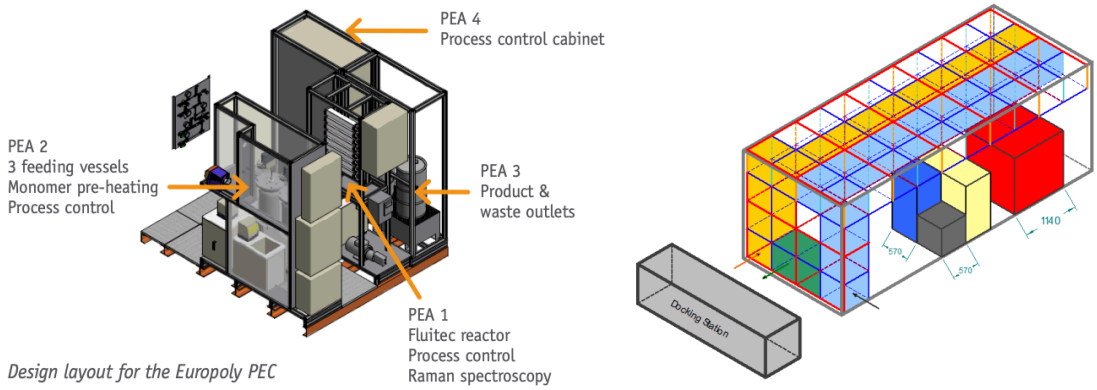


Figure 1 Containerized production unit design from F3 Factory Project Final Report. (Left) Example of design of containerized production for Europoly project. (Right) F3 Factory design guidelines and standards for modular, container based production units

Building on these developments, companies such as Bayer, Unilever, PureAqua or Modul'Data have developed commercial containerized production units across a range of industries, including chemicals, agri-food industry, water treatment, and data centers (Figure 2) These systems address the long-standing challenges of flexibility, adaptability, and robustness that characterize traditional supply chains, which often rely on a small number of fixed production sites (Jabbarzadeh et al., 2016). In addition, mobile and containerized production units support the transition toward sustainable, localized manufacturing by reducing transportation needs and associated environmental impacts.

Despite this industrial progress, research on the design and operation of multi-tier mobile supply chains remains limited. Existing work primarily focuses on single-tier mobile manufacturing or assumes static facility configurations, without considering dynamic inter-tier interactions or capacity relocation. This gap motivates the present research, which investigates multi-tier mobile supply chain networks where multiple layers of mobile units and facilities interact dynamically under uncertainty. The objective is to develop a mathematical modeling and simulation-optimization framework that captures these interdependencies and assesses how configurational flexibility can improve both economic and environmental performance.



Figure 2 : Various application of Mobile Production Facility. Unilever Nano Factory (Top Left), PureAqua Water Treatment Mobile Plant (Top Right), Modul'Data Center (Bottom Left), Bayer Mobile Production Unit (Bottomg Right)

2 Related Literature

Designing resilient and efficient supply networks under uncertainty remains a central challenge in operations research. While traditional deterministic optimization approaches provide computational tractability, they fail to capture stochastic dynamics of real-world systems, leading to poor performance when their solutions are implemented in practice. Simulation–Optimization addresses this limitation by embedding stochastic evaluation directly within the search process, yielding more robust solutions.

A common class of Simulation-Optimization approaches integrates simulation within heuristic or metaheuristic search procedures. For example, Pagès-Bernaus et al. (2019) embedded Monte Carlo simulation into Iterated Local Search framework, showing improved solution quality. When the decision space involves integer-valued variables, such as facility activation or discrete capacity selection, discrete simulation–optimization methods become necessary. In this context, the Adaptive Hyperbox Algorithm (AHA) proposed by Xu et al. (2013) addresses this challenge by iteratively exploring promising regions of the design space using adaptative hyperboxes guided by lower confidence bounds, significantly improving sampling efficiency in noisy, high-dimensional settings.

Another stream of literature focuses on surrogate modeling to reduce the computational burden of simulation. Metamodels based approach approximate the simulation response surface, enabling faster exploration of the design space. Osorio et al. (2015) pioneered such techniques in transportation systems, with subsequent extensions to logistics and digital twin environments. (Cao et al., 2024) However, as stated by Soares et al. (2022) and Snoeck and Winkenbach (2022) these methods are less effective in discrete settings or when response functions exhibit discontinuities, motivating the use of enhanced discrete simulation-optimization techniques such as MetaAHA.

In parallel, recent research has explored more flexible production and supply chain architectures. Traditional supply network design models typically assume fixed facilities, whereas emerging paradigms such as Distributed Manufacturing Systems (Matt et al., 2015) and Mobile Supply Chains (Shahmoradi-Moghadam and Schönberger, 2021), Alman et Zhang (2020) introduce flexibility through modular and relocatable production assets. Review by Alarcon-Gerbier and Buscher (2022) synthesize these developments, and industrial prototypes such as Adamietz et al. (2018) demonstrate the feasibility of containerized configuration plants. However, most existing work focuses on simplified settings, such as single-tier systems or mobile depot routing problems. Comprehensive multi-tier networks under uncertainty remain largely unexplored.

The Physical Internet paradigm has emerged a shift for addressing the economic, environmental, and social inefficiencies of current logistics systems. Montreuil B. (2011). The Physical Internet promotes open, standardized and interconnect logistics networks based on resource sharing, modularization, and interoperability. Building on these principles, Marcotte and Montreuil (2016) introduces the concept of Hyperconnected Mobile Production, in which plug-and-play production modules can be dynamically deployed and reconfigured across networks of open, certified facilities. Recent works have extended this concept to multi-tier supply chains, demonstrating its potential to improve cost efficiency, adaptability, and sustainability through dynamic capacity deployment and resource pooling. (Maurice et al., 2024, 2025) However, the integration of uncertainty and dynamic interactions across multiple tiers in such systems remains largely unexplored.

Motivated by this gap, this work adapts the MetaAHA simulation–optimization framework to a multi-tier mobile supply network design problem under stochastic demand and operations,

demonstrating how simulation-driven search uncovers solutions that deterministic approaches fail to identify.

3 Problem Setting

We consider a multi-tier supply chain integrating mobile production resources within a distributed manufacturing environment, following the Hyperconnected Mobile Supply Chain (HMSC) paradigm introduced in Maurice et al (2024). In this paradigm, production capacity is no longer tied to a fixed facilities and is embodied in mobile, modular unsighted that can be dynamically deployed and shared across a network of open and certified facilities.

The system is structured as a four-tier network in which each tier represents successive stages in the transformation and delivery of products from suppliers to final demand, as illustrated in Figure 3. A key feature of the HMSC paradigm is that multiple tiers of mobile production resources interact dynamically: upstream mobile consolidation units and downstream production modules are jointly deployed and coordinated across facilities, forming an interconnected and reconfigurable production network. Demand arises at geographically dispersed markets and must be satisfied over time. Production relies on mobile, relocatable modules deployed at candidate facilities, while upstream consolidation facilities assemble components into kits supplied to production sites.

At the strategic level, the decision-maker determines the configuration of the network over time, including facility activation, capacity deployment across tiers, and flow assignment across tiers. The network evolves dynamically across discrete time periods, reflecting the relocation and reconfiguration of mobile production resources, and is subject to flow conservation and capacity constraints, ensuring consistency between inbound and outbound quantities throughout the network. The problem is formulated as a bi-objective optimization model to deal with the trade-off between minimizing total economic cost and greenhouse gas emissions, accounting for facility operations, capacity deployment, and transportation activities.

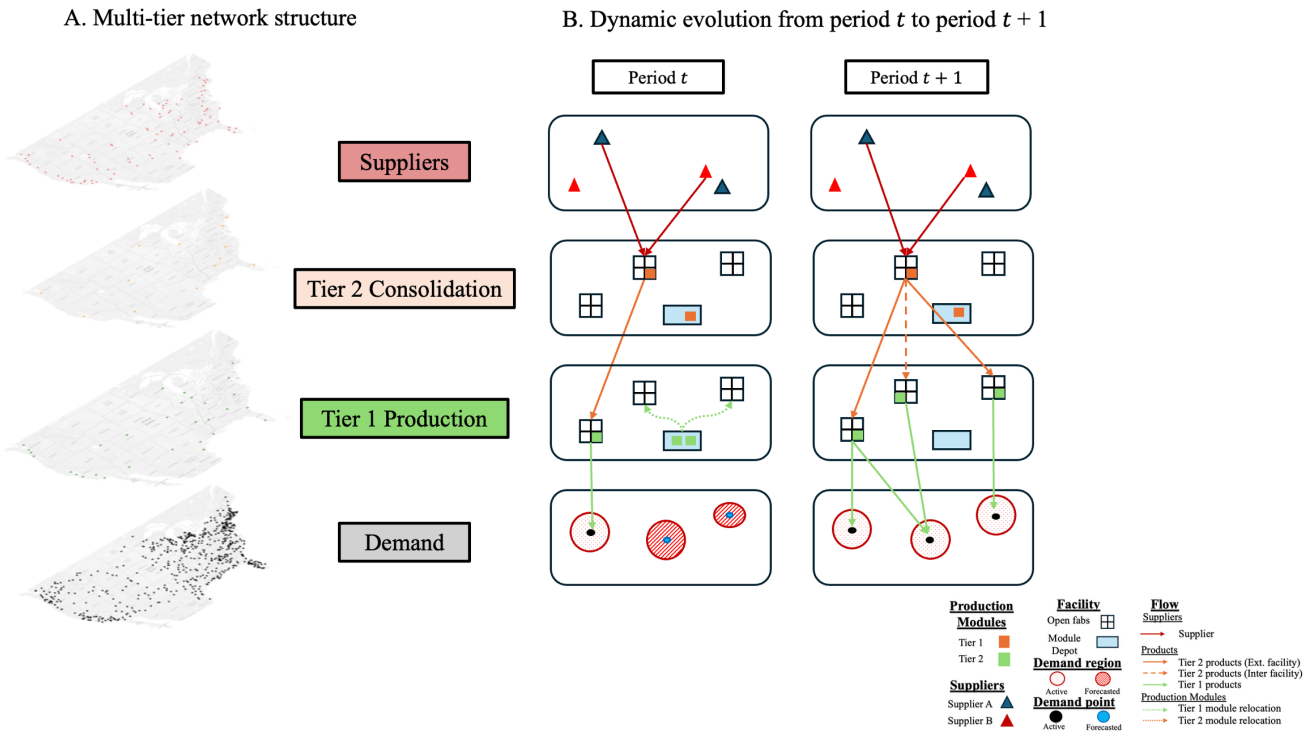


Figure 3: Illustration of the Hyperconnected Mobile Supply Chain Network with its evolution between time periods

4 Methodology

This section presents the simulation-optimization framework developed to solve the multi-period mobile supply chain network design problem formulated in Maurice et al (2024). The combinatorial nature of the time-indexed decision space and the intractability of exact methods under demand uncertainty motivate the proposed approach, which integrates a strategic mixed-integer linear program with a stochastic operational simulator within the MetaAHA algorithm of Snoeck and Winkenbach (2022), itself an extension of the Adaptive Hyperbox Algorithm of Xu et al (2013).

A candidate design is represented by the strategic decision vector $x = (Y, N^1, N^2, S)$, encoding the network configuration across all facilities $f \in \mathcal{F}$, periods $t \in T$, and suppliers $s \in \mathcal{S}$:

$$x = \{0,1\}^{|T|\times|\mathcal{F}|} \times \mathbb{Z}_{\geq 0}^{|T|\times|\mathcal{F}|} \times \mathbb{Z}_{\geq 0}^{|T|\times|\mathcal{F}|} \times \{0,1\}^{|T|\times|\mathcal{S}|}$$

Where $Y_{t,f} \in \{0,1\}$ is the facility activation flag, $N_{t,f}^k \in \mathbb{Z}_{\geq 0}$ is the Tier-k module capacity cap at facility f in period t ($k \in \{1,2\}$), and $S_{t,s} \in \{0,1\}$ indicates whether supplier s is active in period t. These are the strategic decisions set before the horizon begins. All operational decisions, module acquisition, deployment, inter-facility relocation, depot storage, and end-of-horizon liquidation, are delegated to the simulator, which resolves them period by period in response to realized demand.

Given a candidate x, performance is estimated by simulating the deployment over a set of scenarios $\omega \in \Omega$, each representing a realisation of demand noise, facility disruptions, and relocation uncertainty. Within each scenario and period t, the simulator: (i) determines available capacity from the current fleet and facility states; (ii) acquires new modules if the fleet is insufficient to meet expected demand; (iii) assigns demand to open facilities via a nearest-facility greedy heuristic, within the 300-mile service radius; (iv) propagates kit demand upstream to Tier-2 hubs, which source components from the nearest active supplier; and (v) executes inter-period module relocations to rebalance the fleet between facilities.

The expected profit across scenarios is:

$$\overline{\Pi}(\mathrm{x}) \approx \frac{1}{|\Omega|}\sum_w \sum_t (\rho D_t^{served} + \sum_{k,f} v^k N_{t,f}^{sell,k} - C_t^{op} - C_t^{tr} - C_t^{stor} - C_t^{reloc})] \quad (1)$$

where ρ is unit revenue, v^k is the salvage value of a liquidated Tier-k module, and $\$C^{op}, C^{tr}, C^{stor}, C^{reloc}$ are operating, transport, storage, and relocation costs respectively. The bi-criteria scalar objective is $J(x) = -\overline{\Pi}(\mathrm{x}) + \lambda \bar{E}(x)$, where $\lambda \geq 0$ is a carbon weight. At $\lambda = 0$, the objective reduces to pure profit maximization and, increasing λ traces the Pareto frontier over $(\overline{\Pi}(\mathrm{x}), \bar{E}(x))$.

The search follows the MetaAHA framework Snoeck and Winkenbach (2022). A single hyperbox $\mathcal{H}_k$ is maintained at each iteration k, centered on the current incumbent x_k^*, with bounds tightened component-wise as more designs are evaluated. Two candidate generation mechanisms operate at each iteration. First, B designs are sampled uniformly within $\mathcal{H}_k \cap \mathcal{X}$. Second, the MILP is re-solved with hyperbox bounds as side constraints and time limit τ, yielding a structured candidate within the current promising region. All candidates are evaluated stochastically, added to $\mathcal{L}$, and compared on J. The algorithm terminates after $\bar{s}$ consecutive iterations without improvement. The full procedure is given in Algorithm 1.

Algorithm 1 MetaAHA simulation-optimization framework

Input: Instance, λ, batch size B, scenario set Ω, stall limit $\bar{s}$, MILP limit τ
Output: $x^\star$ with estimated profile $(\hat{\Pi}(x^\star), \hat{\hat{E}}(x^\star))$

Solve MILP; set $x_0 \leftarrow x_0^{\text{MILP}}$; estimate $J_\lambda(x_0)$ over Ω $x^\star \leftarrow x_0$, $J^\star \leftarrow J_\lambda(x_0)$, $\mathcal{L} \leftarrow \{x_0\}$, stall $\leftarrow 0$

while *budget not exhausted* ***and*** *stall* $< \bar{s}$ **do**
  // Hyperbox update
  $l^{(d)} \leftarrow \max\{y^{(d)} \in \mathcal{L} : y^{(d)} < x^{\star,(d)}\}$; $u^{(d)} \leftarrow \min\{y^{(d)} \in \mathcal{L} : y^{(d)} > x^{\star,(d)}\}$ // Candidate generation (within $\mathcal{H}_k$)
  Sample B designs from $\mathcal{H}_k \cap \mathcal{X}$; add to $\mathcal{B}$ Solve MILP s.t. $l^{(d)} \le y^{(d)} \le u^{(d)}$, limit τ; add x^{MILP} to $\mathcal{B}$ // Stochastic estimation ($|\Omega|$ paths each)
  foreach $x_i \in \mathcal{B}$ **do**
    Estimate $J_\lambda(x_i)$ over Ω; $\mathcal{L} \leftarrow \mathcal{L} \cup \{x_i\}$
  // Incumbent update
  if $\min_{\mathcal{B}} J_\lambda < J^\star$ **then** $x^\star \leftarrow \arg\min_{\mathcal{L}} J_\lambda$; stall $\leftarrow 0$
  else stall $\leftarrow$ stall $+1$

return $(x^\star, \hat{\Pi}(x^\star), \hat{\hat{E}}(x^\star))$

Algorithm 1: MetaAHA simulation optimization framework pseudo code

5 Case Study

5.1 Problem instance

The framework is applied to a multi-period Hyperconnected mobile supply chain network design problem leveraging real industry data from hotel construction projects across the continental United States. The demand set comprises $|\mathcal{A}| = 288$ markets corresponding to hotel properties scheduled for delivery in 2026 across 41 states, with demand measured in rooms to be outfitted per period (Figure 4).

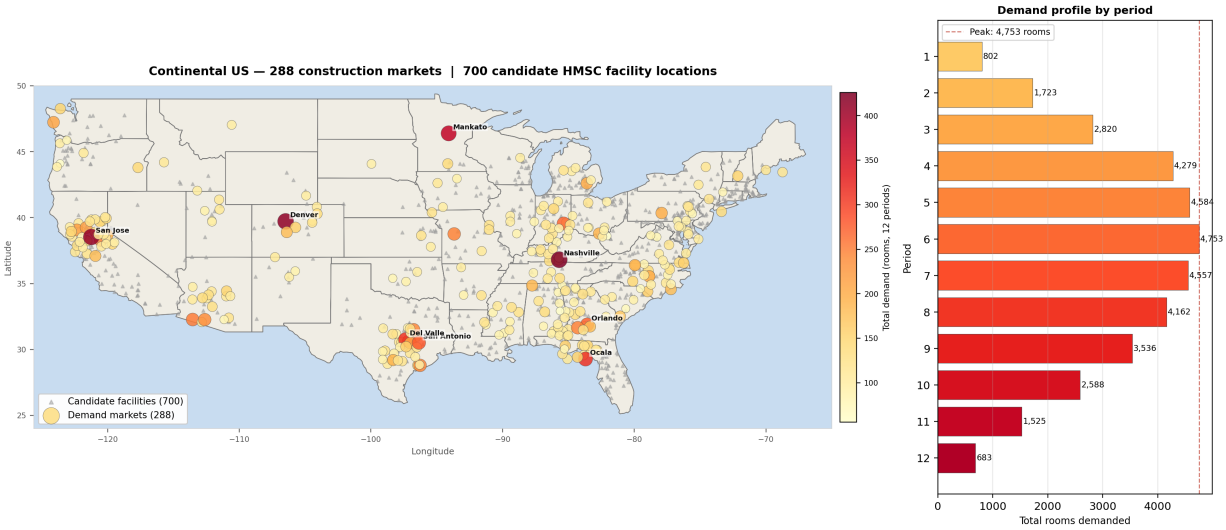


Figure 4: 288 hotel construction markets and 700 candidate HMSC facility locations. Right panel: aggregate demand profile by period,

A total of $|\mathcal{F}| = 700$ candidate facility locations are drawn from existing logistics infrastructure distributed across the continental US, selected from 781 candidate sites to maximise geographic coverage while ensuring every market is reachable within a 300-mile service radius. The planning horizon covers $|T| = 12$ periods of equal duration. Each market follows a project life cycle: demand evolves according to a sinusoidal profile over an active window that varies by project, reflecting typical construction ramp-up and wind-down dynamics. Aggregate demand peaks at approximately 4,754 rooms in period 6, with a total of 36,020 rooms over the full horizon.

The supply chain comprises Tier-1 production modules (capacity 25 rooms per module per period) and Tier-2 support modules (capacity 25 rooms per module per period), each subject to

facility-specific deployment caps, and deployed in tandem. A central depot with capacity ($Q^{depot} = 30$ modules) supports inter-period storage. The economic and stochastic parameters are summarized in Table 1. Revenue reflects a premium furnished room. Distance-based relocation costs create a genuine trade-off between fleet mobility and geographic consolidation. Salvage values are below acquisition cost to incentivize redeployment over liquidation except in low-demand periods. Demand uncertainty is modeled with a coefficient of variation $\sigma_d = 0.20$, reflecting realistic month-to-month variability in hotel renovation scheduling.

Table 1: Economic and operational parameters

Parameter	*Unit*	*Description*
ρ	\$k/room	Revenue per room served
$c_{capex\,1}$	\$k/module	Tier-1 acquisition (one-time)
$c_{capex\,2}$	\$k/module	Tier-2 acquisition (one-time)
$c_{op\,1}$	\$k/module/period	Tier-1 operating cost
$c_{op\,2}$	\$k/module/period	Tier-2 kit hub operating cost
b_1, d_1	\$k; \$/mi	Tier-1 relocation (base + dist.)
b_2, d_2	\$k; \$/mi	Tier-2 relocation (base + dist.)
v_1	\$k/module	Tier-1 salvage value
v_2	\$k/module	Tier-2 salvage value
h_1	\$k/module/period	Tier-1 depot holding cost
h_2	\$k/module/period	Tier-2 depot holding cost
c_{sup}	\$/kit/km	Supplier to Tier-2 transport
Q_{depot}	30 modules	Central depot capacity
σ_d	-	Demand volume CV
p_{full}	-	Full facility outage probability
$p_{partial}$	-	Partial capacity outage probability
$e_{op\,1}$	kg CO2eq/mod.	Tier-1 operating emissions
$e_{op\,2}$	kg CO2eq/mod.	Tier-2 operating emissions
e_{tr}	kg CO2eq/mile	Module transport emissions

5.2 Algorithmic Configuration

MetaAHA is configured with $B = 10$ candidate designs per iteration evaluated against a fixed common random numbers pool of $|\Omega| = 30$ scenarios shared across all candidates. The solution space is encoded as a 1,080-dimensional compact vector with 360 dimensions per echelon (Tier-1, Tier-2, and supplier selections). The AHA hyperbox simultaneously bounds all three echelons, confining the search to geographically coherent configurations. The strategic MILP candidate is generated at each iteration with a time limit of $\tau = 300$ seconds and a MIP gap tolerance of 5%. Final performance is reported over 200 independent Monte Carlo scenarios drawn from the same demand distribution ($\sigma_d = 0.20$). A service level threshold of 60% is used to assess solution acceptability across scenarios.

5.3 Managerial Insights

To evaluate the operational value of mobility, we compare the MetaAHA deployment strategy, as implemented in MiTek's Distributed Construction™ Methodology, evaluated across 200 independent demand scenarios against a conventional static modular construction network composed of only three Tier 1 production facilities located in Atlanta, Dallas and Los Angeles, each dimensioned to serve the mean demand of markets within its 300-mile area. The objective is to assess whether dynamic module deployment improves market responsiveness, capital utilization, and robustness under uncertain demand. Table 2 reports the resulting economic and operational performance of both network configurations.

Table 2: HMSC mobile network vs. static fixed network (local demand-based sizing across 3 hubs).

	Expected Profit	Cost/Revenue Ratio	Market Coverage	Global Warming Potential (kt CO_2eq)
Static	-	0.89	40%	91
HMSC Mobile	+118.9%	0.85	97.4%	132

The comparison reveals two distinct effects of network mobility. First, the geographic reach. The mobile network covers 97.4% of the markets versus 40% for the static network, showing the ability to deploy modules where construction activity is highest. Second, capital efficiency. As illustrated in Figure 5, the HMSC network deploys 134 modules progressively as demand builds, acquiring 41 modules at peak in period 4, performs, a total of 131 inter-facility relocations occur across the horizon as construction demand shifts geographically and liquidating 56 modules at horizon end, while the static network commits to its capacity from period 1 regardless of demand. Despite deploying more production capacity overall, HMSC achieves a near-identical cost-to-revenue ratio, demonstrates that inter-facility relocation, progressive deployment and end-of-horizon liquidation effectively amortize the large capital investment across the planning horizon. The combined effect yields a profit improvement of 118.9% over the static baseline. Beyond profitability, the framework exposes the trade-off between financial performance and environmental impact. Running the framework across a grid of carbon weights $\lambda \in \{0, 10, 50, 100, 200, 500\}$ traces a Pareto frontier over the $(\overline{\Pi}, \bar{E})$ space. At $\lambda = 0$, the profit-maximising plan generates a GWP of approximately 132 kt CO_2eq over the horizon. Decision-makers can select any point on the frontier to reflect their organization's carbon pricing or sustainability commitments, with the framework quantifying the exact profit cost of each reduction target. The evolution of the network between period 6 and 8 is illustrated in Figure 6.

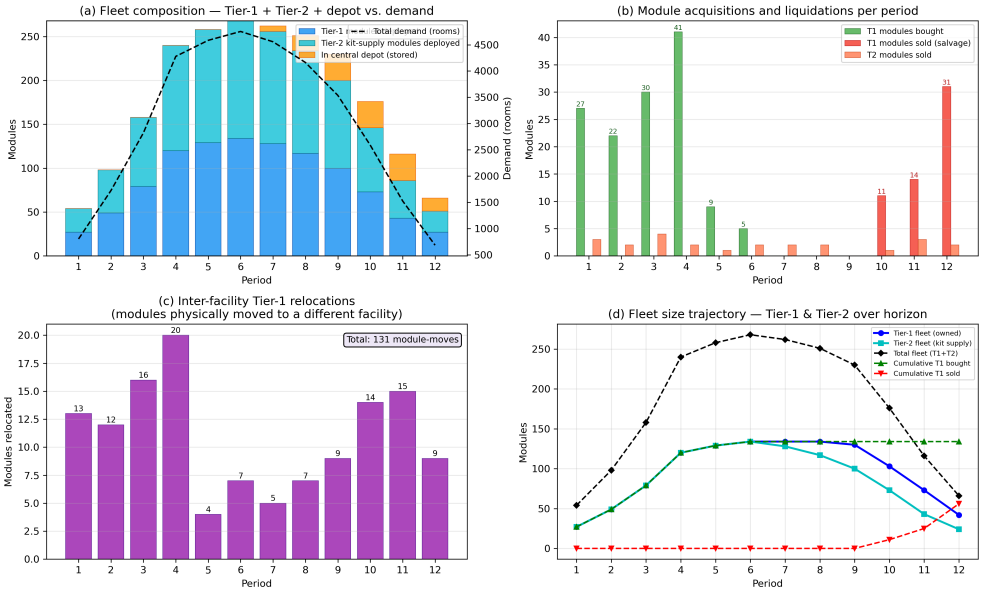


Figure 5 : Multi-tier HMSC modules lifecycle: fleet size over time (a), monthly acquisitions and liquidations (b), inter-facility relocations (c), and cumulative fleet trajectory (d)

The HMSC deployment plan remains profitable in all 200 demand scenarios tested, including those with facility outages and demand surges, with a coefficient of variation of 4.8%.This robustness is the core value of the simulation-optimization approach: rather than a plan optimized under mean demand that underperforms when reality deviates from forecasts, the framework produces a design stress-tested against hundreds of realistic futures before deployment.

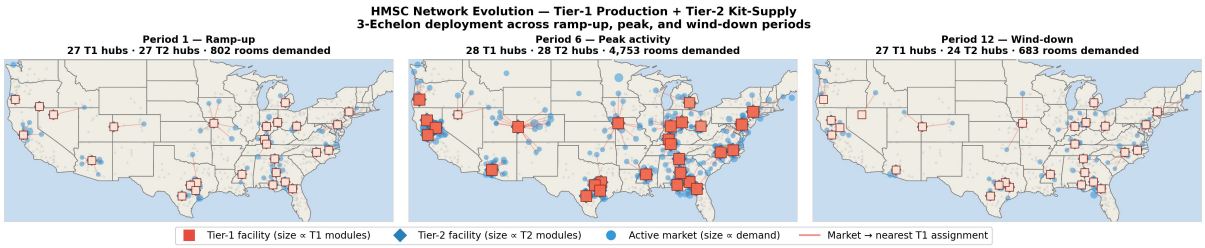


Figure 6: Evolution of the HMSC network between period 6 and 8 (Tier 1 only)

To rule out geographic bias and assess capital commitment, Experiments 3 and 4 stress-test the HMSC advantage. Experiment 3 restricts both methods to the 112 markets within the static network's 300-mile service areas, evaluating whether dynamic redeployment creates value independently of broader geographic reach. HMSC still outperforms on every indicator, generating 24.3% higher profit and coverage. Experiment 4 tests what a traditional static network would require to match HMSC's 97.4% coverage: a set-cover optimization identifies a minimum of 10 fixed hubs, committing 222 modules from period 1 regardless of demand, requiring 66% more capital than HMSC's 134-module peak fleet, and yet yielding less than half the profit and 87% higher emissions. Table 3 summarizes both comparisons.

Table 3: Robustness of HMSC advantage — within-footprint (Exp. 3) and coverage parity (Exp. 4).

		Expected Profit	Cost/Revenue Ratio	Market Coverage	Global Warming Potential (kt CO_2eq)
Exp. 3	Static	$21.2M	0.89	100%	70
	HMSC	$26.4M	0.85	99.5%	72
Exp. 4	Static	$21.9M	0.961	94.1%	247
	HMSC	$48.3M	0.85	97.4%	132

Together, Experiments 3 and 4 establish that HMSC mobility creates value both within and beyond the static network's geographic footprint: dynamic deployment yields higher profit and coverage even on the same markets, and achieves comparable coverage at 66% lower capital commitment and 47% lower emissions than a static alternative.

6 Conclusion

This paper presented MetaAHA, a simulation-optimization framework for multi-period Hyperconnected Mobile Supply Chain network design under demand and operational uncertainties as envisioned by MiTek's Distributed Construction Methodology. By combining a strategic MILP with a full operational simulator across a three-echelon network, the framework produces deployment plans that are stress-tested against realistic demand futures rather than optimized for a single forecast.

Applied to our case study data, MetaAHA identifies a plan generating that is 118.9% more profitable than the static case and covering 97.4% of the market, at a lower cost-to-revenue ratio while limiting the induced Global Warming Potential. Experiments 3 and 4 further confirm this advantage holds within the static netwrok's own geographic footprint and that achieving comparable coverage statically would require 66% more capital and produce 87% higher emissions. The key practical finding is that HMSC mobility creates value only when exercised strategically: the deployment plan anticipates demand shifts rather than reacting to them.

For future research, a formal validation on smaller instances against an exact solver would establish the optimality gap of MetaAHA and provide theoretical guarantees currently absent from the heuristic framework.

References

- Adamietz R., T. Giesen, P. Mayer, A. Johnson, R. Bibb, C. Seifarth (2018): Reconfigurable and transportable container-integrated production system. *Robotics and Computer-Integrated Manufacturing*, v53, 1–20.
- Alarcon-Gerbier E., U. Buscher (2022): Modular and mobile facility location problems: A systematic review. *Computers & Industrial Engineering*, v173, 108734.
- Allman A., Q. Zhang (2020): Dynamic location of modular manufacturing facilities with relocation of individual modules. *European Journal of Operational Research*, v286, no2, 494–507.
- Cao X., M. Yao, Y. Zhang, X. Hu, C. Wu (2024): Digital Twin Modeling and Simulation Optimization of Transmission Front and Middle Case Assembly Line. *Computer Modeling in Engineering & Sciences*, v139, no3, 3233–3253.
- Ivanov D. (2022): Viable supply chain model: integrating agility, resilience and sustainability perspectives—lessons from and beyond the COVID-19 pandemic. *Annals of Operations Research*, v319, no1, 1411–1431.
- Jabbarzadeh A., B. Fahimnia, J.-B. Sheu, H. S. Moghadam (2016): Designing a supply chain resilient to major disruptions and supply/demand interruptions. *Transportation Research Part B: Methodological*, v94, 121–149.
- Marcotte S., B. Montreuil (2016): Introducing the Concept of Hyperconnected Mobile Production. *Proceedings of the 14th IMHRC*, 2016.
- Matt D. T., E. Rauch, P. Dallasega (2015): Trends towards Distributed Manufacturing Systems and Modern Forms for their Design. *Procedia CIRP*, v33, 185–190.
- Maurice J., S. S. Kwon, B. Montreuil (2025): Implementation of Hyperconnected Mobile Supply Chains for Large Scale Networks.
- Maurice J., S. S. Kwon, B. Montreuil (2024): Strategic Network Design for Hyperconnected Mobile Supply Chains.
- Montreuil B. (2011): Toward a Physical Internet: meeting the global logistics sustainability grand challenge. *Logistics Research*, v3, no2-3, 71–87.
- Osorio C., G. Flötteröd, C. Zhang (2015): A Metamodel Simulation-based Optimization Approach for the Efficient Calibration of Stochastic Traffic Simulators. *Transportation Research Procedia*, v6, 213–223.
- Pagès-Bernaus A., H. Ramalhinho, A. A. Juan, L. Calvet (2019): Designing e-commerce supply chains: a stochastic facility–location approach. *International Transactions in Operational Research*, v26, no2, 507–528.
- Shahmoradi-Moghadam H., J. Schönberger (2021): Coordinated allocation production routing problem for mobile supply chains with shared factories. *Computers & Chemical Engineering*, v155, 107501.
- Snoeck A., M. Winkenbach (2022): A Discrete Simulation-Based Optimization Algorithm for the Design of Highly Responsive Last-Mile Distribution Networks. *Transportation Science*, v56, no1, 201–222.
- Soares Do Amaral J. V., J. A. B. Montevechi, R. D. C. Miranda, W. T. D. S. Junior (2022): Metamodel-based simulation optimization: A systematic literature review. *Simulation Modelling Practice and Theory*, v114, 102403.
- Xu J., B. L. Nelson, L. J. Hong (2013): An Adaptive Hyperbox Algorithm for High-Dimensional Discrete Optimization via Simulation Problems. *INFORMS Journal on Computing*, v25, no1, 133–146.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Autonomous Transport and the Physical Internet: Towards a Shared Research Agenda

Elisah van Kempen[1], Matthias Santing[1], Hannah Onverwagt[1], Jaco van Meijeren[1]
1 TNO, The Hague, The Netherlands
Corresponding author: elisah.vankempen@tno.nl

Abstract: *The Physical Internet (PI) and Connected Cooperative Automated Mobility (CCAM) are two innovative paths within the logistics sector that share similar goals and ambitions but do not have an explicit link yet. This paper introduces a preliminary shared research agenda based on past PI and CCAM research projects, the two roadmaps and expert interviews. The research agenda points towards three focus areas: Standardization and digitization of logistics processes, Innovative planning methods for decentralization and optimization, and asset sharing and impact on business models. Within these topics, the goals of PI and CCAM are combined in a joint research agenda strengthening the two roadmaps. For both the PI and CCAM communities, this paper can stimulate closer research alignment by addressing synergies between their respective roadmaps, identifying gaps for collaboration, and supporting shared objectives and ambitions for a more sustainable and efficient logistics sector.*

Keywords: *Autonomous Transport, Connected Cooperative Automated Mobility, Physical Internet, Research Agenda, PI Nodes, Logistics Automation, Asset Management.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan: ☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories), ☒ Transportation Equipment, ☒ PI Networks, ☒ System of Logistics Networks, ☐ Vertical Supply Consolidation, ☐ Horizontal Supply Chain Alignment, ☐ Logistics/Commercial Data Platform, ☐ Access and Adoption, ☐ Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

1 Introduction

The freight transport sector is under increasing pressure to deliver high levels of efficiency, sustainability and adaptability to changing circumstances such as growing transport volumes and labour shortages. Global logistics networks remain fragmented with limited interoperability across both organizational and national borders. Both technological improvements as well as a complete rethinking of the design and operation of freight transport systems at a network level are required to address these flaws of the logistics sector. This paper discusses two major innovation paths that respond to the challenges for the logistics sector

The Physical Internet (PI) aims for an efficient and sustainable logistics system by proposing a fundamental redesign of logistics systems (Montreuil et al., 2012). In the Physical Internet Roadmap these goals are described as amongst others: standardization of processes, procedures and loading units, secure and efficient protocols, pooling of resources and services and lastly clear frameworks of benefits for every stakeholder (ALICE-ETP, 2020). The PI roadmap anticipates an evolution on the long term in which nodes and flows in the logistics network are increasingly automated and decentralized or self-organizing (ALICE-ETP, 2020).

In parallel, research on Connected Cooperative Automated Mobility (CCAM) focuses on the development, introduction and deployment of autonomous vehicles. The introduction of autonomous vehicles and digitally coordinated fleets is expected to contribute to more efficient,

safe and sustainable transport, while mitigating personnel shortages and strengthening the competitiveness of European industries. However, surprisingly in European roadmaps on CCAM (ERTRAC, 2022; ERTRAC, 2024), there is not an explicit link between the goals and ambitions of the CCAM agenda and those of the Physical Internet.

Despite their shared objectives, the PI and CCAM communities have transitioned along separate research, policy and innovation paths. CCAM pays limited attention to PI concepts, while PI initiatives sees automation mainly as an optional enabler than an important design principle. This results in a high risk that current autonomous freight solutions are being developed in closed, siloed freight systems, exactly the kind of systems that PI tries to overcome. While autonomous transport is part of PI, the PI roadmap does not explicitly address all CCAM-related dimensions, highlighting the need for closer alignment between the two domains. PI systems that do not take autonomous transport into account will not unlock the full potential of autonomous vehicles within their intelligent nodes and real-time decision making.

Therefore, this paper proposes a joint research agenda indicating research topics where the Physical Internet community and the CCAM community can find synergies and can collaborate to contribute to the shared vision of a safe, efficient, and sustainable logistics system.

2 Approach

For developing the preliminary shared research agenda on autonomous transport and the Physical Internet, we followed the steps as displayed in Figure 1. First, we started by analysing the Physical Internet Roadmap and the ERTRAC CCAM roadmap. We complemented this by material and experiences from various PI projects - a.o. SOLiD, SOL Port, Talking Trucks, Future Proof Smart Logistics (Fransen et al., 2023) - and CCAM projects - a.o. Dutch projects CATALYST Living Lab, SAVED, TKI HTSM CAT4Yards and EU projects MAGPIE and MODI (Van Kempen et al., 2025). Through analysing and discussing these in the author team we identified gaps and opportunities to bring both communities and topics together. We validated these with logistics industry experts outside of the author team. This did not lead to new topics, only more refinements on the three topics already identified. Ultimately, we summarized these in the preliminary shared research agenda. We explicitly position the shared research agenda as *preliminary* as the research topics are not exhaustive and we want to invite the IPIC community to think along into further detailing, prioritizing and identifying concrete steps to collectively build on the research agenda and close research gaps in future research.

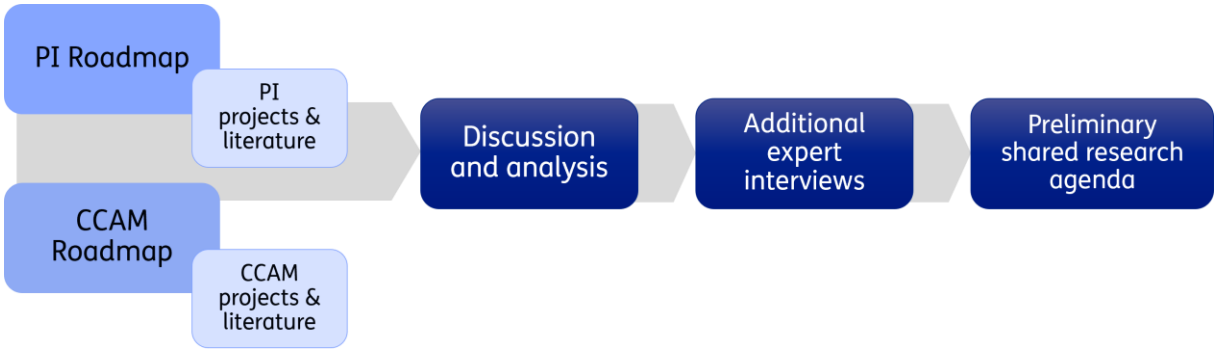


Figure 1 Approach in developing a preliminary shared research agenda

3 Preliminary shared research agenda

The three identified gaps, validated by internal and external experts are: i.) Standardization and digitization of logistics processes, ii.) Innovating planning methods for decentralization and optimization and iii.) Asset sharing and implications for business models. This chapter proposes joint future research topics for the CCAM and PI community for each of the identified gaps which are summarized in Table 1.

3.1. Standardization and Digitization of Logistics processes

The standardization and digitization of logistics processes have emerged as fundamental steps toward enabling the integration of autonomous transport, as observed across multiple applied research projects (van Kempen et al., 2026). Within the logistics sector, differences exist between companies with respect to the maturity of standardized procedures and digital processes. For example, in certain logistics yards, operations still depend on radio communication or Excel-based planning with numerous exceptions, while other yards make use of more standardized processes and advanced yard management systems. CCAM requires a mature digital environment, as driverless vehicles rely on tasks being communicated in a standardized digital format. Companies intending to move toward CCAM implementation must therefore begin by standardizing and digitizing their processes. Next to enabling CCAM possibilities, these steps can simplify workflows and give transport companies competitive and operational advantages (Mishchenko et al., 2024). According to Perman et al. (2025), the technologies that depend on standardizing data (such as AI and PI) also reduce inefficiencies.

Some future joint efforts for Physical Internet and CCAM in the development of standardization and digitization of logistics processes contain: 1.) Develop open standards for data and messaging in autonomous logistics; 2.) Create digital connections protocols between autonomous vehicles and logistics nodes, for example automated check-in/ check-out procedures; 3.) PI-compliant digital infrastructure should enable software to construct routing shipments autonomously through a network of hubs and other logistics nodes; 4.) Using the standard PI containers and modular handling systems in combination with the autonomous vehicles resulting in automated and consistent loading and unloading.

By firstly digitizing and standardizing logistics processes, therefore creating a blueprint in which autonomous systems can operate efficiently everywhere, one can enable the PI's vision of interconnected autonomous logistics nodes (ALICE-ETP, 2020).

3.2. Innovating planning methods for decentralization and optimization

The second research theme focuses on the development of innovative planning methods that support decentralization and optimization across both PI and CCAM applications. PI research mainly focuses on macro-level network structure and CCAM research often focuses on vehicle level (e.g. navigation, safety etc.). Only few studies merge these two for specific use cases, applied to, for example, the military domain (Fransen et al., 2024). In the Physical Internet, planning will become more decentralized and dynamic with many independent agents (transport companies, hubs, vehicles). Hopman et al. (2022) developed a framework to analyse trade-offs between centralized and decentralized control structures. One of their key findings is that container hinterland transport by road can be improved through self-organizing, decentralized planning approaches.

To combine the possibilities of PI and CCAM, research is needed on new planning and optimization methods that can combine the PI's open networks with autonomous transport capabilities. Currently, in the field of CCAM an opportunity exists to research how *fleets* of autonomous freight vehicles will be coordinated, as currently relatively quite some focus is on how to coordinate a single vehicle.

Effective decentralization is important for both PI and of (scalable) autonomy fleets of vehicles. This requires distributed decision-making algorithms that can evaluate, negotiate and coordinate transportation options with multiple stakeholders. Furthermore, automated planning at a logistics level is seen as an important key in Physical Internet (ALICE-ETP, 2020). Integrating vehicle automation into such planning processes has the potential to further enhance efficiency.

Within the PI, there is a need for methods to pool resources, for which decentralized planning has a high potential to increase operational performance (van Meijeren et al., 2025). However, the management of autonomous vehicles as shared, pooled resources introduces additional complexity, and achieving efficient large-scale utilization of autonomous vehicles is still underexplored. Stein and Prandtstetter (2023) suggest that the use of shared autonomous vehicles could enable more ad-hoc, on-demand transportation models, moving beyond classical transport logistics.

Next research steps with regards to this discussed gap contain: 1.) Designing multi-agent algorithms where autonomous vehicles, hubs or other nodes and shipments are agents that collaborate to a certain goal (e.g. minimize cost); 2.) Develop planning platforms or data spaces that allow secure data sharing between involved autonomous vehicles; 3.) Investigate strategies for planning automated vehicles that serve multiple companies and different loads on the same day.

Combining these research topics enables the logistics sector to transition from siloed operations to optimized, self-organizing, and autonomous logistics, supporting the next steps in PI and CCAM applications.

3.3. Asset sharing and implications for business models

The third theme in the preliminary shared research agenda focuses on asset sharing and its implications for business models. The introduction of the Physical Internet and Autonomous Transport represents is a broader transition that extends beyond technology or logistics alone. Moving toward scalable deployment of autonomous and PI-enabled logistics systems therefore requires the exploration of new logistics concepts and innovative business models.

In the CAT4Yards project the ex-ante impacts of the implementation of Connected Automated Transport at three different logistics yards was assessed (van Kempen et al., 2026). Depending on the use case and the specific CAT design, autonomous transport yields a considerably lower Total Cost of Ownership (TCO) compared to manned operations. However, investments costs might be challenging for transport companies. Asset sharing is a logistics operating model that might be part of a future with autonomous trucking (Fritschy and Spinler, 2019). Asset sharing will provide opportunities to share investment costs and optimize transport. So rather, than each transport company investing in its own autonomous trucks (which might sit idle for a part of the day), a pool of autonomous trucks could be used by various transport companies. The Physical Internet proposes to pool resources and assets in

open, connected and shared networks (ALICE-ETP, 2020). Hence, from an asset sharing perspective, the transition to autonomous transport and the Physical Internet will converge.

However, making asset sharing work raises questions regarding governance, sharing costs, risks and benefits. Ensuring that all collaborating transport companies can use the autonomous trucks at their preferred time requires advanced planning methods (see section 3.2). Developing and advancing decentralized planning methods alone will not be sufficient. Also, clear boundary conditions and agreements must be made among the participating transport companies on how costs and benefits are shared. For horizontal collaboration in logistics, governance mechanisms are important to coordinate shared activities and resources and mitigate the risk of opportunistic behaviour between cooperation partners (Raue & Wieland, 2015; Moazzeni & Sgarbossa, 2026).

Furthermore, asset sharing and the transition to autonomous transport also has implications for business models of the actors involved, including OEMs, transport companies, shippers and automation providers (Fritschy & Spinler, 2019). It can be expected that roles within the ecosystem, as well as established ways of value creation, delivery, and capture, will evolve as these transitions unfold. New roles may emerge, for example in orchestrating shared assets or coordinating autonomous fleets, while entirely new actors may enter the ecosystem. This is already observable in the context of CCAM, where automation providers increasingly act as intermediaries between OEMs and logistics companies, reshaping traditional industry boundaries. Understanding how these actors behave toward one another and how coordination and incentives are organized will be critical for enabling a successful transition toward autonomous, PI-enabled logistics systems.

Table 1 Key research topics for integrating Physical Internet and Autonomous Transport

Research topic	**Research opportunities**
Standardization and digitization of logistics processes	• Develop open standards for data messaging to autonomous vehicles • Create digital connection protocols between autonomous vehicles and logistics nodes • Using PI modular loading units for uniform loading and unloading of autonomous trucks
Innovating planning methods for decentralization and optimization	• Develop multi-agent algorithms • Develop decentralized planning platforms and data spaces • Investigate strategies for asset sharing planning methods
Asset sharing and implications for business models	• Governance models of shared autonomous truck fleets in PI networks. • Sharing costs and benefits of shared autonomous fleets in PI networks. • Implications for business models: Emerging roles and value chain configurations in the ecosystem.

4 Conclusions and considerations

4.1 Conclusions

This paper brings together two innovation pathways in the transportation sector – Physical Internet (PI) and Cooperative, Connected and Automated Mobility (CCAM) – into a preliminary joint research agenda. Three key overlapping research themes are identified: i) the standardization and digitization of logistics processes, ii) innovative planning methods for decentralization and optimization, and iii) asset sharing and its implications for business models. These themes represent areas with strong potential synergies between the PI and CCAM roadmaps, within which relevant research challenges and opportunities have been identified. While this agenda does not aim to be exhaustive, it provides an initial framework for aligning research efforts across both domains. For the PI and CCAM communities, this paper offers a first step toward closer research alignment, supporting stronger, coordinated innovation efforts to achieve shared objectives and ambitions.

4.2 Considerations

For further research, several considerations should be taken into account. The analysis of previous applied research projects and the PI and CCAM roadmaps was validated with a limited number of experts beyond the author team. This validation provides an important first confirmation of the relevance of the identified themes and supports the value of the work presented in this paper. At the same time, the limited scope of this validation means that additional themes may yet be identified where PI and CCAM research could and should further align to create additional synergies for improving transport systems. The three themes discussed here should therefore be seen as an initial but meaningful step toward a shared research agenda. As such, this paper is intended as a discussion starter, and the IPIC community is warmly invited to think along, reflect on the presented directions, and contribute to the continued development of this joint agenda.

5 Acknowledgements

This paper is co-funded by EU project MODI, and an indirect result of the already finished project CAT4Yards (co-funded by TKI-HTSM).

References

- ALICE-ETP. (2020). Roadmap to the Physical Internet. Brussels: ALICE. Retrieved from https://www.etp-logistics.eu/wp-content/uploads/2022/11/Roadmap-to-Physical-Intenet-Executive-Version_Final-web.pdf
- ERTRAC. (2022). Connected, Cooperative and Automated Mobility Roadmap. Brussels: ERTRAC. Retrieved from https://www.ertrac.org/wp-content/uploads/2022/07/ERTRAC-CCAM-Roadmap-V10.pdf
- ERTRAC. (2024). Connected, Cooperative and Automated Mobility Roadmap. Update of Chapter 2 "Agenda 2030" on Innovation Use Cases. Brussels: ERTRAC. Retrieved 04 2024, from https://www.ertrac.org/wp-content/uploads/2023/12/ERTRAC-CCAM-Roadmap-Chapter-2-Update-2024.pdf
- Fransen, R., Sprenger, L., Lee, W.-J., & van Meijeren, J. (2023). Performance of Self-Organizing Logistics: A Practical Comparison between Centralized and Decentralized Logistics. IPIC.

- Fritschy, C., & Spinler, S. (2019). The impact of autonomous trucks on business models in the automotive and logistics industry–a Delphi-based scenario study. Technological Forecasting and Social Change, 148, 119736.
- Hopman, M., Fransen, R., van Meijeren, J., & Zubin, I. (2022). A Framework on Centralised to Decentralised Logistics Control Structures Applied in Two Case Studies. Computational Logistics.
- Mishchenko, D., Iryna, S., Babinin, D., Vdovichen, A., & Verstiak, O. (2024). Standardisation of data in logitics and business operations. Multidisciplinary Reviews.
- Moazzeni, S., & Sgarbossa, F. (2026). Enabling collaborative logistics in rural contexts with digital technologies: a systematic literature review. International Journal of Production Research, 64(7), 2805-2828. doi:https://doi.org/10.1080/00207543.2025.2610764
- Montreuil, B., Meller, R. D., & Ballot, E. (2012). Physical internet foundations. IFAC Proceedings Volumes, 45(6), 26-30.
- Perman, A., Dovletgeldi, S., Begench, A., Bayramgeldi, G., & Tahyr, A. (2025). Digital Transformation in Logistics: Key drivers, challenges and sustainable solutions for the modern supply chain. North American Academic Research, 8(4), 152-160. doi:10.5281/zenodo.15241435
- Raue, J., & Wieland, A. (2015). The interplay of different types of governance in horizontal cooperations: A view on logistics service providers. The International Journal of Logistics Management, 26(2), 401–423. doi:https://doi.org/10.1108/IJLM-08-2012-0083
- Stein, S., & Prandtstetter, M. (2023). Can the Physical Internet Pave the Way to a Mobility of Entities. IPIC.
- van Kempen, E., Onverwagt, H., & Santing, M. (2026). Connected Automated Transport at Yards; Requirements and impacts for logistics. Den Haag: TNO.
- van Meijeren, J., Fransen, R., van Leeuwen, C., & Sprenger, L. (2025). (De)central planning approaches for asset sharing between multiple connected fleets for airside baggage transport. IPIC.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Place your logo here

Dynamic Pricing and Bayesian Demand Learning for Loading Dock Management

Hamrah Kor1[1,2], Lele Zhang[1,2], Mark Fackrell[1,2], Russel G. Thompson[3]
1 School of Mathematics and Statistics, The University of Melbourne, Parkville 3010, Australia.
2 ARC Training Centre in Optimisation Technologies, Integrated Methodologies, and Applications (OPTIMA), Australia.
3 Department of Infrastructure Engineering, The University of Melbourne, Parkville 3010, Australia
Corresponding author: h.kor@student.unimelb.edu.au

Abstract:

In the context of the Physical Internet (PI), loading docks function as shared, open-access nodes where fixed-capacity constraints and high variability in booking requests often lead to congestion and suboptimal resource utilisation. This paper proposes a mathematical model for a loading dock booking system where the facility manager dynamically sets prices for time slots to maximise revenue. A key feature of the model is its integration of Bayesian learning: the system updates its beliefs about unknown customer price sensitivities based on observed booking choices across available time slots. Because customers always select one of the available options, the system infers demand parameters from discrete-choice observations even without direct refusals. We formulate the problem as a finite-horizon Markov decision process with a mixed discrete-continuous state space representing beliefs about demand parameters and solve it using a hybrid epsilon-greedy policy that transitions from exploration to exploitation as confidence accumulates. Numerical experiments calibrated on real booking data from a Sydney loading dock facility demonstrate that the learning agent identifies the true demand parameters for all time slots within 20-29 days. A sensitivity analysis across booking volume and confidence thresholds confirms this result is robust.

Keywords: Dynamic pricing, Bayesian demand learning, loading dock management, Physical Internet (PI), PI nodes, Markov decision process, discrete choice model, urban logistics, Logistics-as-a-Service.

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),☐ Transportation Equipment, ☐ PI Networks, ☐ System of Logistics Networks, ☐ Vertical Supply Consolidation, ☐ Horizontal Supply Chain Alignment, ☒ Logistics/Commercial Data Platform, ☒ Access and Adoption, ☐ Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐ Flash Video, ☐ In-Person presentation*

1. Introduction

Loading dock management is a fundamental part of urban logistics, requiring the precise coordination of vehicle flows at key hubs such as retail centres, urban precincts, and shipping ports. However, these facilities have become critical bottlenecks, with a severe temporal mismatch: demand is highly congested during peak morning periods, while capacity remains

underutilised during afternoon hours (Sadeghi, 2024). This imbalance leads to site congestion and potential traffic issues at the entrance, while also generating noise and emissions that conflict with dense urban environments and strict regulatory standards (Alho et al., 2022; Dalla Chiara & Cheah, 2017). Furthermore, these docks must navigate land-use conflicts in settings where space is increasingly scarce. This environment creates a complex revenue management problem: how can facility managers implement dynamic pricing for time slots to maximise revenue, while smoothing temporal utilisation and learning customer preferences under unknown demand patterns?

Traditional approaches to loading dock management rely on fixed-fee schedules, first-come, first-served allocation, or simple rule-based priority systems (McLeod & Cherrett, 2011; Mor et al., 2020). These approaches ignore customer price responsiveness and forego potential revenue from peak-hour slots where demand substantially exceeds supply. Dynamic pricing, adjusting prices in real time based on observed demand, has transformed revenue management in airlines, hotels, and car-sharing platforms, but its application to urban logistics remains limited. The central difficulty is that demand parameters governing customer price sensitivity are unknown a priori and must be learned from observed choices (Friesz et al., 2008; Lei & Ouyang, 2017).

This paper addresses that challenge by modelling the facility manager as a Bayesian decision-maker who maintains and updates beliefs about price sensitivity across time slots as bookings arrive. The update rule exploits a structural feature of the environment: customers always select one of the available slots rather than leaving without booking, so each observed choice carries information not only about the chosen slot but also about all not-chosen alternatives, enabling efficient learning from relatively few daily observations. We formulate this as a finite-horizon Markov decision process (MDP) in which the state captures current capacity utilisation, arriving request characteristics, and the manager's belief distribution over demand parameters. Prices follow a hybrid epsilon-greedy policy that balances exploration and exploitation.

The model is calibrated using real booking data from a loading dock facility in Sydney, Australia. We define the static as a fully informed benchmark agent that knows the true demand parameters and sets prices optimally from day one, and cumulative revenue difference as the total revenue loss of the learning agent relative to the static over the horizon. Numerical experiments demonstrate rapid convergence: the learning agent identifies the true demand parameters for all time slots within 20 to 29 days and outearns the static benchmark cumulatively within the 30-day horizon. A sensitivity analysis across booking volume (daily request rate) and confidence thresholds (the posterior certainty required before switching from exploration to exploitation) confirms that both convergence speed and post-convergence revenue recovery are robust across low to high daily booking volumes and a range of manager-controlled parameter settings.

The contributions of this paper are threefold. First, we provide a rigorous MDP formulation of the loading dock pricing problem that accounts for capacity constraints, demand uncertainty, and multi-slot Bayesian inference from forced-choice observations. Second, we develop and analyse a hybrid epsilon-greedy learning policy with adaptive per-slot confidence thresholds calibrated to actual booking reservations. Third, we present a comprehensive numerical study using real facility data, including a sensitivity analysis that characterises how model parameters affect convergence speed, cumulative regret, and post-convergence revenue performance.

The remainder of the paper is organised as follows. Section 2 reviews related work in dynamic pricing. Section 3 gives the MDP model including the problem description, MDP components, and pricing policies. Section 4 presents numerical results. Section 5 concludes the paper with some ideas for future research.

2. Literature review

The challenge of dynamic pricing when customer price sensitivity is unknown and must be inferred from observed choices has been studied from several angles in operations research. Araman and Caldentey (2009) showed that a monopolist facing demand uncertainty must navigate a path between exploiting current beliefs and probing the market for better information. Harrison et al. (2012) formalised this trade-off in a Bayesian framework, treating pricing decisions explicitly as signals of demand rather than merely as revenue levers. Den Boer and Zwart (2014) showed that greedy policies can achieve near-optimal regret in parametric demand models, while Russo and Van Roy (2014) offered Thompson sampling as a tractable alternative to full Bayesian optimisation. Chen and Mersereau (2015) confirmed that the value of learning is highest when the prior is diffuse, and the horizon is long, both of which characterise our setting.

Our model differs from most of this literature in what the manager observes. Rather than sales volumes, we see only which slot a customer chose from an available menu, placing our work in the multinomial logit (MNL) choice tradition (Gallego & Wang, 2014). Agrawal et al. (2019) studied online learning for MNL-based assortment problems and derived regret bounds for Thompson sampling, but their setting retains an outside option. In our model, every customer books something, which means each observation carries richer information: choosing slot t is simultaneously a positive signal for t and a negative signal for every other slot on the menu.

On the parking and loading dock side, Legros and Fransoo (2024) formulated an MDP for an on-street lot shared between passenger vehicles and delivery trucks, combining dynamic pricing with admission control under finite-capacity constraints. This is structurally the closest antecedent to our model, though demand parameters are treated as known throughout. For loading docks specifically, McLeod and Cherrett (2011) showed, through simulation, that advance booking reduces double parking and idle bay time in an urban freight context. Mor et al. (2020) reinforced this from the distributor side, embedding booking into a routing model and demonstrating gains in a Lisbon case study. None of these papers treat demand parameters as unknown and learn them from booking observations. That is the gap this paper fills in. Within the PI framework, loading docks function as critical shared nodes where goods are consolidated, transferred, and dispatched across hyperconnected logistics networks (Montreuil et al., 2013). Ballot (2019) argued that PI nodes must evolve beyond static infrastructure toward intelligent, price-responsive service points capable of dynamically allocating capacity in response to real-time demand signals. Our model directly operationalises this vision by equipping a PI node operator with a Bayesian learning mechanism that adapts prices as demand patterns are revealed through observed bookings.

3. Model

We formulate the loading dock pricing problem as a finite-horizon MDP in which a facility manager dynamically prices M identical loading bays across $T = 24$ discrete hourly time slots $\mathcal{T} = \{1,2, \dots, T\}$ spanning an operating day. Prices are uniform across bays within each time slot, and customers book a time slot rather than a specific bay; a customer requiring consecutive slots is treated as occupying a contiguous block. While the mathematical model supports multi-slot requests, numerical experiments in this paper focus on single-slot bookings to establish a baseline for the learning mechanism. The objective is to maximise revenue over a horizon of N sequential booking arrivals within each day, repeated across multiple operating days as the manager progressively learns the unknown demand parameters from observed booking choices.

3.1 State, decision variables, and exogenous information

3.1.1 State variables

The system state is observed at each decision epoch n, where each epoch corresponds to the arrival of one booking request. We begin with the components of the state variable. Let $R_{n,t} \in \{0,1/M,2/M,\dots,1\}$ denote the fraction of bays reserved in slot t after processing request n, where M is the total number of bays per slot. The system starts empty, so $R_{0,t} = 0$ for all $t \in \mathcal{T}$, and the reserved capacity is non-decreasing across requests (cancellations are not permitted). The vector $\boldsymbol{R}_n = (R_{n,1}, \dots, R_{n,T})$ collects the reservation fractions across all slots. A slot is available for booking only if $R_{n,t} < 1$, and we write the set of available slots as

$$\mathcal{A}_n = \{t \in \mathcal{T} \mid R_{n,t} < 1\}. \tag{1}$$

Let $Q_n \in \mathcal{T}$ denote the requested start time of the n-th arriving customer, and $L_n \in \{1, \dots, L_{\max}\}$ the requested duration in consecutive slots, where $L_{\max}$ is the maximum permissible booking duration. The pair $\boldsymbol{I}_n = (Q_n, L_n)$ characterises the n-th request. This information is exogenous and stochastic, reflecting customer requirements beyond the manager's control. While the model supports multi-slot bookings, the numerical experiments in Section 4 restrict attention to single-slot bookings (i.e., $L_{\max} = 1$).

Let $\boldsymbol{\theta}_t$ denote the unknown demand parameter vector for slot t. Let $\Theta = \{\boldsymbol{\theta}^1, \dots, \boldsymbol{\theta}^K\}$ be a discrete set of K candidate parameter vectors. For each slot t, let $P_{n,t}^k = \mathbb{P}(\boldsymbol{\theta}_t = \boldsymbol{\theta}^k \mid \mathcal{H}_n)$ denote the posterior probability that the true parameter equals candidate k, where $\mathcal{H}_n$ is the history of all booking observations up to request n. The belief distribution over candidates for slot t is the vector $\boldsymbol{B}_{n,t} = (P_{n,t}^1, \dots, P_{n,t}^K) \in \Delta^{K-1}$, where $\Delta^{K-1}$is the $(K-1)$-dimensional probability simplex. The full belief state is $\boldsymbol{B}_n = (\boldsymbol{B}_{n,1}, \dots, \boldsymbol{B}_{n,T})$.

The complete system state at epoch n is given by the tuple $S_n = (\boldsymbol{R}_n, \boldsymbol{I}_n, \boldsymbol{B}_n)$ with state space $\mathcal{S} = \mathcal{R} \times \mathcal{I} \times \mathcal{B}$, where $\mathcal{R} = \{0,1/M, \dots,1\}^T$, $\mathcal{I} = \mathcal{T} \times \{1, \dots, L_{\max}\}$, and $\mathcal{B} = (\Delta^{K-1})^T$.

3.1.2 Decision Variables

Let $[x_{\min}, x_{\max}]$ be the feasible price range. At each epoch n, the manager observes S_n and posts a price vector $\boldsymbol{x}_n = (x_{n,1}, \dots, x_{n,T}) \in [x_{\min}, x_{\max}]^T$, where fully reserved slots are assigned a price of zero by convention: $x_{n,t} = 0$ whenever $R_{n,t} = 1$. A policy π is characterised by a decision function X^π that maps each state to a price vector, so that $\boldsymbol{x}_n = X^\pi(S_n)$.

3.1.3 Exogenous Information

The next arriving customer, request $n+1$, observes the posted prices $\boldsymbol{x}_n$ and selects a slot $D_{n+1} \in \mathcal{A}_n$. By accepting slot D_{n+1} at price $x_{n,D_{n+1}}$, the customer implicitly rejects the posted price of every other available slot. We assume all customers book one of the available slots rather than leaving without a booking, so D_{n+1} always corresponds to a confirmed reservation. The exogenous information arriving at epoch $n+1$ is therefore $W_{n+1} = (I_{n+1}, D_{n+1})$, and the revenue collected from request $n+1$ is

$$C(S_n, \boldsymbol{x}_n, W_{n+1}) = x_{n,D_{n+1}}. \tag{2}$$

The objective is to maximise the total expected revenue over the finite horizon,

$$\max_{\pi} \mathbb{E}\left[\sum_{n=0}^{N-1} x_{n,D_{n+1}}\right]. \tag{3}$$

3.2 Demand and Choice Model

We model the booking probability for a time slot t at posted price $x \geq 0$ using a logistic demand function. Given a candidate parameter vector $\boldsymbol{\theta}^k = (\theta_1^k, \theta_2^k)$, the demand function is

$$d_{\boldsymbol{\theta}^k}(x) = \frac{1}{\theta_1^k + \exp\left(\theta_2^k \cdot x\right)}, \tag{4}$$

where $\theta_1^k > 0$ controls the baseline booking probability and $\theta_2^k > 0$ determines customers' price sensitivity. This functional form is utilised as it provides a flexible parameterisation where θ_1 represents the inherent attractiveness of a time slot independently of the price-sensitivity (θ_2), facilitating stable parameter identification. While logit models usually reflect individual couriers, we focus on aggregate time-slot demand. This simplifies the learning task for the facility manager, as identifying individual couriers' utilities is significantly more challenging and variable than learning aggregate demand parameters. This formula ensures $d_{\boldsymbol{\theta}^k}(x) \in (0,1)$ for all $x \geq 0$ and is strictly decreasing in price, which is consistent with standard demand theory. If the true parameter vector $\boldsymbol{\theta}_t$ were known, the probability that a customer chooses slot $t \in \mathcal{A}_n$ under prices $\boldsymbol{x}_n$ would be given by a multinomial logit model in which every customer selects one of the available slots and the option of leaving without booking is assumed unavailable, that is

$$\mathbb{P}(D_{n+1} = t \mid \boldsymbol{\theta_t}, \boldsymbol{x}_n, \mathcal{A}_n) = \frac{d_{\boldsymbol{\theta}_t}(x_{n,t})}{\sum_{j \in \mathcal{A}_n} d_{\boldsymbol{\theta}_j}(x_{n,j})}. \tag{5}$$

Since the true parameters $\boldsymbol{\theta}_t$ are unknown, the MDP works with the expected demand for slot t at price x under the current belief $\boldsymbol{B}_{n,t}$, that is

$$\bar{u}_{n,t}(x) = \sum_{k=1}^{K} P_{n,t}^k \cdot d_{\boldsymbol{\theta}^k}(x) = \sum_{k=1}^{K} P_{n,t}^k \cdot \frac{1}{\theta_1^k + \exp\left(\theta_2^k \cdot x\right)}, \tag{6}$$

which serves as the current demand estimate for pricing decisions.

3.3 State Transitions and Bayesian Updating

After observing the customer's choice $D_{n+1} = t^*$, the resource state and the belief state are updated simultaneously. The resource state evolves as

$$R_{n+1,t} = \begin{cases} R_{n,t} + \dfrac{1}{M} & \text{if } t \in \{D_{n+1}, \dots, D_{n+1} + L_{n+1} - 1\}, \\ R_{n,t} & \text{otherwise,} \end{cases} \tag{7}$$

showing that the booked slot and all $L_{n+1} - 1$ consecutive slots required by the customer's requested duration have their reserved capacity increased by one bay. Under the single-slot restriction adopted in the numerical experiments ($L_{\max} = 1$), equation (7) reduces to incrementing only slot D_{n+1}.

Belief updating exploits the information embedded in the complete choice set. Observing that a customer chose t^* from $\mathcal{A}_n$ provides a positive signal for t^* and a negative signal for every other available slot, because each observed choice is simultaneously an acceptance of the selected slot and an implicit rejection of all alternatives. To compute the likelihoods, we first write the expected utility of slot j under the current belief at the posted price as $\bar{u}_{n,j} = \bar{u}_{n,j}(x_{n,j})$. The likelihood that slot t^* is chosen under the hypothesis $\boldsymbol{\theta}_{t^*} = \boldsymbol{\theta}^k$ is

$$\mathcal{L}(D_{n+1} = t^* \mid \boldsymbol{\theta}_{t^*} = \boldsymbol{\theta}^k, \boldsymbol{x}_n, \mathcal{A}_n) = \frac{d_{\boldsymbol{\theta}^k}(x_{n,t^*})}{d_{\boldsymbol{\theta}^k}(x_{n,t^*}) + \sum_{\substack{j \in \mathcal{A}_n \\ j \neq t^*}} \bar{u}_{n,j}}, \tag{8}$$

while for a non-chosen slot $t \in \mathcal{A}_n, t \neq t^*$, the likelihood that slot t is not selected under the hypothesis $\boldsymbol{\theta}_t = \boldsymbol{\theta}^k$ is

$$\mathcal{L}(D_{n+1} = t^* \mid \boldsymbol{\theta}_t = \boldsymbol{\theta}^k, \boldsymbol{x}_n, \mathcal{A}_n) = 1 - \frac{d_{\boldsymbol{\theta}^k}(x_{n,t})}{d_{\boldsymbol{\theta}^k}(x_{n,t}) + \sum_{\substack{j \in \mathcal{A}_n \\ j \neq t}} \bar{u}_{n,j}}. \tag{9}$$

Applying Bayes' rule, the posterior belief for every slot $t \in \mathcal{A}_n$ is updated as

$$P^k_{n+1,t} = \frac{\mathcal{L}(D_{n+1} \mid \boldsymbol{\theta}_t = \boldsymbol{\theta}^k, \boldsymbol{x}_n, \mathcal{A}_n) \cdot P^k_{n,t}}{\sum_{j=1}^{K} \mathcal{L}\left(D_{n+1} \mid \boldsymbol{\theta}_t = \boldsymbol{\theta}^j, \boldsymbol{x}_n, \mathcal{A}_n\right) \cdot P^j_{n,t}}, \tag{10}$$

while beliefs for fully booked slots remain unchanged. Together, the resource and belief transitions constitute the transition function S^M, and the next state $S_{n+1} = S^M(S_n, \boldsymbol{x}_n, W_{n+1})$.

3.4 Pricing Policies

We consider two policies that differ in how they trade off exploitation and exploration.

Pure exploitation

At each decision epoch n, the exploitation policy maximises expected revenue by solving, for each slot $t \in \mathcal{A}_n$ independently, the single-slot revenue problem $\max_{x \geq 0} \bar{C}_{n,t}(x)$, where

$$\bar{C}_{n,t}(x) = x \cdot \bar{u}_{n,t}(x) = x \cdot \sum_{k=1}^{K} P^k_{n,t} \cdot d_{\boldsymbol{\theta}^k}(x). \tag{11}$$

The optimal price x_t^* satisfies the first-order condition $d\bar{C}_t/dx|_{x_t^*} = 0$, which we evaluate numerically using bounded scalar optimisation (Brent, 1973).

We consider a *static benchmark* that has perfect knowledge of the *true* slot-specific demand parameters $\boldsymbol{\theta}_t$ and computes a fixed optimal price, that is,

$$x_t^{\text{static}} = argmax_{x \geq 0}\{x \cdot d_{\boldsymbol{\theta}_t}(x)\}. \tag{12}$$

The price remains constant across all epochs since the static benchmark has the perfect knowledge in advance and requires no learning. The corresponding revenue serves as an informed benchmark for evaluating the learning agent's performance.

Hybrid epsilon-greedy policy

The hybrid epsilon-greedy policy operates in two phases, transitioning from exploration to permanent exploitation on a per-slot basis as posterior confidence accumulates.

Phase 1, Exploration. Let $[x_{\min}, x_{\max}]$ denote the feasible price range and let $\varepsilon_n = \max(\varepsilon_{\min}, \varepsilon_{\text{start}} \cdot \gamma^n)$ denote the exploration rate at epoch n, where $\varepsilon_{\text{start}} \in (0,1]$ is the initial exploration rate, $\gamma \in (0,1)$ is a decay factor, and $\varepsilon_{\min} > 0$ is a floor that prevents exploration from vanishing entirely. For each unlocked slot t, defined as any slot for which $\max_k P^k_{n,t} < \tau_t$, the policy selects a uniformly random price from $[x_{\min}, x_{\max}]$ with probability ε_n and the exploitation price $x^*_{n,t}$ with probability $1 - \varepsilon_n$, that is,

$$x_{n,t} = \begin{cases} \mathcal{U}(x_{\min}, x_{\max}) & \text{with probability } \varepsilon_n, \\ x_t^* & \text{with probability } 1 - \varepsilon_n. \end{cases} \tag{13}$$

This phase ensures broad price experimentation in early epochs, when beliefs are diffuse and the value of information is high, and gradually shifts toward exploitation as beliefs concentrate.

Phase 2, Permanent exploitation. Once the posterior confidence for slot t, defined as $\max_k P^k_{n,t}$, exceeds the slot-specific threshold τ_t, that slot permanently transitions to the exploitation phase:

it is *locked* at the joint revenue-maximising price $x^*_{n,t}$ and never explores again, regardless of the global ε_n. Full system convergence is declared when all T slots have been locked.

4. Numerical analysis

4.1 Data and Calibration

The model is calibrated using 47,421 booking records from MobileDock at a facility in Barangaroo, Sydney, from January to July 2023. Each day is divided into $T = 24$ hourly time slots, yielding an average daily booking of $N = 310$ requests. Demand is highly skewed: the 10:00 slot accounts for 15.9% of daily bookings while the least-demanded slot accounts for less than 0.3%. The facility operates $M = 50$ bays per slot. The feasible price range is set to [\$0, \$20]. The exploration decay factor is set to $\gamma = 0.994$ with $\varepsilon_{start} = 0.60$ and $\varepsilon_{min} = 0.001$, and the per-slot confidence threshold to $\tau_t = 0.99$ for all t. To assign ground-truth demand parameters, we rank all 24 slots by their historical booking share and map each rank to one of $K = 10$ candidate parameter vectors $\boldsymbol{\theta}^1, \ldots, \boldsymbol{\theta}^{10}$, where $\boldsymbol{\theta}^1$ corresponds to the lowest observed demand intensity and $\boldsymbol{\theta}^{10}$ to the highest. Slots with similar booking shares are assigned the same candidate vector. The initial belief for every slot is set to the uniform distribution over all K candidates, that is, $P^k_{0,t} = 1/K$ for all t and k, reflecting no prior knowledge of demand parameters. Initial prices $x_{0,t}$ are drawn uniformly from $[x_{\min}, x_{\max}]$. The convergence of the posterior belief toward the true $\boldsymbol{\theta}_t$ for each slot is verified in Section 4.2.

4.2 Results

Figure 1 presents the simulation results across the 30-day horizon. Each day begins where the previous day ended: the belief state accumulated from all prior booking observations is carried forward, so the system builds on its learning continuously across the 30-day period.

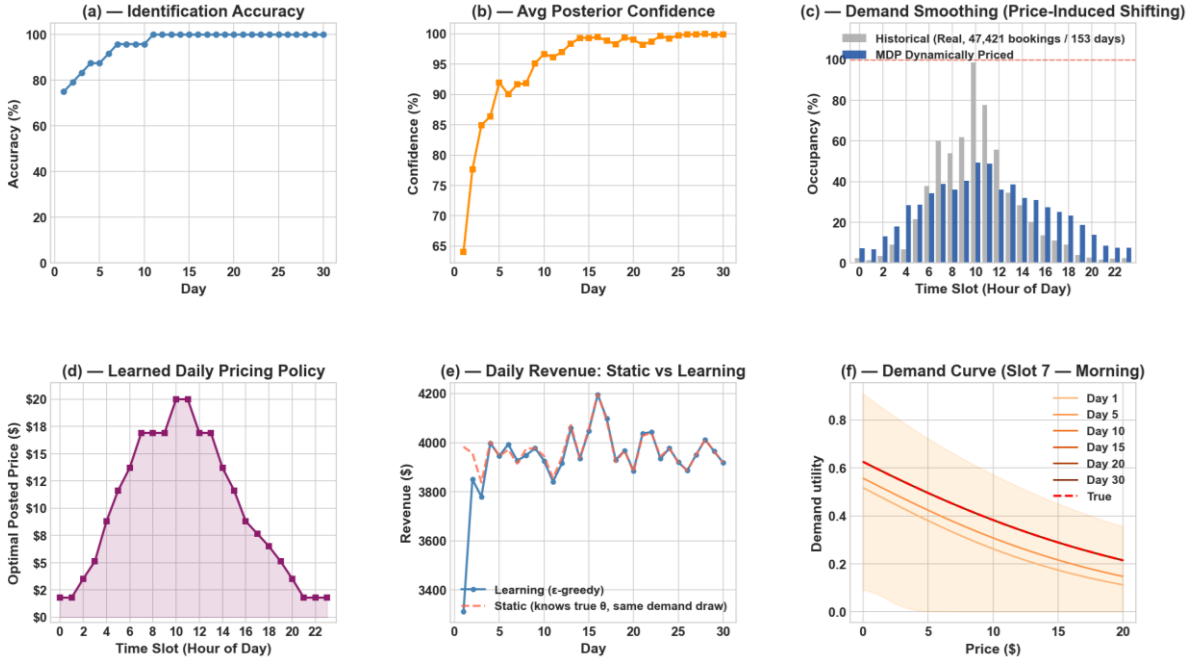
Identification accuracy, defined as the fraction of slots for which the posterior-modal candidate matches the true demand parameter θ_t, as shown in Figure 1(a), climbs from approximately 75% on day 1 to 100% by day 11, with low-traffic slots being the last to resolve due to their weak likelihood signals. Average posterior confidence, defined as the mean over all slots of $\max_k P^k_{n,t}$, the highest posterior probability assigned to any single candidate, (Figure 1(b)) follows a similar pattern, rising from 64% on day 1 to 99% by day 25. Figure 1(c) illustrates the booking demand over a day before and after dynamic pricing is implemented. Under historical patterns in the absence of dynamic pricing, demand concentrates heavily in slots 7–12, with slot 10 near full capacity while afternoon slots sit at 10-20% utilisation. The hybrid pricing policy substantially flattens this profile by pricing peak slots high enough to redirect excess demand into adjacent hours.

Figure 1(d) shows the final prices. The arch shape reflects the demand structure: slots 10-11 hit the \$20 cap, slot 7 settles around \$17, and prices drop toward \$2 for low-demand slots. The high price at slot 10 redistributes excess demand away from the congested peak toward underutilised afternoon slots, producing a more evenly distributed utilisation profile across the day. This flattening is desirable as it reduces congestion and idle capacity simultaneously, improving overall facility efficiency.

Figure 1(e) shows the daily revenue comparison. As expected in a standard revenue management setting, the static benchmark initially outearns the learning agent because it possesses perfect information. The learning agent incurs 'regret' during the early exploration phase (Days 1–10). However, as identification accuracy reaches 100% (Day 11), the learning

agent's revenue converges to the static benchmark, achieving nearly identical performance for the remainder of the horizon.
Figure 1(f) shows demand curve learning for slot 7, where the solid line represents the posterior mean estimated demand curve and the shaded band represents the 95% credible interval across candidate parameter vectors. The estimated curve aligns closely with the ground truth by day 15, though the uncertainty band narrows gradually rather than collapsing sharply, consistent with this slot's moderate observation rate.

Figure 1: Numerical results



4.3 Sensitivity Analysis

We conduct three sensitivity tests on booking volume (SA1), confidence threshold (SA2), and seed robustness (SA3) and report the sensitivity test results in Table 1. Tests SA1 and SA2 are evaluated over 5 seed replications and 30 days, while SA3 considers 20 seeds. Table 1 reports four metrics across scenarios: *Days to Exploit* is the number of days until all T slots are locked (that is, full system convergence); *Bookings to Exploit* is the total number of booking arrivals accumulated up to that point; *Cumulative Revenue Difference* is the total revenue of the static benchmark minus hybrid policy that of the static benchmark summed over all 30 days; and *Revenue Gap vs Static* is the post-convergence percentage revenue difference between the hybrid and static policies, measured after all slots have locked.

Booking volume has the most significant effect on learning speed. When there are $N = 150$ requests per day, the hybrid policy requires nearly the full 30-day window to learn the demand parameters $\boldsymbol{\theta}_t$, as low-traffic slots rarely accumulate enough observations to exceed threshold $\tau_t = 0.99$. Convergence improves steadily with volume, falling from 24.8 days at the baseline ($N = 310$) to under 10 days at $N = 1{,}000$, though the post-convergence revenue gap widens to -1.95% at the lowest volume, reflecting premature locking before beliefs have fully concentrated. The cumulative revenue difference is positive across all scenarios, meaning the static benchmark outearns cumulatively the learning policy.

The confidence threshold sensitivity (SA2) reveals a precision-speed trade-off. Lower τ_t accelerates convergence but increases the post-convergence revenue gap, while the baseline $\tau_t = 0.99$ achieves the tightest gap of 0.08% at the cost of a longer learning period. The seed

robustness results (SA3) confirm that these findings are stable across 20 replications, with moderate standard deviations reflecting genuine demand randomness rather than sensitivity to initialisation.

Table 1: Sensitivity analysis, key results (5 seeds × 30 days per scenario)

Scenario	Parameter	Days to Exploit (mean±std)	Bookings to Exploit (mean±std)	Cum. Rev. Diff. (Static-Hybrid, $, mean±std)	Rev. Gap vs Static (%)
SA1 - Booking volume	N = 150	29.4 ± 1.2	4,410 ± 180	1,114 ± 1,124	-1.95%
	N = 310 (baseline)	24.8 ± 4.5	7,688 ± 1406	881 ± 1,337	-0.08%
	N = 500	18.2 ± 6.3	9,100 ± 3153	915 ± 1,549	0.08%
	N = 800	12.2 ± 2.8	9,760 ± 2229	944 ± 1,115	0.00%
	N = 1,000	9.8 ± 3.4	9,800 ± 3429	1,486 ± 1,284	-0.00%
SA2 - Confidence threshold	τ_t = 0.90	18.2 ± 6.9	5,642 ± 2142	658 ± 918	0.21%
	τ_t = 0.95	18.2 ± 6.1	5,642 ± 1885	939 ± 1,114	-0.13%
	τ_t = 0.99 (baseline)	24.8 ± 4.5	7,688 ± 1406	881 ± 1,337	-0.08%
SA3 - Seed robustness	20 seeds	23.9 ± 5.1	7,394 ± 1585	921 ± 956	-0.28%

5. Conclusion

This paper developed and analysed a Bayesian dynamic pricing model for loading dock management, calibrated on real facility data from a Sydney operator. The facility manager maintains a belief distribution over unknown demand parameters $\boldsymbol{\theta}_t$ for each time slot and updates these beliefs through a finite-horizon MDP as booking requests arrive. To balance learning and revenue generation, we propose a hybrid epsilon-greedy policy with per-slot adaptive confidence thresholds τ_t, which transitions from exploration to permanent exploitation on a slot-by-slot basis as posterior confidence accumulates. The exploration phase generates price variation that produces informative likelihood signals, enabling the posterior to concentrate progressively around the true $\boldsymbol{\theta}_t$ for all slots, with full parameter identification achieved within 20 to 29 days across the baseline setting period.

A notable feature of the hybrid epsilon-greedy policy is that random exploration prices are drawn uniformly from $[x_{\min}, x_{\max}]$, which on average exceed the theoretical optimum for low-demand slots but approximate it more closely for high-demand slots already priced near the cap. For high-demand slots, this means exploration operates near the revenue-maximising price, generating strong likelihood signals at low cost. For low-demand slots, exploration occasionally sets prices far above the optimum, producing informative observations but at a temporary revenue cost. This asymmetry is self-correcting: high-demand slots accumulate sufficient posterior mass quickly and lock early, while low-demand slots take longer to converge but contribute little to total revenue in the interim. The net effect is that the policy concentrates its learning cost where it matters least and converges fastest where it matters most.

The sensitivity analysis confirms that the main findings are robust across booking volumes of 150 to 1,000 per day and confidence thresholds of 0.90 to 0.99. Booking volume is the most important operational parameter: facilities with low demand ($N = 150$) may require the full 30-day window to converge, while high-volume facilities ($N = 800$) achieve full convergence within two weeks.

Future work will extend the model in several key directions. We will incorporate an outside or ‘no-booking’ option to reflect the possibility of customers rejecting all offered prices, which would introduce a critical trade-off between revenue maximization and customer retention. Further extensions include handling multi-slot bookings and alternative pricing policies with

stronger theoretical guarantees. A practically motivated direction is fixed pricing for low-traffic slots, where sparse observations make Bayesian learning ineffective within a realistic horizon.

References

- Agrawal S., V. Avadhanula, V. Goyal, A. Zeevi (2019): *MNL-Bandit: A Dynamic Learning Approach to Assortment Selection.* Operations Research, v67, no5, 1453–1485.
- Alho A., S. Oh, R. Seshadri, G. Dalla Chiara, W.H. Chong, T. Sakai, L. Cheah, M. Ben-Akiva (2022): *An Agent-Based Simulation Assessment of Freight Parking Demand Management Strategies for Large Urban Freight Generators*. Research in Transportation Business & Management, v43, 100804.
- Araman V.F., R. Caldentey (2009): *Dynamic Pricing for Nonperishable Products with Demand Learning*. Operations Research, v57, no5, 1169-1188.
- Ballot E. (2019): The Physical Internet (pp. 719-734). Springer International Publishing.
- Brent, R. P. (1973). Algorithms for minimization without derivatives. Prentice-Hall.
- Chen L., A. Mersereau (2015): Analytics for Operational Visibility in the Retail Store: The Cases of Censored Demand and Inventory Record Inaccuracy. Retail Supply Chain Management, Springer, Boston, MA.
- Dalla Chiara, G., & Cheah, L. (2017). *Data stories from urban loading bays*. European Transport Research Review, 9, 1-16.
- Den Boer A.V., B. Zwart (2014): *Simultaneously Learning and Optimizing Using Controlled Variance Pricing.* Management Science, v60, no3, 770-783.
- Friesz, T.L., Mookherjee, R., Holguín-Veras, J., & Rigdon, M.A. (2008). *Dynamic pricing in an urban freight environment*. Transportation Research Part B-methodological*, 42*, 305-324.
- Gallego G., R. Wang (2014): *Multiproduct Price Optimization and Competition Under the Nested Logit Model with Product-Differentiated Price Sensitivities*. Operations Research, v62, no2, 450-461.
- Harrison J.M., N.B. Keskin, A. Zeevi (2012): *Bayesian Dynamic Pricing Policies: Learning and Earning Under a Binary Prior Distribution*. Management Science, v58, no3, 570-586.
- Legros B., J.C. Fransoo (2024): *Admission and Pricing Optimization of On-Street Parking with Delivery Bays*. European Journal of Operational Research, v312, no1, 138–149.
- Lei C., Y. Ouyang (2017): *Dynamic Pricing and Reservation for Intelligent Urban Parking Management.* Transportation Research Part C: Emerging Technologies, v77, 226-244.
- McLeod F., T. Cherrett (2011): *Loading Bay Booking and Control for Urban Freight. International* Journal of Logistics Research and Applications, v14, no6, 385–397.
- Montreuil B., R.D. Meller, E. Ballot (2013): *Physical Internet foundations*. Studies in Computational Intelligence, 472.
- Mor A., M.G. Speranza, J.M. Viegas (2020): *Efficient Loading and Unloading Operations via a Booking System.* Transportation Research Part E: Logistics and Transportation Review, v141.
- Powell, W. (2022). Sequential Decision Analytics and Modeling: Modeling with Python. Emerald Publishing Limited.
- Russo D., B. Van Roy (2014): *Learning to Optimize via Posterior Sampling*. Mathematics of Operations Research, v39, no4, 1221–1243.
- Sadeghi, H. (2024). Enhancing Loading Dock Booking Systems. Master of Philosophy thesis, The University of Melbourne. http://hdl.handle.net/11343/355542

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE

Station Simulation for Hyperloop Systems

Lukas Eschment[1], Irina Yatskiv (Jackiva)[2], Jurijs Tolujevs[2], Thomas Schüning[1,3] and Walter Neu[1,3]

1. Institute of Hyperloop Technology, University of Applied Sciences Emden/Leer, Emden, Germany
2. TSI, Transport and Telecommunication Institute, Riga, Latvia
3. School of Mathematics and Science, Carl von Ossietzky University of Oldenburg, Oldenburg, Germany

Corresponding author: lukas.eschment@hs-emden-leer.de

Abstract: *Hyperloop is increasingly discussed as a high-speed, low-carbon complement to rail and road that can absorb part of the expected growth in passenger and freight demand while supporting decarbonisation and resilience objectives. Its success, however, depends on how closely technology and operations are matched for all subsystems, this includes stations which operate as high-throughput Physical Internet nodes embedded in logistics networks, rather than only focussing on tube and vehicle technologies alone. This paper develops a discrete-event simulation model of a hyperloop station that translates high-level concepts and stakeholder requirements into quantitative station-level performance indicators. The AnyLogic-based model captures interactions between station layout, convoy operations, variable demand and disruptions, using synthetic pod-arrival scenarios and detailed dwell-time formulations for different airlock and airdock technologies. Key outputs include throughput, queue lengths, station occupancy, convoy formation times and resilience metrics, which are mapped to a multi-level framework which is employing a TOPSIS decision-making approach. Preliminary results reveal non-linear relationships between station sizing and capacity, identifying convoy formation and boarding technology as critical secondary bottlenecks, and highlighting the importance of capacity reserves and operational control strategies for maintaining robust PI-compatible hyperloop station operations.*

Keywords: *Physical Internet, Hyperloop, Station Design, Simulation, Demand, Capacity, Disruption*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☒ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☒ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

1 Introduction

European freight and passenger transport demand is projected to grow substantially over the coming decades, while climate targets require large reductions in greenhouse gas emissions and improved resilience to disruptions. High-speed, electric and automated modes such as hyperloop have been proposed as part of a broader green and digital transformation of the transport system [European Commission, 2020; International Transport Forum, 2019]. Hyperloop combines near-vacuum tubes, magnetically levitated pods and high-speed

operations to offer aircraft-like travel times with energy efficiency similar or lower than rail [Hyperloop Development Program, 2025; Hyper4Rail Consortium, 2024; Duin et al., 2025].

In the Physical Internet (PI) vision, transport networks are organised as open, modular and standardised logistics systems, in which flows of goods and passengers are routed through nodes and links using common protocols [Montreuil, 2011; Sarraj et al., 2014]. To participate in such a system, hyperloop must be embedded into existing and future PI-compatible networks via stations that act as high-speed hubs and interfaces to other modes. These stations must handle large and fluctuating volumes while maintaining strict operational, safety and resilience constraints also considering vacuum operations and automated boarding technologies.

Existing hyperloop research has focused strongly on vehicle and guideway design, energy consumption and corridor-level demand, whereas detailed modelling of station operations under stochastic demand and disruption remains scarce [Rajendran and Srinivas, 2020; Wu and Margreiter, 2023; Hyper4Rail Consortium, 2024]. At the same time, PI-oriented decision frameworks require quantitative key performance indicators (KPIs) that capture how node-level technologies and layouts influence network performance and stakeholder objectives. This paper addresses this gap by answering the following research questions:

1. How do hyperloop-specific station technologies and layouts affect throughput, dwell time and resilience under stochastic demand and disruption?
2. Which combinations of station design parameters and convoy operations provide robust performance for PI-compatible hyperloop nodes?
3. How can station-level simulation be integrated into a multi-level decision framework to inform technology selection and concept-of-operations decisions?

The study builds on a multi-level decision framework for hyperloop implementation developed in earlier work (Eschment L. et al., 2025) and expands it through a detailed station simulation model that generates PI-relevant KPIs for the node and network level. This paper reviews related work on hyperloop and PI station modelling, introduces the simulation methodology, setup and parameterization and discusses implications for technology and design decisions and PI-aligned hyperloop implementation to derive KPIs for stakeholder focussed decision-making [Eschment et al., 2025; McGinnis and Huang, 2021].

2 State of the art and related work

2.1 Physical Internet and hyperloop nodes

The PI conceptualises freight and passenger transport as a set of interoperable, modular and open logistics systems, in which goods are encapsulated into standardised units and routed through shared networks of PI hubs and corridors [Montreuil, 2011; Ballot et al., 2014; Sarraj et al., 2014]. Within this paradigm, nodes such as terminals, stations and warehouses become critical infrastructure elements whose design must support high throughput, automation, interoperability and resilience across operators and modes [McGinnis and Huang, 2021; Puskás et al., 2021].

Hyperloop fits naturally into PI thinking because it offers high-speed, fully electric line-haul services that can connect major metropolitan areas and logistics clusters, while interfacing with existing rail, road and air networks at PI-compatible hubs. Concepts such as the CargoTube freight hyperloop and hybrid passenger–freight deployments rely on standardised containers that can be transferred between hyperloop and other modes via automated handling systems at

stations. Station performance of directly affects PI-level KPIs including corridor capacity, travel and lead times, and the utilisation of high-value vehicles and infrastructure through its network focussed transport system approach.

2.2 Simulation of high-throughput terminals and hyperloop stations

Discrete-event and agent-based simulation are widely used in transport planning to analyse complex terminal and hub operations before physical implementation. Studies of airport terminals, baggage handling systems and urban interchanges in Riga and other cities have shown how microscopic simulations can quantify queue formation, operations, evacuation times and service levels under different layouts and operational policies, providing decision support for safety and capacity planning [Yatskiv et al., 2016; Savrasovs, 2022].
In the hyperloop domain, recent work has applied pedestrian micro-simulation to passenger flows in conceptual stations as well as mesoscopic models to analyse high-speed network operations. However, existing hyperloop station models typically focus on either passenger flow or analytic throughput approximations and often assume deterministic or average demand and process times. They rarely incorporate explicit models and parameters of emerging hyperloop-specific technologies, such as airlock or airdock boarding systems, nor do they link station indicators systematically to multi-level decision frameworks or PI KPIs [Rajendran and Srinivas, 2020; Wu and Margreiter, 2023].
The present work complements these studies by developing a process-level discrete-event simulation of station operations that explicitly models key hyperloop technologies, integrates stochastic demand and disruption, and outputs KPIs that can be used by a multi-level decision framework. It deliberately incorporates three technological levels: the network, node (stations and hubs) and physical handling layer.

3 Methodology

3.1 Hyperloop station system and PI-specific complexities

Hyperloop stations differ from conventional rail or airport terminals in several respects that materially affect their operational behaviour and motivate the need for dedicated simulation models. First, the near-vacuum guideway requires specialised interface technologies—such as pumping airlocks, opening airlocks or localised airdock systems—that maintain pressure integrity while enabling rapid boarding and alighting [Koot and Fu, 2024; Hyper4Rail Consortium, 2024]. These devices introduce additional processing steps and constraints beyond those in atmospheric terminals, including pressurisation cycles, gate-valve sequencing and safety interlocks.

Second, hyperloop operations are expected to rely on frequent, small vehicles that can be grouped into convoys on the main line to achieve high capacity while preserving operational flexibility. Convoy formation before departure and dispersal within or near stations create complex interactions between arrival patterns, docking positions and departure slots, which differ from conventional fixed-train formations. Third, conceptual business-case studies foresee peak design demands of above 20 000 passengers per hour in the busiest European stations, corresponding to around 500 pods per hour with 40 passengers each [Hyper4Rail Consortium, 2024]. Such intensities, combined with stringent safety and pressure constraints, require precise coordination of flows at PI nodes. Hyperloop specific technologies posing complexities for station design are outlined in Figure 1.

Finally, hyperloop stations must operate as intermodal PI hubs that connect to high-speed rail, local public transport and airports, while ensuring short and predictable transfer times. This places additional requirements on land-use efficiency, accessibility and robustness under disruptions, which are reflected in the KPI set of the multi-level decision framework.

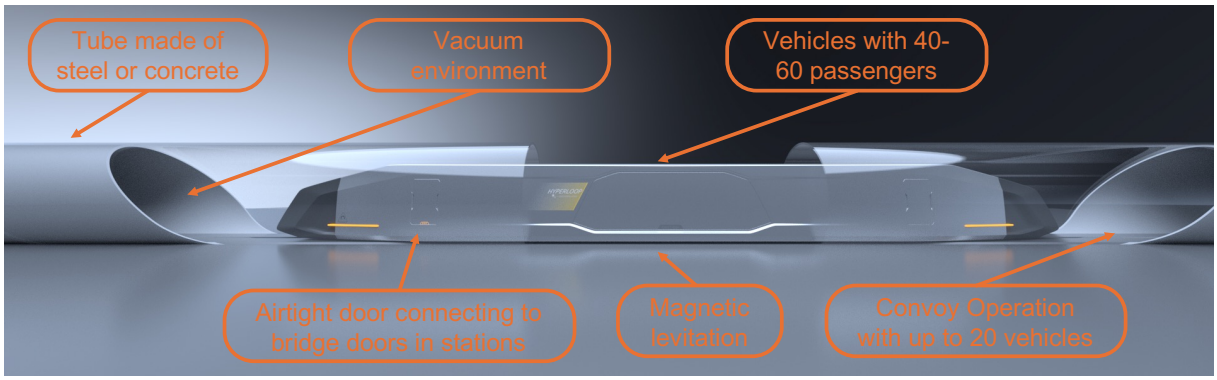


Figure 1: Distinct features of hyperloop influencing operations and station design.

3.2 Development of simulation tools for station operations

To investigate the movement, processing, and queuing dynamics of capsules within station environments, two universal simulation models were developed. These models are distinguished by their logic regarding capsule access to processing points. **Figure 2** illustrates the model structure featuring **sequential access**, while **Figure 3** depicts the structure for **parallel access** to the processing points (station cells). These structural diagrams represent the animation frames generated by the operational models. Capsules shown in a light shade represent units in a waiting state for which the processing stage has been completed.

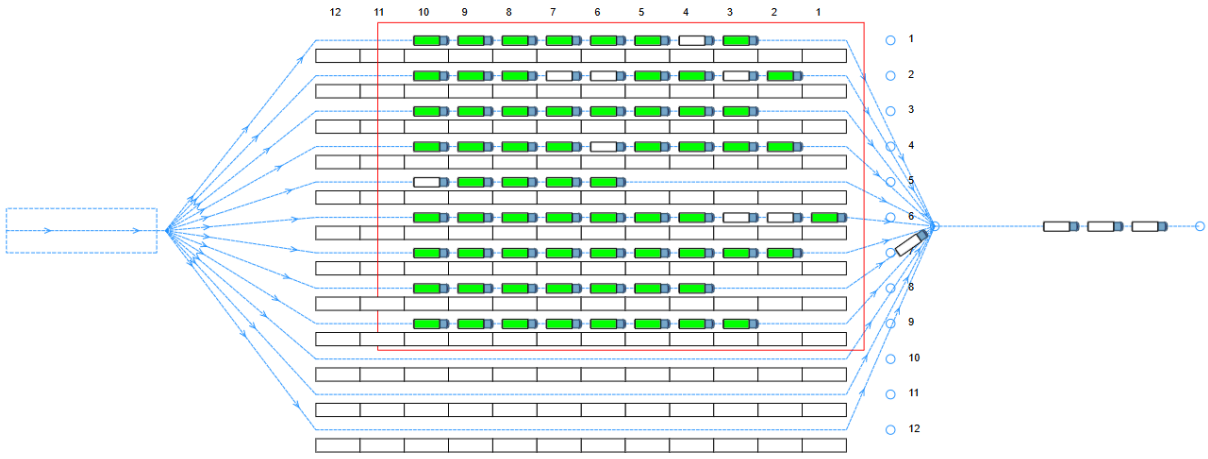


Figure 2: Placement of capsules in a sequential access model.

In the sequential access model in Figure 2, the arrival of a new capsule initiates a procedure to identify a line with available capacity at its distal end. The capsule occupies a position either at the head of the line or immediately behind a capsule currently undergoing processing. In the latter case, a capsule is restricted from exiting the line until all preceding units have departed. Assuming capsules arrive at a specific line at a given interval τ_a and the discrete processing time is τ_s, logical analysis yields two distinct operational scenarios:

Case 1: $\tau_a > \tau_s$. A new capsule arrives at the line only after the primary cell (Cell 1) has been vacated. Consequently, secondary cells (e.g., Cell 2) are never utilized, rendering the total line length irrelevant to the system's throughput.

Case 2: $\tau_a < \tau_s$. A new capsule occupies a position behind a unit that is still being processed. Although preceding capsules eventually begin to exit, incoming units continue to populate the line until its maximum capacity is reached. The final capsule in the sequence can only depart once all preceding positions are cleared, after which the entire line becomes vacant and can be reused. This cycle creates a self-oscillating process (self-excitation), where the line entrance is periodically blocked before the entire resource is released.

If the values τ_a are τ_s are stochastic variables, various operational scenarios may arise within the line, examples of which are illustrated in Figure 2. Under these conditions, two distinct negative effects are observed to varying degrees: (1) Underutilization of resources: A significant portion of the primary cells (those with lower index numbers) remains vacant. (2) Exit blocking: Certain capsules that have completed their processing stage remain in a waiting state, as they are physically unable to depart the line due to preceding units.

Figure 2 demonstrates an example of a station configuration consisting of 9 lines, each equipped with 10 cells, providing a total capacity of 90 processing units for capsule handling.

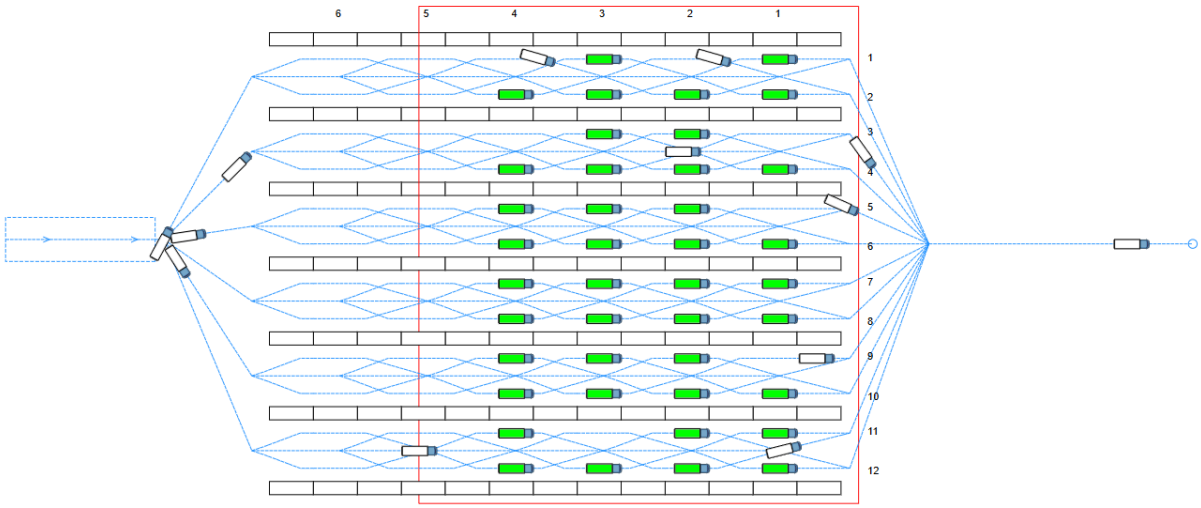


Figure 3: Placement of capsules in a parallel access model.

In the **parallel access** model, in Figure 3, an arriving capsule can immediately occupy any available cell and vacate it promptly upon the conclusion of the processing stage. Figure 3 illustrates an example featuring 12 lines with 4 cells each, providing a total capacity of 48 processing units. On the left side of the diagram, four capsules are shown entering the station area, as an available cell has been identified for each.

These models are designated as universal due to their extensive parameter sets, which allow for the description of diverse operational modes and station scenarios. The station's configuration is defined by the number of lines and the number of cells per line. The maximum dimensionality for each model is determined by the graphical structures shown in Figure 2 and Figure 3. For the sequential model in Figure 2, the maximum dimensionality is 12x12, whereas for the parallel model in Figure 3, it is 12x6.

Both models were developed using the Process Modeling Library of the AnyLogic simulation software. The input and output capsule flows are simulated using the Discrete Event paradigm, while the movement and processing of capsules within the station space are implemented via multi-agent modeling based on complex state charts.

3.3 Simulation setup and parameterization

A simulation scenario is defined as a comprehensive description of the properties and model parameters configured for the simulation of processes at a specific station. Table 1 presents different parameters and technical nuances & logic for future scenario formulation.

By emphasizing "transit times" as the proxy for physical distance, Table 1 clarifies that while the visual model is a template, the mathematical model is spatially accurate. In aspect of capsule proceeding technologies, it preserves the distinction between "upstream" (airlocks) and "at-bay" (airdocks), which is a critical engineering difference in hyperloop station design.

Table 1: Parameters,technical nuances and logic for future scenario formulation

Parameters	Description of Technical Nuances and Logic
Station Geometry	The fundamental operational logic must align with the schemes in **Figure 2** or **Figure 3.** Geometric constraints are represented via **calculated transit times** between any coordinate pair in the logical scheme.
Capsule Processing Technologies	Based on **pumping airlocks**, **opening airlocks**, and **localized airdock** concepts. Pumping and opening airlocks are typically located upstream of the station to handle entire pods, while airdocks serve as localized pressure interfaces at each bay, functioning similarly to airport jet bridges but engineered for near-vacuum environments. Process times for airlock and airdock operations are calibrated from conceptual engineering studies of various boarding technologies and decomposed into phases: pressure equalization, door operations, passenger or cargo exchange, and re-pressurization. Processing time is an aggregate of phases of the selected technology defined via constants or stochastic distributions.
Capsule Arrival Flows	Supports three modes: **(1) Schedule-based** (deterministic timestamps); **(2) Stochastic** (inter-arrival distributions); and **(3) Demand-responsive** (closed-loop system) for peak capacity testing.
Convoy Arrival and Departure	Arrival Logic managed via user-defined schedules using Convoy IDs and synchronized arrival timestamps and includes: • Sequential Access: Distributed randomly across lines or queued for a single line until full capacity is available. • Parallel Access: Supports simultaneous occupation of cell groups, limited only by total station capacity. Departure Logic supports preservation of original inbound convoy structures. Enables formation of new convoys based on deterministic or stochastic sizing.
Internal Station Failures	Hardware outages for specific cells or lines requires a model extension: (1) **Failure/Recovery Schedule:** A table defining the downtime and restoration intervals for cells or lines. (2) **Status Matrix:** A variable matrix reflecting the real-time operational status (Active/Failed) of each cell or line.

Capsule Arrival Flows supports 3 modes. Let’s consider it in detail:

(1) Schedule-based arrivals: Each capsule enters the station according to a user-defined schedule (manual table or auxiliary program). Each entry includes arrival time, capsule type, cell processing time, and other user-defined parameters. This method allows for the modeling of two failure types: (a) capsule delays (via adjusted arrival times) and (b) cell dwell-time delays (via extended processing times).

(2) **Stochastic arrivals:** Theoretical or empirical probability distributions are used to model inter-arrival times and processing durations. This approach is suitable when the statistical properties of the flow are time-invariant. The objective is to assess process stability when flow parameters deviate from nominal values within narrow stochastic limits.
(3) **Demand-responsive (Closed system):** The input flow depends on the availability of cells. A new capsule is generated only when a free cell is identified. In queuing theory, such systems are classified as closed systems. Since all available cells remain occupied at any given time, this is referred to as saturated mode. This modeling objective is to determine the station's maximum throughput capacity.

Convoy modeling is best implemented via user-defined schedules. Simultaneous arrival times are assigned to the respective capsules, with a specific parameter identifying the convoy ID. Processing methods for sequential access: (a) capsules are distributed randomly among parallel lines, or (b) the entire convoy is directed to a single line once sufficient capacity becomes available and for parallel access: both options remain valid; however, simultaneous occupation of a cell group is only possible if the convoy size does not exceed the station's total cell count. For departures, the model supports: (a) maintaining or re-forming the original inbound convoy, (b) forming a new convoy with a deterministic or stochastic size.

While simple processing delays are handled via input flow parameters, simulating the failure of specific cells or entire lines requires a model extension: (1) Failure/Recovery Schedule and (2) Status Matrix. The failure and recovery cycle can be explicitly defined by the user or generated using reliability-based distribution laws.

4 Results

4.1 Station sizing, capacity and land-use implications

Simulation experiments confirm that station throughput scales with the number of lanes and docking bays, though a pronounced saturation effect occurs beyond critical input thresholds.

- Capacity Benchmark: For a demand of 20,000 passengers/hour, a configuration of 9 lanes and 10 bays is sufficient to maintain acceptable queue lengths and dwell times, provided nominal performance is met.
- Infrastructure Trade-offs: Adding lanes beyond this point yields marginal capacity gains while significantly increasing infrastructure complexity and station footprint.
- PI-Level Performance: Land-use KPIs—specifically footprint per unit of peak throughput—indicate that compact layouts with high utilization are more effective for PI integration than simply increasing station scale.

4.2 Convoy formation and handling technologies

The interaction between convoy operations and docking technologies is a primary determinant of system performance.

- The Convoy Bottleneck: While larger convoys increase mainline capacity, they impose restrictive synchronization requirements at platforms. Beyond a specific length, mainline gains are offset by increased station dwell times, particularly under high demand variability.
- Handling Technology Comparison:

    - Pumping Airlocks: Provide robust pressure control but suffer from cyclical delays (pump-down/venting) and become critical at high vehicle frequencies.
    - Opening Airlocks: Reduce cycle times by venting residual air into the tube, though at the cost of higher mechanical stress on pumps and seals.
    - Localized Airdocks: Enable the shortest cycles and smallest footprints. However, they introduce a high volume of safety-critical moving parts and exhibit sharper performance degradation during failures [Koot and Fu, 2024].
- Resilience Strategy: Analytical data suggests airdocks are superior for footprint and speed, but require more robust redundancy and maintenance strategies to mitigate failure impacts.

4.3 Disruptions, resilience and recovery behavior

The station transitions from a **quasi-steady state** to a **congested regime** based on stochastic variability and disruption intensity.

- Delay Propagation: Under peak loads, disruptions such as airlock failures trigger cascading delays that persist long after resolution. This is exacerbated when stations lack capacity reserves.
- Strategic Correlations: Certain control strategies, such as maximizing convoy size to improve mainline throughput, directly increase station footprint requirements and reduce shock-absorption capacity.
- PI-Oriented Frameworks: These findings underscore the need for a shift in hyperloop evaluation. Rather than focusing on average throughput, decision frameworks must prioritize resilience KPIs:
    - Maximum Queue Length (Spatial impact)
    - Time to Recovery (TTR) (Temporal impact)
    - Fraction of Demand Served on Time (SLA fulfillment)

5 Discussion

5.1 Contribution to Physical Internet node design

By explicitly modeling hyperloop-specific station technologies and convoy operations under stochastic demand and disruption, this simulation provides a quantitative foundation for designing nodes that function as seamless interfaces within PI networks. The derived KPIs (throughput, dwell-time distributions, resilience indicators, and land-use metrics) align with the multi-level decision frameworks required by diverse stakeholders.

- For technology developers: The results facilitate a comparative analysis of handling technologies regarding capacity, robustness, and feasibility.
- For operators: The model provides evidence-based strategies for convoy management and buffer sizing to optimize the balance between utilization and reliability.
- For policymakers: These metrics assess how hyperloop integration supports broader PI objectives, specifically regarding decarbonization, infrastructure resilience, and efficient land-use at multimodal hubs.

5.2 Comparative advantages over existing approaches

Unlike conventional transport models, this research introduces three distinct methodological advances:

- Domain-Specific Complexity: It addresses hyperloop-unique constraints such as near-vacuum docking interfaces and high-frequency convoy synchronization, rather than extrapolating from traditional rail or aviation models.

- Stochastic Resilience Modeling: By incorporating stochastic demand and disruption, the model captures non-linear cascading effects and "congested regimes" that deterministic approaches fail to identify.
- Integrated Decision Framework: Simulation outputs are fed into a TOPSIS framework, ensuring that microscopic design decisions remain synchronized with network-level PI-compatible planning and multi-stakeholder policy objectives.

5.3 Limitations and Future Research

While robust, the current model employs several abstractions:

- Agent Dynamics: Passenger behavior is modeled via aggregate distributions; future iterations will integrate explicit agent-based pedestrian dynamics to assess safety and accessibility.
- Logistics Granularity: Current focus remains on pod-level throughput; detailed cargo-handling sequences and customer-specific logistics will be addressed in subsequent work.
- Economic Modeling: CAPEX and OPEX are currently assessed via proxies (e.g., footprint). Future research will couple this simulation with network-scale tools to evaluate complete lifecycle emissions and revenue streams.
- Empirical Validation: As emerging hyperloop demonstrators provide data, the model will be calibrated to serve as a standardized, PI-oriented evaluation tool for operational reliability and evacuation procedures.

6 Conclusions

Simulation experiments highlight that resilience to disruptions depends strongly on capacity reserves and operational control strategies. The simulation experiments underscore that convoy formation and boarding technologies represent the critical bottlenecks in station design. Furthermore, system resilience is shown to be highly dependent on capacity reserves and dynamic operational control. By integrating these insights into a multi-level TOPSIS framework, stakeholders can evaluate station concepts based on a holistic set of PI objectives—including modularity, high asset utilization, and network-wide synchronicity—rather than nominal capacity alone.

This study demonstrates that microscopic station simulation is an essential prerequisite for network-level hyperloop analyses, providing the operational evidence necessary to design PI-ready, high-speed nodes within future global transport systems.

References

- Awasthi A., S.S. Chauhan, H. Omrani (2011): Application of fuzzy TOPSIS in evaluating sustainable transportation systems, Expert Systems with Applications, v38, no10, 12270–12280, https://doi.org/10.1016/j.eswa.2011.04.005.
- Ballot E., B. Montreuil, R.D. Meller (2014): The Physical Internet: The Network of Logistics Networks.
- Duin, H. et al.: Supporting the Planning of Dedicated CargoTube Links Through Simulation. In: McNally, C., Carroll, P., Martinez-Pastor, B., Ghosh, B., Efthymiou, M., Valantasis-Kanellos, N. (eds) Transport Transitions: Advancing Sustainable and Inclusive Mobility. TRA-conference 2024. Lecture Notes in Mobility. Springer, Cham. (2025). DOI: 10.1007/978-3-031-89444-2_47.

- European Commission (2020): Sustainable and Smart Mobility Strategy – Putting European transport on track for the future, European Commission, COM(2020) 789 final, https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52020DC0789.
- Eschment L. et al. (2025): Resilient and sustainable transport networks: a novel decision-making framework for optimized logistics solutions, Proceedings of IPIC 2025.
- Hyperloop Development Program (2025): The European Hyperloop Development Program: A Vision for the Future of Transport, Delft, https://europe-diplomatic.eu/science/the-european-hyperloop-development-program-a-vision-for-the-future-of-transport/.
- Hyper4Rail Consortium (2024): Deliverable 2.3: Concept of Operations, Hyper4Rail Consortium.
- International Transport Forum (2019): Transport System Resilience: Summary and Conclusions, OECD/ITF, Paris, https://www.itf-oecd.org/sites/default/files/docs/transport-system-resilience.pdf.
- Koot J., Y.F. Fu (2024): Feasibility Study of a Bridge Door Airlock Design at Hyperloop Stations, Delft Hyperloop – EHW full-scale report.
- Macharis C., A. Bernardini (2015): Review of multi-criteria analysis for transport project evaluation: principles and applications, Transport Policy, v37, 177–191, https://doi.org/10.1016/j.tranpol.2014.11.002
- Mattsson L.-G., E. Jenelius (2015): Vulnerability and resilience of transport systems – A discussion of recent research, Transportation Research Part A: Policy and Practice, v81, 16–34, https://doi.org/10.1016/j.tra.2015.06.002.
- McGinnis L.F., K. Huang (2021): Design and simulation of a Physical Internet logistics hub, Proceedings of IPIC 2021.
- Montreuil B. (2011): Towards a Physical Internet: Meeting the Global Logistics Sustainability Grand Challenge, Logistics Research, v3, no2–3, 71–87, https://doi.org/10.1007/s12159-011-0045-x.
- Puskás, N., Abunnasr, Y., Naalbandian, S. (2021): Assessing deeper levels of participation in nature-based solutions in urban landscapes – A literature review of real-world cases, Landscape and Urban Planning, v.210, 104065, https://doi.org/10.1016/j.landurbplan.2021.104065.
- Rajendran S., S. Srinivas (2020): A simulation-based approach to provide insights on the operations of a Hyperloop network, Transportation Research Part A: Policy and Practice, v132, 115–139, https://doi.org/10.1016/j.trip.2020.100092
- Savrasovs M. (2022): Simulation as a decision support tool for airport planning: Riga International Airport case study, Transport, v36, no6, 474–485, https://doi.org/10.3846/transport.2021.16198.
- Sarraj R., E. Ballot, S. Pan, B. Montreuil (2014): The Physical Internet: review, analysis and future research agenda, International Journal of Production Research, v52, no15–16, 4055–4080, https://doi.org/10.1108/IJPDLM-12-2016-0353
- Wu H.-T., M. Margreiter (2023): Pedestrian Flow Simulation and Capacity Analysis of Stations for Ultra-High-Speed Ground Transportation (Hyperloop), 8th International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS 2023), 1–6, https://doi.org/10.1109/MT-ITS56129.2023.10241765.
- Yatskiv I., M. Savrasovs, V. Gromule, V. Zemljanikins (2016): Passenger Terminal Safety: Simulation Modelling as Decision Support Tool, Procedia Engineering, v134, 459–468, https://doi.org/10.1016/j.proeng.2016.01.068.
- Zavadskas E.K., A. Mardani, Z. Turskis, A. Jusoh, K.M. Nor (2016): Development of TOPSIS Method to Solve Complicated Decision-Making Problems—An Overview on Developments from 2000 to 2015, International Journal of Information Technology & Decision Making, v15, no3, 645–682, https://doi.org/10.1142/S0219622016300019.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE
Physical Internet Center

Designing Scalable Multi-Tier Networks for Physical Internet (PI)-Enabled Hyperconnected Urban Logistics

Praveen Muthukrishnan[1], Walid Klibi[2] and Benoit Montreuil [1,3]

1. Physical Internet Center, Supply Chain and Logistics Institute, H. Milton Stewart School of Industrial & Systems Engineering, Georgia Institute of Technology, Atlanta, USA
2. Kedge Business School, Bordeaux, France
3. Coca-Cola Material Handling & Distribution Chair

Corresponding author: pmuthukr3@gatech.edu

Abstract: *Urban freight systems require scalable methods for identifying suitable hub structures across multiple spatial tiers, yet fully integrated multi-tier design models become difficult to solve when candidate sets are large and demand is highly heterogeneous. This paper proposes an iterative, multi-tier, coverage-based screening for pre-designing PI-enabled urban freight networks. The framework combines tier-wise coverage screening, modular hub capability and capacity representation, KPI-based evaluation, and iterative top-down and bottom-up refinement across area, local and proximity tiers. The methodology is demonstrated on Metropolitan Atlanta, comprising 1,248 proximity zones, 100 local cells, and 10 urban areas, with candidate pools exceeding 166,000* proximity *hubs, 66,000 local hubs, and 2,000 area hubs. Results show that the framework achieves several orders of magnitude reduction in candidate space while maintaining cross-tier structural consistency, providing a tractable basis for subsequent detailed network design.*

Keywords: *Physical Internet, multi-tier freight networks, hyperconnected urban logistics, scalable logistics networks, coverage-based screening; modular hub capacity*

Physical Internet (PI) Roadmap Fitness*: ☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories), ☒ PI Networks, ☒ System of Logistics Networks*

Targeted Delivery Mode-s*: ☒ Paper, ☒ In-Person presentation*

1 Introduction

Urban freight systems face growing pressure to become efficient, flexible and sustainable. The Physical Internet (PI) paradigm offers a promising direction for redesigning metropolitan systems as open, multi-tier, hyperconnected, and interoperable networks in which resources are shared across organizational boundaries (Montreuil, 2011). In practice, this vision can be operationalized through a three-tier hierarchical structure in which area, local and proximity hubs operate across three spatial scales (zones, cells and areas) to coordinate freight movements throughout a metropolitan region (Liu et al., 2025; Montreuil et al., 2018).

Translating this architecture into a deployable metropolitan network requires solving a tightly coupled set of problems, involving territory design, hub location and role assignment, service network design, and network flow optimization. These decisions exhibit strong cross-tier interdependencies, as hub role assignment at one tier conditions the feasible service coverage, inter-tier connections, and flow patterns at other tiers. This challenge is further compounded by the heterogeneity of candidate facilities, which differ substantially in accessibility, cost,

capacity potential and operational suitability. As a result, candidate space quickly becomes large and structurally complex to handle through direct optimization.

This motivates the need for a dedicated pre-design stage that reduces the candidate space before the remaining design problems are addressed in detail. The purpose of the pre-design is not to determine the final network, but to screen and retain a structurally coherent set of hub candidates that preserves the essential features needed for downstream design. These include tier-wise service requirements, inter-tier connectivity requirements, and indicative capacity needs. At this stage, the emphasis is placed on territorial structure and aggregate inter-hub flow assessment, while detailed routing and exact flow allocation are deferred to later stages once the design space has become tractable.

In response, this paper develops a coverage-based framework for the pre-design of PI-enabled multi-tier metropolitan freight networks. The contribution is threefold. First, the paper extends coverage-based candidate screening beyond predominantly single tier or partially integrated multi-tier formulations by introducing an iterative multi-tier logic, in which candidate evaluation is informed by top-down structuring and bottom-up feedback. Second, it moves beyond static or aggregate early-stage capacity treatment by incorporating a modular hub capacity representation aligned with PI principles of scalable and progressive deployment. Third, while candidate reduction is commonly addressed using a limited set of cost or coverage indicators, the framework develops a broader KPI-based assessment structure for evaluating retained candidate sets across efficiency, serviceability, equity, resiliency, and sustainability. The framework is illustrated on a case representative of a large metropolitan region, using Metropolitan Atlanta as the experimental testbed.

The remainder of the paper is organized as follows. Section 2 reviews the relevant literature on hierarchical facility location, multi-echelon network design, and modular capacity deployment. Section 3 defines the freight system and planning problem considered. Section 4 presents the proposed framework. Section 5 reports the computational study and discusses the main findings. Section 6 concludes with implications, limitations, and future research directions.

2 Literature Review

Table 1 Comparison of related studies across key design dimensions

Reference	Hierarchical / Multi Tier	Multi-Echelon Location Routing	Modular Capacity	Multi-Criteria Evaluation	Large Candidate Pool Screening
Farahani et al. (2014)	✓	Partial	-	-	-
Ortiz-Astorquiza et al. (2018)	✓	Partial	-	-	-
Janjevic et al. (2021)	✓	✓	-	-	-
Ben Mohamed et al. (2023)	✓	✓	-	-	-
Jena et al. (2015)	-	-	✓	-	-
Silva et al. (2015)	Partial	-	✓	-	-
Liu et al. (2026)	Partial	Partial	✓	Partial	-
Behbabani et al. (2019)	-	-	-	✓	-
Dastgoshade et al. (2022)	-	Partial	-	✓	-
Klibi and Martel (2012)	Partial	Partial	-	✓	-
Muthukrishnan et al. (2021)	-	-	✓	-	✓
This paper	✓	✓	✓	✓	✓

As summarized in Table 1, the proposed framework is positioned at the intersection of four related literature streams: *hierarchical facility location*, *multi-echelon freight network design*, *modular capacity planning*, and *multi-criteria network evaluation*. A first relevant stream is the literature on hierarchical and multi-level facility location. Existing studies provide important

building blocks by addressing hierarchical structure, multi-echelon coordination, modular deployment, or broader performance assessment. However, these dimensions are typically treated separately or only partially combined, with limited attention to metropolitan-scale candidate screening under strong cross-tier dependence. The figure shows that no single study brings together multi-tier screening, cross-tier structural consistency, modular hub capacity, and KPI-based multi-criteria evaluation in a tractable pre-design setting. This paper addresses that gap through a coverage-based framework that reduces the candidate space while preserving essential service, connectivity, and capacity structure required for downstream network design.

3 Problem Definition

3.1 Pre-Design Problem Scope

A logistics service provider (LSP), or a consortium of LSPs, deploying a PI-enabled multi-tier metropolitan freight network faces a large, heterogenous pool of hub candidates from which a reduced yet structurally coherent subset must be identified before detailed network design can proceed. Rather than determining the final metropolitan network in full operational detail, the aim is to preserve key structural properties such as territorial serviceability, feasible inter-tier connectivity, and a reasonable aggregate transportation and facility cost structure required for subsequent stages in which routing, flow allocation and refined operational decisions are addressed explicitly. The problem is accordingly formulated as a multi-criteria, structure preserving screening problem, preceding and informing full network design.

3.2 Multi-Tier Freight System Structure

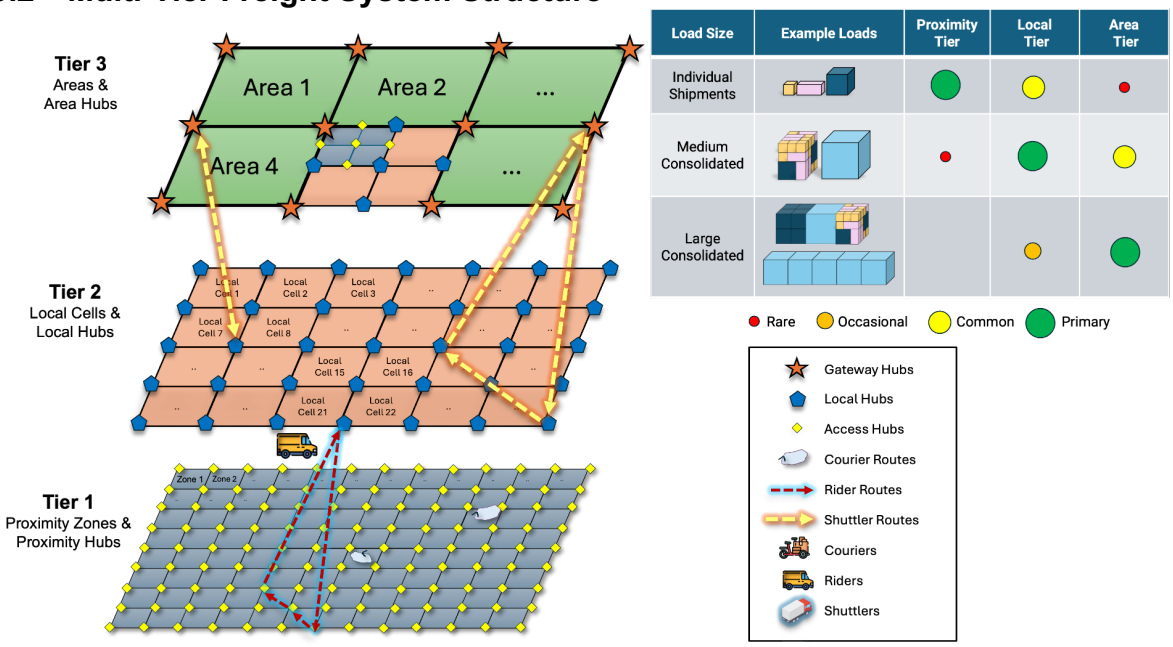


Figure 1: Multi-Tier Hyperconnected Urban Freight Network – Topology, Operational classes and shipment entry-tier relationships based on consolidation level

We consider a PI-enabled urban freight system organized as a three-tier network composed of area, local and proximity hubs. The metropolitan space is represented through a predefined spatial hierarchy in which demand locations are modeled as proximity zones, grouped into local cells, which are further grouped into areas as shown in Figure 1. This hierarchy is obtained upstream through clustering procedures designed to balance geographic contiguity, compactness, demand density and travel time accessibility. In the present study, this hierarchy

is treated as fixed input. Within this hierarchy, proximity hubs serve as transshipment and consolidation point to localized demand. Local hubs serve as intermediate consolidation, redistribution, and transfer facilities. Area hubs provide the interface between the metropolitan freight system and external regional or intercity freight networks (Garbers et al., 2025).

Together, these hub tiers support freight movement across spatial scales through coordinated consolidation, transfer and deconsolidation operations. While the hierarchy provides the organizing structure of the network, it does not impose a strict one-to-one correspondence between hub tier and hub function. Higher-tier hubs may, where appropriate, perform functions associated with lower tiers, and multiple tier roles may therefore be assigned to the same physical location. Similarly, shipment flows are not required to traverse all tiers sequentially. Depending on shipment characteristics, consolidation conditions, and service requirements, they may bypass an intermediate tier or connect directly to a higher-tier hub.

3.3 Flow Representation

Metropolitan freight demand is heterogeneous in activity type, shipment size, consolidation level, and movement direction. Flows may be generated by or attracted to residential, commercial, fulfillment, and industrial locations, and may correspond to intracity, intercity inbound, or intercity outbound movements. To represent this diversity within a common modeling structure, shipments are expressed as equivalent π-modular containers (Montreuil et al., 2014) Consistent with the Physical Internet containerization logic, π-containers are distinguished by function into packaging, handling, and transport containers, each of which may exist at multiple modular dimensions. In this pre-design model, small, medium, and large π-containers may act as handling or transport units depending on their consolidation level and the hub tier through which they interact with the metropolitan network.

Small π-containers capture parcel, tote, or low-volume modular flows, such as e-commerce parcels and small retail orders. Medium π-containers capture case, carton, grouped parcel, or moderately consolidated flows, such as store replenishment and pre-bundled urban shipments. Large π-containers capture highly consolidated modular freight units, such as fulfillment-center or industrial shipments. These categories are used as modular abstractions rather than direct equivalents of conventional parcels, pallets, or trailers. This representation is summarized in Figure 1, which shows how different shipment forms are associated with entry relationships across the area, local, and proximity tiers of the network.

3.4 Operational Layer

The entry tier of a flow is defined by its container function, modular scale, and degree of pre-consolidation. Small handling-oriented loads typically enter through proximity hubs, medium flows through local or area hubs, and larger consolidated flows primarily through area hubs. Figure 1 summarizes this representation by illustrating shipment forms and entry tier associations. Fully consolidated truckload shipments for a specific destination directly interface with higher-tier PI nodes and are treated as external to the pre-design scope.

Freight movement across the three-tier hierarchy is structured through tier-aligned route classes. Courier routes provide first-mile and last-mile connectivity between proximity zones and hubs, typically using light vehicles such as cargo bikes and small vans. Rider routes connect proximity hubs to local hubs and support intermediate consolidation across local cells, generally using larger urban vehicles such as vans or small trucks. Shuttler routes connect local hubs to area hubs and carry the most consolidated upper-tier flows, typically using medium- or heavy-duty trucks. These route classes define principal operational interfaces between successive tiers rather than explicit paths to be optimized at the pre-design.

3.5 Modular Capacity Representation

To represent scalable and progressive deployment, hub modularity is modeled along two complementary dimensions: *functional* and *capacity modularity.* Functional modularity specifies a set of logistics operations a hub can perform, drawn from a common library $\mathcal{F}$ defined for the network. This library encompasses operations such as handling different π-container sizes, consolidation and deconsolidation, sorting, cross-docking, buffering, and interfacing with tier-appropriate vehicle classes. Let $F_h \subseteq \mathcal{F}$ denote the set of functions activated at hub h, determined by its role. Similar functions appear across all tiers, what varies is their processing intensity, technological sophistication, and installed scale. For example, transfer at a proximity hub may involve low-volume parcel handling with simple staging and light-vehicle interfaces, whereas transfer at an area hub may involve much higher throughput, larger handling units, more advanced equipment, and tighter synchronization with upper-tier freight services. Capacity modularity defines the amount of such capability deployed, including space, equipment, staging area, labor, and automation support. capacity modules. This distinction matters because hub evolution may occur either through expansion of existing operations or through activation of additional operational roles.

For modeling purposes, installed capacity is represented through a tier-specific equivalent processing capacity rather than through each operational dimension separately, an abstraction that preserves the essential capacity structure while remaining tractable at the pre-design stage. Let N_h denote the number of deployed capacity modules at hub h, and K_h the effective processing capacity of one module at the corresponding tier. Capacity feasibility is enforced through: $W_h = \sum_{f \in F_h} w_{hf} \leq K_h \cdot N_h$, where w_{hf} denotes the workload contribution of function f at hub h, and W_h denote the total workload across all activated functions. Functions outside F_h require no capacity, directly reflecting the tier-specific activation logic.

To represent the operational burden imposed on candidate hubs, workload is decomposed into three components: upward flow generated within the hub's service territory, downward flow arriving for local distribution, and transfer flow passing through the hub. This distinction matters because the induced workload components consume hub resources differently. Upward and downward distribution flows require handling, sorting, and transfer across mover classes, whereas transfer flow is more closely associated with cross-docking and staging operations. Treating all assigned flow as a single undifferentiated volume would therefore distort module requirements, particularly at topologically central hubs where a significant share of activity may be transfer-oriented rather than origin- or destination-based. At the pre-design stage, these components are estimated from aggregate demand and indicative flow patterns rather than resolved routing, consistent with the structural scope of the problem.

To address this pre-design problem, we develop an iterative, multi-tier screening procedure that builds on the above representation to reduce the candidate space while preserving cross-tier coherence in coverage, workload, and modular capacity requirements.

4 Methodology

4.1 Framework Overview

The proposed framework operationalizes this logic through four interacting components: tier-wise coverage-based screening, modular-capacity-aware feasibility assessment, KPI-based performance evaluation, and iterative cross-tier refinement. The screening logic is coverage based but unlike single tier approaches (Muthukrishnan et al., 2021), a candidate is retained

only if it provides adequate same tier coverage as well as viable upward and downward integration within the broader hub structure.

Within each iteration, the procedure follows a tier-ordered progression from area to local to proximity. At each tier, screening is governed by operational admissibility rules including tier-specific service thresholds, minimum redundancy requirements, workload bounds compatible with modular capacity, and feasible inter-tier links under prescribed separation thresholds. These workload bounds also capture bypass-induced burden at the aggregate pre-design level. When estimated flows bypasses an intermediate tier, the corresponding handling, transfer, staging, or re-containerization workload is assigned to the retained hub where it next interacts with the network, rather than to the skipped tier.

Because lower-tier performance may reveal upstream structural inadequacies, the framework iterates between a top-down screening pass and a bottom-up feedback pass in which KPI outcomes spanning efficiency, serviceability, resiliency, equity, and sustainability identify whether upstream selections require revision. When revision is warranted, the subsequent top-down pass restarts from the highest tier flagged, allowing structural adjustments to propagate downward. This alternation continues until the retained candidate sets and their KPI profiles stabilize, yielding a reduced candidate pool that is progressively aligned with cross-tier service, load, and accessibility conditions.

4.2 Key Performance Evaluation Metrics

Complementing the coverage-based screening model, we assess retained candidate sets through a KPI framework spanning five dimensions: *efficiency, serviceability, resiliency, equity, and sustainability at each tier*. Table 2 summarizes the selected metrics along with their mapping to the performance dimensions. These KPIs serve two roles: they support comparison of alternative reduced candidate sets beyond basic feasibility, and they provide the signals used for iterative cross-tier refinement. Poor downstream performance in any dimension therefore triggers revision of upstream selections and assignments.

Table 2: Key Performance Indicators and their performance dimensions

Metric	EF	SE	RS	EQ	SU
Number of Selected Hubs	○		○		○
Total System Cost	○				
Cost per Shipment	○				○
Distance per Shipment	○	○			
Travel time per Shipment		○		○	
Maximum travel time to assigned hubs		○		○	
Average travel time to assigned hubs		○	○		
Average number of zones assigned to hubs		○	○		
Average number of hubs assigned to zones	○		○		
Fraction of unused capacity			○		
Variation in load distribution across hubs assigned to zones				○	
Variation of travel time to assigned hubs across zones				○	
Demand-weighted Gini Index				○	
Max/Min Service Time Ratio					○
Total transportation distance					○
Emissions per Shipment					○
Population exposed to freight activity near selected hubs					○
Total Infrastructure Emissions (Hubs + Modules Footprint)					○

EF: Efficiency **SE:** Serviceability **RS:** Resilience **EQ:** Equity **SU:** Sustainability

4.3 Iterative Approach

Let $\mathcal{H}^{A}_{(r)}, \mathcal{H}^{L}_{(r)}, \mathcal{H}^{P}_{(r)}$ denote the retained area, local and proximity hub candidate sets at iteration r. At each iteration, screening at a given tier is conditioned not only on same-tier coverage and capacity considerations, but also on aggregate flow representation derived from

the immediately lower tier. Thus, area hub screening relies on local-tier flow estimates, local hub screening requires proximity-tier, and proximity hub screening requires unit zone demand.

4.3.1 Top-Down Screening Pass

Each iteration begins at the highest tier whose KPI profile indicates a need for structural revision, after which the screening proceeds downward through the remaining tiers. The pass may start at the area, local, or proximity tier depending on where revision is required, but once initiated, it follows a top-down order. This order is useful *operationally* because upper-tier hubs establish the main service and transfer for lower tiers, and *computationally* because fixing higher-tier selections restricts the feasible service, assignment, and connectivity configurations of lower tiers, thereby reducing the downstream search space. At each tier, the model determines hub selection**,** assignment, and capacity decisions under a common constraint structure comprising coverage**,** assignment activation**,** assignment capacity**,** flow conservation**,** and capacity deployment, and KPI-related constraints.

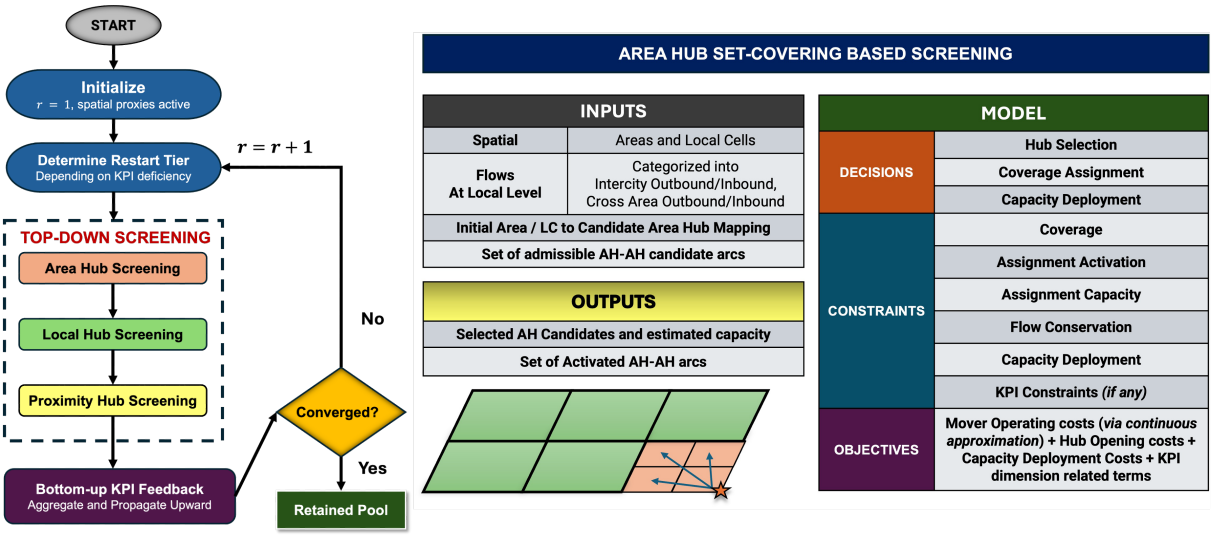


Figure 2: Iterative Multi-Tier Pre-Design Approach with Illustrative Area-Hub Screening Model

A key feature of the procedure is that the aggregate flow representation used in screening is progressively refined across iterations. In the first iteration, since retained lower tier hubs are not yet known, screening relies on spatial proxies drawn from the next lower spatial level. For example, area hub screening initially uses local cell centroids to approximate the local-tier hub structure, and local hub screening analogously relies on unit-zone-based proxies for proximity-tier activity. From the second iteration onward, these proxies are replaced by flow estimates derived from the retained lower-tier hub structure from the previous iteration. Thus, area screening uses flows derived from retained local hubs, local screening uses flows derived from retained proximity hubs, and proximity screening relies on unit-zone demand.

4.3.2 Bottom-Up Feedback

Once the top-down pass is completed, KPI outcomes from the lower tiers are aggregated and propagated upward. This feedback does not involve detailed integrated re-optimization, rather, it indicates whether the retained downstream structure creates burdens or misalignments that call the upstream configuration into question. At the proximity tier, the framework returns indicators such as coverage difficulty, service-time patterns, load concentration, assignment imbalance, and spatial equity gaps across unit zones. At the local tier, it returns consolidation burden, transfer pressure, module stress, and evidence of misalignment with the area hub structure. Each KPI is compared against a pre-specified admissible bound, and a tier is flagged for revision whenever one or more governing indicators fall outside those limits. The next

iteration then restarts from the highest tier so flagged, allowing upstream structural adjustments to propagate through the lower tiers.

4.3.3 Convergence Check

At the end of each iteration, the framework evaluates whether both the retained candidate structure and associated KPI profile have stabilized. Convergence may be declared when successive changes in retained candidate sets and KPI values fall below prescribed tolerances, otherwise, the next iteration restarts from the highest tier whose KPI profile indicates structural inadequacy. Formally if $\mathcal{H}_{(r)}^{t}$ denote the retained candidate set at tier $t \in \{G, L, A\}$, convergence may be declared when

$$\frac{\left|\mathcal{H}_{(r)}^{t} \Delta \mathcal{H}_{(r-1)}^{t}\right|}{\max\left(1, \left|\mathcal{H}_{(r-1)}^{t}\right|\right)} \leq \epsilon_{H} \quad \forall t \in \{G, L, A\}$$

$$\left\|K_{(r)}^{t} - K_{(r-1)}^{t}\right\| \leq \epsilon_{K} \quad \forall t \in \{G, L, A\}$$

where $K_{(r)}^{t}$ denote the KPI vector at tier t, and ϵ_H and ϵ_K are user-specified tolerances. In the computational study, ϵ_H is set at the level of a small relative change in retained candidates (in the order of 1-5%), while ϵ_K is set to reflect only the marginal changes in normalized KPI values (for example, on the order of 10^{-2} to 10^{-3}, depending on scaling).The procedure terminates once the retained candidate pool stabilizes or maximum iterations is reached.

5 Case Study: Metropolitan Atlanta

We apply the proposed framework to the Metropolitan Atlanta region to illustrate its ability to screen multi-tier hub candidate pools under realistic metropolitan contexts. Metro Atlanta provides a relevant testbed because it exhibits substantial spatial heterogeneity in demand distribution, land use, and logistics activity, together with a complex urban structure that is well suited to three-tier hierarchy. The study region comprises 11 counties and is represented through 1,248 census tracts representing proximity zones, 100 local cells and 10 areas. Freight Data is constructed by disaggregating the Freight Analysis Framework (FAF5) Experimental County Flows to the tract-level using tract-specific activity indicators including population, establishment counts, employee count and per-capita income (BTS, 2024). These variables are used as allocation weights to distribute county level freight totals across tracts and to approximate differences in shipment size composition across locations, thereby reflecting variation in residential, commercial and industrial activity within the metropolitan region.

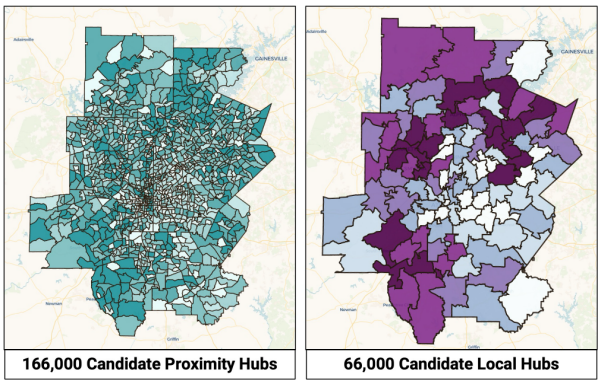


Figure 3: Proximity Zones and Local Cells with Boundary-Derived Candidate Sets for Proximity and Local Hubs

Based on this spatial and demand structure, tier-specific candidate sets are then generated for area, local and proximity hubs. Proximity and local hub candidates are derived from feasible points along the boundaries of the corresponding spatial units, reflecting their role as interface

locations between adjacent zones and local cells. This choice reflects the fact that candidate locations near shared boundaries can potentially support interface functions across multiple adjacent zones or cells, as shown in Figure 3. Area hub candidates, by contrast, are initialized from real estate data to capture facilities with characteristics consistent with upper-tier metropolitan and intercity freight function. This procedure yields approximately 166,000 proximity (PH), 66,000 local (LH) and 2,000 area candidate hubs (AH).

At each tier, the screening model is further parameterized by two design controls: μ, the minimum number of hubs required to cover each spatial component, and Ψ, the maximum number of spatial component hub can serve simultaneously. Operationally, μ controls service resiliency by determining how many alternative hubs are available to given zone, cell or area, whereas Ψ controls assignment concentration by limiting the number of spatial components assigned to any single hub.

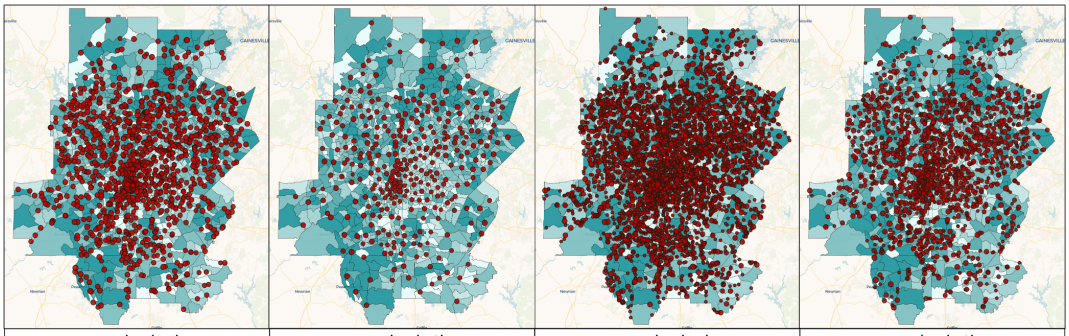


Figure 4: Retained Proximity Hub Structures Under Alternative (μ,Ψ) Settings

For urban decision makers, low values of μ and high values of Ψ correspond to leaner but centralized structures, while higher μ and lower Ψ produce more distributed configurations. In practice, tested values are chosen to span a plausible planning range from a sparse coverage, low redundancy design to more robust but operationally heavier structures. Figure 4 illustrates this effect by showing the retained proximity hubs under different (μ, Ψ) settings and Table 3 highlights the trade-off between accessibility, redundancy, model size, computational effort and cost using robust daily demand scenario.

Table 3: Sensitivity of screening outcomes to alternative (μ,Ψ) settings under a robust daily-demand scenario, showing the tradeoff between accessibility, redundancy, model size, computational effort and cost

Metric	(μ, Ψ)			
	(1,1)	(1,3)	(4,1)	(4,3)
Model Solving Time (in sec)	40	574	45	70
Number of Hubs Selected	1248	463	4992	1748
Average Number of Modules	31	82	8	22
Estimated Total System Cost (*Includes Hub Operations + Capacity + Induced Transportation Cost*)	681,273	468,738	1,451,585	551,549
Maximum distance to assigned hubs (in km)	9.2	8.8	9.9	7.8
Average distance to assigned hubs (in km)	1.9	1.7	1.8	1.4
Average number of zones assigned to hubs	1	3	4	4
Average number of hubs assigned to zones	1	1	1	3

The numerical results make the trade-offs induced by (μ, Ψ) more concrete. Relative to (1,1), the setting (1,3), reduces the number of selected hubs from 1,248 to 463, a decrease of about 63%, and lowers the estimated total system cost by 31%. It also reduces the average distance to assigned hubs from 1.9 km to 1.7 km, or roughly 11%, but at the expense of a substantial increase in solve time from 40 to 574 seconds, more than a 13x increase, and a rise in the average number of deployed modules from 31 to 82. By contrast, increases the number of selected hubs to 4,992, roughly four times the (1,1) case, and raises estimated total system cost

by 113%, reflecting the burden of stronger redundancy under limited assignment flexibility. The setting (4,3) provides the strongest accessibility and assignment flexibility, reducing the average distance to assigned hubs to 1.4 km, about 26% below the (1,1) case, and increasing the average number of hubs assigned to each zone from 1 to 3. Overall, the results show that varying (μ, Ψ) shifts the balance among accessibility, redundancy, deployment burden, and computational effort in non-monotonic ways. Following the iterative approach mentioned in Section 4.3 above to the convergence, Figure 5 presents a representative pre-designed three-tier urban network for metropolitan Atlanta. The left figure shows the spatial clustering into areas, local cells, and proximity zones; the middle shows the retained three-tier hub structure superimposed on these spatial units; and the right augments this representation with the inter-tier connectivity activated by the screening.

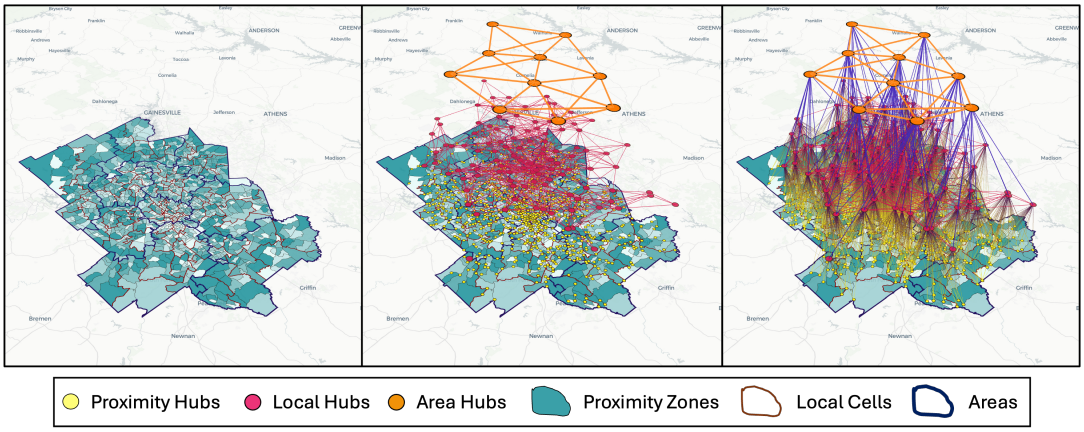


Figure 5: Representative Pre-Designed Three-Tier Hyperconnected Urban Network for Metropolitan Atlanta

6 Conclusion and Future research

This paper proposed an iterative, multi-tier, coverage-based screening framework for reducing large hub candidate pools in PI-enabled metropolitan freight systems. Methodologically, the framework extends coverage-based screening from a single-tier setting to a metropolitan hierarchy by integrating tier-wise coverage logic, modular hub capability and capacity representation, KPI-based evaluation, and iterative top-down and bottom-up refinement. Computationally, it provides a tractable mechanism for reducing very large candidate spaces by orders of magnitude while preserving the key service, connectivity, flow, and modular-capacity features required for downstream design. From a managerial standpoint, it offers a structured basis for comparing alternative pre-design configurations and the trade-offs they induce among accessibility, redundancy, cost, retained network size, and computational effort. The Metropolitan Atlanta case study demonstrates the applicability of the framework in a large and heterogeneous urban freight context, showing its ability to generate reduced candidate structures while maintaining cross-tier structural consistency.

Future research can extend the framework in several directions. First, although the screening model captures bypass-induced workload at an aggregate pre-design level, it abstracts from detailed routing, temporal peaks, queuing, and operational interactions among consolidation, transfer, staging, and re-containerization. These aspects can be examined by coupling the retained hub structures with routing, flow-allocation, while also testing sensitivity to alternative zone, cell and area partitions. Second, since cost savings depend on modular deployment assumptions, richer module-cost functions and sensitivity analyses are needed to assess when additional modules remain economically justified. Third, the observed trade-offs among accessibility, redundancy, cost, retained network size, and computational effort suggest opportunities for multi-objective formulations or adaptive parameter selection strategies.

References

- Behbahani, H., Nazari, S., Jafari Kang, M., & Litman, T. (2019). A conceptual framework to formulate transportation network design problem considering social equity criteria. Transportation Research Part A: Policy and Practice, 125, 171–183. https://doi.org/10.1016/j.tra.2018.04.005
- Ben Mohamed, I., Klibi, W., Sadykov, R., Şen, H., & Vanderbeck, F. (2023). The two-echelon stochastic multi-period capacitated location-routing problem. European Journal of Operational Research, 306(2), 645–667. https://doi.org/10.1016/j.ejor.2022.07.022
- BTS. (2024). Freight Analysis Framework Version 5 (FAF5): Experimental County-Level Estimates. https://www.bts.gov/faf/county
- Dastgoshade, S., Shafiee, M., Klibi, W., & Shishebori, D. (2022). Social equity-based distribution networks design for the COVID-19 vaccine. International Journal of Production Economics, 250. https://doi.org/10.1016/j.ijpe.2022.108684
- Farahani, R. Z., Hekmatfar, M., Fahimnia, B., & Kazemzadeh, N. (2014). Hierarchical facility location problem: Models, classifications, techniques, and applications. In Computers and Industrial Engineering (Vol. 68, Number 1, pp. 104–117). Elsevier Ltd. https://doi.org/10.1016/j.cie.2013.12.005
- Garbers, A., P. Muthukrishnan, & Montreuil, B. (2025). Pragmatic Hyperconnected Freight Systems Design: A Data-Driven Case Study in the U.S. Southeast Megaregion. 11th International Physical Internet Conference. https://repository.gatech.edu/server/api/core/bitstreams/0efed6e1-6487-4980-b202-d45365aa8120/content
- Janjevic, M., Merchán, D., & Winkenbach, M. (2021). Designing multi-tier, multi-service-level, and multi-modal last-mile distribution networks for omni-channel operations. European Journal of Operational Research, 294(3), 1059–1077. https://doi.org/10.1016/j.ejor.2020.08.043
- Jena, S. D., Cordeau, J. F., & Gendron, B. (2015). Dynamic facility location with generalized modular capacities. Transportation Science, 49(3), 484–499. https://doi.org/10.1287/trsc.2014.0575
- Klibi, W., & Martel, A. (2012). Modeling approaches for the design of resilient supply networks under disruptions. International Journal of Production Economics, 135(2), 882–898. https://doi.org/10.1016/j.ijpe.2011.10.028
- Liu, X., Klibi, W., & Montreuil, B. (2026). Modular and Mobile Capacity Planning for Hyperconnected Supply Chain Networks. http://arxiv.org/abs/2601.11107
- Liu, X., P. Muthukrishnan, & Montreuil, B. (2025). Network Design and Capacity Management in Hyperconnected Urban Logistic Networks. 11th International Physical Internet Conference. https://doi.org/https://doi.org/10.35090/gatech/13235
- Montreuil, B., Buckley, S. M., Faugère, L., Khir, R., & Derhami, S. (2018). Urban Parcel Logistics Hub and Network Design: The Impact of Modularity and Hyperconnectivity. https://api.semanticscholar.org/CorpusID:53636972
- Montreuil, B., Ballot, E., & Tremblay, W. (2014). Modular Design of Physical Internet Transport, Handling and Packaging Containers. In *Progress in Material Handling Research* (Vol. 13). https://minesparis-psl.hal.science/hal-01487239v1
- Muthukrishnan, P., Faugère, L., & Montreuil, B. (2021). Potential Access Hub Network Design Optimization in Hyperconnected Urban Logistics. https://doi.org/https://doi.org/10.35090/gatech/7901
- Ortiz-Astorquiza, C., Contreras, I., & Laporte, G. (2018). Multi-level facility location problems. In European Journal of Operational Research (Vol. 267, Number 3, pp. 791–805). Elsevier B.V. https://doi.org/10.1016/j.ejor.2017.10.019
- Silva, W. A., Jena, S. D., & Jeswani, K. (2023). Mathematical formulations for multi-period network design with modular capacity adjustments. In INFOR (Vol. 61, Number 4, pp. 445–489). Taylor and Francis. https://doi.org/10.1080/03155986.2023.2229208

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE
Physical Internet Center

Digital Shadow-Enabled Validation of Deployable π-Nodes for Distributed Construction Networks

Miguel Campos[1,2], Benoit Montreuil[1,2], Leon McGinnis[1,2], Todd Ullom[3], Mingze Li[1,2], Zhihan Liu[1,2], Julien Maurice[1,2], Jorge Garcia[1,2], Mahsa Mahdavisharif[1,2], Jihye Jung[1,2], Kamyar Fatemifar[1,2], Negin Nadji[1,2] and Reem Ashima[1,2]

1. Georgia Tech Physical Internet Center, Supply Chain & Logistics Institute
2. School of Industrial & Systems Engineering, Georgia Institute of Technology, Atlanta, USA
3. MiTek Inc.

Corresponding author: mcampos@gatech.edu

Abstract: *This paper proposes a Digital Shadow–enabled framework for execution-level validation of deployable RapidPlants (RPs) envisioned as manufacturing π-nodes within the Physical Internet. A high-fidelity discrete-event agent-based simulation platform is calibrated using task-level execution data from a pilot DistributedConstruction™ assembly case. Calibration focuses on a consistent subset of subassembly tasks to derive stochastic representations of processing variability and disruption effects. A two-stage simulation framework combines deterministic replay of observed execution with Monte Carlo forecasting from intermediate system states, generating performance distributions and confidence intervals for key indicators such as makespan and overtime. Results show that the Digital Shadow produces accurate and progressively improving forecasts as execution unfolds, with stable convergence toward observed outcomes. The study highlights the importance of capturing moderate execution delays alongside disruption events and demonstrates the feasibility of uncertainty-aware validation of deployable Rapid Plants prior to physical implementation.*

Keywords: *Physical Internet; Digital Shadow; Digital Twin; Discrete-Event Agent-Based Simulation; Distributed Construction; Modular Construction; Manufacturing π-Node; Monte Carlo Experimentation; Execution-Level Validation*

Physical Internet (PI) Roadmap Fitness*:☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories)*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

1 Introduction

The Physical Internet (PI) envisions a transformation of logistics and production systems into open, modular, and interoperable networks in which standardized nodes collaborate to move, transform, and deliver goods efficiently (Montreuil, 2011; Ballot et al., 2013). Within this paradigm, manufacturing facilities evolve from isolated production units into deployable π-nodes capable of integrating dynamically into distributed supply networks. This transformation is particularly relevant in offsite construction environments, where building components and modules are assembled in manufacturing-type facilities before shipment to a jobsite.

Within this broader domain, MiTek's DistributedConstruction™ methodology combines optimized processes, technology, and a component manufacturing network to assemble smart panels and full modules offsite before delivery to building sites. In this methodology, RapidPlants™ operate as temporary deployable assembly facilities where module assembly occurs near project demand using the DistributedConstruction™ approach.

A central challenge in this context is the validation of project execution feasibility prior to physical deployment. Distributed construction environments are characterized by high variability, project-specific configurations, and frequent disruption events, which collectively limit the predictive accuracy of traditional planning artifacts such as labor schedules and takt-based production programs. As a result, decision makers lack reliable tools to assess whether a planned offsite construction facility using Distributed Construction will achieve its intended performance and targets once deployed.

Digital Twin research has made significant progress in conceptualizing and implementing virtual representations of physical systems, particularly at the machine, process, and production-system levels (Grieves and Vickers, 2017; Uhlemann et al., 2017; Kritzinger et al., 2018; Cimino et al., 2019; Tao et al., 2019). However, the literature continues to report implementation, scalability, and model-maintenance challenges when moving from localized Digital Twin applications toward reusable, facility-scale simulation infrastructures capable of representing execution dynamics under uncertainty (Sharma et al., 2022; Zhang et al., 2022; Sun et al., 2025; Overbeck et al., 2024). This gap is especially critical in distributed and modular construction settings, where execution-level variability and disruption propagation play a dominant role in determining system performance.

This paper extends prior work on PI-inspired agile assembly-center simulation and hyperconnected assembly-factory design (Campos et al., 2023; Bao et al., 2024) by proposing a Digital Shadow–enabled approach for validating deployable RapidPlants™ as manufacturing π-nodes. The approach leverages a high-fidelity discrete-event agent-based simulation platform calibrated using real execution data, enabling the generation of stochastic performance outcomes and confidence intervals through Monte Carlo experimentation. By comparing simulated results with observed execution outcomes, the proposed framework provides a quantitative basis for validating the predictive capability of the Digital Shadow and assessing the feasibility of production systems prior to deployment.

In doing so, this work contributes to the advancement of PI-enabled manufacturing systems by bridging the gap between planning and execution, providing a data-driven mechanism for uncertainty-aware validation, and supporting the design and deployment of reliable and scalable π-node production facilities.

2 Background

2.1 Physical Internet and Distributed Construction Systems

The Physical Internet proposes a systemic redesign of logistics and production networks based on modularity, standardization, and interconnectivity, enabling assets and facilities to operate as shared and interoperable resources (Montreuil, 2011; Ballot et al., 2013). Within this framework, manufacturing facilities are conceptualized as π-nodes that can be dynamically integrated into distributed networks, supporting flexible and scalable production configurations. In distributed construction, this vision materializes through deployable

assembly facilities (RapidPlants) that can be located near project sites to reduce transportation costs and improve responsiveness.

In this paper, the term offsite construction refers to the assembly of building components, including roof trusses, wall panels, and full 3D modules, in a manufacturing-type facility before shipment to a jobsite. Modular construction refers more specifically to the traditional offsite construction approach in which large portions of a building are produced offsite, often without fundamentally redesigning the production system for productivity. Distributed Construction refers to MiTek's methodology, which combines optimized processes, technology, and MiTek's component manufacturing network to assemble smart panels and full modules offsite before delivery. In this context, a RapidPlant is a temporary facility where modules are assembled using the DistributedConstruction methodology.

Recent work in modular integrated construction has explored digital integration across the construction lifecycle, including IoT-enabled BIM platforms that enable tracking, coordination, and rule-based decision support (Zhai et al., 2019). Extensions of these platforms incorporate blockchain-enabled Digital Twin collaboration for fit-out operations in modular integrated construction (Jiang et al., 2023), as well as out-of-order synchronization frameworks for off-site fit-out operations in smart prefabrication yards (Ding et al., 2025). These contributions demonstrate the increasing maturity of digital infrastructures supporting distributed and modular construction, yet they remain largely focused on information integration and coordination rather than the validation of execution-level performance within manufacturing facilities.

2.2 Digital Twin, Digital Shadow, and Execution-Level Modeling

Digital Twin (DT) technologies have been widely recognized as a cornerstone of Industry 4.0, enabling real-time monitoring, analysis, and control of physical systems (Grieves and Vickers, 2017; Tao et al., 2019; Javaid et al., 2023). The classification proposed by Kritzinger et al. (2018) distinguishes between Digital Models, Digital Shadows (DS), and Digital Twins based on the directionality of data exchange. While Digital Twins enable bidirectional interaction between physical and virtual systems, Digital Shadows are characterized by one-way data flow from the physical system to the digital representation.

In practice, many Digital Twin implementations remain constrained to machine-level or subsystem-level applications, with limited scalability to full-facility representations capable of capturing complex operational dynamics (Cimino et al., 2019; Sharma et al., 2022). Digital Shadows provide a practical intermediate stage between offline Digital Models and full bidirectional Digital Twins, enabling one-way physical-to-digital state reflection and decision support before closed-loop control becomes feasible (Kritzinger et al., 2018; Negri et al., 2020; Legato et al., 2026). However, their application at the execution level of manufacturing facilities, where task-level variability, resource constraints, and disruption effects interact, remains underdeveloped.

2.3 Simulation and Uncertainty in PI Systems

Simulation-based approaches have long been used to analyze logistics and manufacturing systems, particularly in contexts characterized by uncertainty and complex interactions (Montreuil et al., 2015; Goyal et al., 2016; Kim et al., 2020). In the PI and hyperconnected logistics domains, discrete-event and agent-based simulation models have been applied to

evaluate mobility webs, urban fulfillment systems, parcel logistics networks, and hyperconnected megacity logistics systems under varying configurations (Hakimi et al., 2012; Goyal et al., 2016; Kim et al., 2021; Shaikh et al., 2023; Kaboudvand and Montreuil, 2024).

Monte Carlo experimentation enriches these approaches by enabling the generation of performance distributions rather than deterministic outputs, providing insights into variability, risk, and confidence intervals. Despite these advances, there remains a lack of approaches that combine task-level execution modeling, empirical calibration from real operations, and stochastic validation of full-facility performance within a unified framework. This limitation motivates the development of a Digital Shadow approach capable of linking planning data, observed execution behavior, and stochastic simulation to support robust validation of production systems.

3 Digital Shadow for Execution-Level Validation

This work adopts the Digital Model-Digital Shadow-Digital Twin framework proposed by Kritzinger et al. (2018) and focuses on the Digital Shadow paradigm as a mechanism for validating manufacturing systems prior to deployment. In this context, a Digital Shadow is defined as a data-driven virtual representation of a manufacturing facility that ingests planning data and leverages historical execution data for calibration to generate simulated performance outcomes without direct feedback to the physical system. This definition distinguishes it from Digital Twins, which are typically designed for real-time control and optimization, and instead emphasizes its role as a validation and decision-support tool.

The Digital Shadow implemented in this study was developed using the commercial simulation software Anylogic version 8.9.7, and it operates at the execution level of a RapidPlant, capturing the internal dynamics of task processing, resource allocation, and workflow interactions. The model represents the facility as a hierarchical system composed of centers, stations, and workspaces, within which tasks are executed according to predefined precedence relationships and resource constraints. Production programs derived from labor schedules are used to initialize the simulation, enabling the representation of takt-based execution logic while allowing for stochastic variability in task durations and disruption events.

A central feature of the proposed approach is its data-driven calibration using real execution data collected from a pilot assembly operation. Task-level observations, including processing times and disruption durations, are mapped to their corresponding planned task instances, enabling the estimation of stochastic parameters that reflect both baseline execution variability and disruption induced delays. Processing-time variability is modeled using multiplicative factors derived from the ratio between observed and planned task durations, while disruption effects are modeled as additive stochastic time components conditioned on task level disruption incidence.

To evaluate system performance under uncertainty, the Digital Shadow employs Monte Carlo experimentation, executing multiple simulation replications in which processing and disruption parameters are sampled from empirically calibrated distributions. This approach produces distributions of key performance indicators, such as project makespan and overtime, enabling the assessment of variability and risk. By comparing observed execution outcomes with simulated confidence intervals, the Digital Shadow provides a quantitative basis for validating its predictive capability and, consequently, for assessing the feasibility of deployable manufacturing π-nodes prior to implementation. The proposed approach explicitly incorporates

task-level execution data and disruption observations into the calibration of stochastic simulation parameters, enabling validation at the operational level of facility execution. This capability enables the Digital Shadow to serve not only as a descriptive representation of past execution, but as a predictive tool for assessing the feasibility of planned production systems under uncertainty.

4 Case Study and Data Calibration

This study is based on a pilot RapidPlant execution conducted within a distributed construction context with an industry partner, in which detailed task-level execution data were collected during the assembly of a unit within a pilot RapidPlant operation. The objective is to evaluate the capability of the proposed Digital Shadow to transform deterministic planning inputs into stochastic task representations capable of reproducing execution level performance.

Execution data was collected for 229 tasks, from which a subset of 64 unique tasks corresponding to the subassembly process was selected for calibration. This restriction ensures consistency in task type, resource allocation, and execution context, thereby reducing confounding variability associated with heterogeneous operations. Figure 1 illustrates the discrepancy between planned and actual task durations for the calibrated subset where the tasks are presented in their actual execution order, while Figure 2 illustrates a representative comparison between the real system and its simulated counterpart.

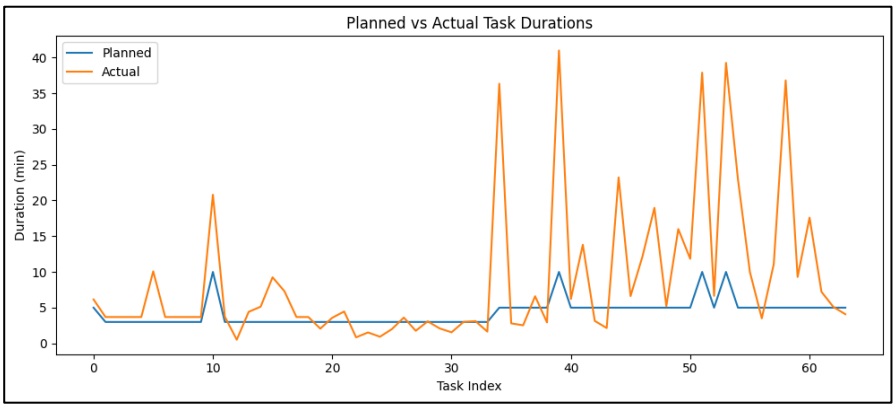


Figure 1. Planned vs. actual task durations

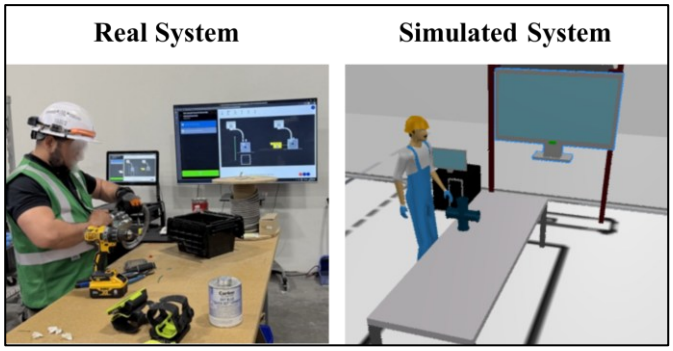


Figure 2. Real vs. simulated system subassembly station

The planning baseline is defined by a labor schedule specifying task sequences, precedence relationships, and planned durations. These planned durations represent the initial estimates

used to define the production program. The results show that the initial deterministic estimates were poor predictors of realized execution performance, particularly in the later stages of the process where disruptions and delays dominate, mainly because workers were allowed to postpone subassemblies that had incomplete or damaged parts towards advancing with the ready ones, leaving the most problematic subassemblies towards the end of the execution. The purpose of the calibration process is therefore not to refine the original schedule itself, but to construct a systematic transformation from these deterministic estimates into stochastic task durations that more accurately reflect observed execution behavior.

To quantify the discrepancy between planned and actual execution, baseline processing multipliers are defined as

$$\alpha_i = \frac{T_i^{\text{obs}} - T_i^{\text{disr}}}{T_i^{\text{plan}}} \quad (1)$$

where T_i^{obs} is the total observed task duration, T_i^{disr} is the recorded disruption time, and T_i^{plan} is the planned duration. The resulting multipliers represent baseline execution behavior excluding explicitly recorded disruption effects.

The empirical distribution of α_i is strongly right-skewed. For the calibrated subset, the observed statistics are a minimum of 0.17, a 10th percentile of 0.51, a median of 1.21, a mean of 1.58, a 90th percentile of 3.57, and a maximum of 4.59. These values show that the vast majority of task times were underestimated in the planning phase. Those that were overestimated were so in a much smaller amount as compared to those that were underestimated. These are not necessarily delays but rather a potentiality between underestimated duration by planners and underestimation of stochasticity and/or disruption potential by planners (with the average potentially correct), and probably an underestimation of learning curve effects.

It is important to note that the separation between baseline processing time and disruption effects is subject to measurement uncertainty. In practice, some long observed task durations may include disruption effects that were not fully captured in the recorded disruption intervals. As a result, the baseline multiplier distribution may contain residual contamination from disruption-induced delays, introducing a degree of estimation error.

Alternative parameterizations were evaluated during calibration, revealing that formulations overly influenced by extreme observations produced unstable forecast behavior. The percentile-based parameterization adopted in this study provided a balanced and robust representation of execution variability and was therefore retained. Accordingly, baseline variability is modeled using the triangular distribution $\alpha \sim \text{Triangular}(0.51, 1.21, 3.57)$, where the parameters correspond to the empirical 10th, 50th, and 90th percentiles of the calibrated subset. This choice provides a robust representation of lower-bound, typical, and upper-range execution conditions while limiting the influence of extreme values.

Precisely modeling infrequent events as disruptions and their resolutions is not a well-studied problem, and it requires knowing potentiality, impact and resolution strategies in advance. For this reason, we propose a simplified approach used to demonstrate the DS in which the disruption realized time until going back to normal operation is assumed regardless of the resolution implemented. Therefore, disruption effects are modeled separately to isolate their

contribution to task duration. For tasks affected by disruptions, disruption multipliers are defined as

$$\delta_i = \frac{T_i^{\text{disr}}}{T_i^{\text{plan}}} \tag{2}$$

The disruption dataset is much smaller and more variable than the baseline dataset. For the available observations, the minimum, median, and maximum values are 0.20, 0.79, and 6.07, respectively. Because of the limited number of disruption observations, higher-order statistics such as percentiles or mode are less reliable. For this reason, disruption magnitude is modeled using a triangular distribution based directly on the observed range, $\delta \sim \text{Triangular}(0.20, 0.79, 6.07)$, while disruption occurrence is modeled as a Bernoulli process with probability $p_d = 0.13$, estimated from the observed proportion of disrupted tasks in the calibrated subset. Within the Digital Shadow, the stochastic duration of each task is computed as

$$T_i = T_i^{\text{plan}} \cdot (\alpha_i + \mathbb{I}_i(p_d)\delta_i) \tag{3}$$

where α_i is sampled from the baseline distribution, δ_i is sampled from the disruption distribution, and $\mathbb{I}_i$ is a Bernoulli random variable indicating whether a disruption occurs for task i with probability p_d.

The calibration process was conducted iteratively by comparing the forecasting behavior of alternative stochastic parameterizations. While this process proved effective and insightful for identifying stable and accurate formulations, it also highlights an opportunity for future work in simulation optimization-based parameter tuning. With larger historical datasets, such methods could systematically search for parameter values that minimize forecast error while improving robustness to sampling variability and measurement uncertainty.

5 Monte Carlo Experimentation and Validation

To evaluate the predictive capability of the Digital Shadow, a two-stage simulation framework was implemented, combining deterministic replay of the observed execution with stochastic forecasting based on the calibrated task duration model, as illustrated in Figure 3.

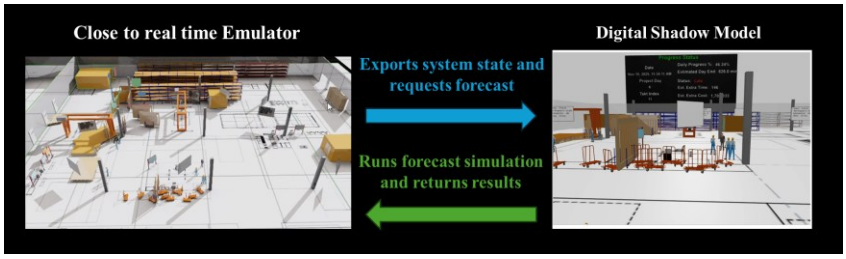


Figure 3. Emulator and Digital Shadow feedback loop

In the first stage, an execution model reproduces the pilot case using the actual observed task durations. This model operates in accelerated time, at approximately 12 seconds of simulated time per second of wall-clock time, allowing the execution process to be replayed in a close-

to-real-time manner. The purpose of this stage is to recreate the evolution of the system state as it occurred during the pilot, including task completions, active work, and disruption effects.

At each simulated hour from the start of execution, the current system state is extracted and used to initialize the Digital Shadow which runs in virtual time. These intermediate states include information on completed tasks, tasks in progress, and the remaining workload. From each of these states, a Monte Carlo forecasting experiment is performed. For each forecast point, 50 independent replications are executed, in which only the remaining tasks are simulated stochastically according to the duration formulation defined in Section 4. Completed tasks retain their observed durations from the execution replay, ensuring consistency between realized and forecasted portions of the process. The outputs of each Monte Carlo experiment are aggregated to estimate project makespan, including the mean and the corresponding 95% confidence interval. Repeating this process at successive forecast points produces a sequence of forecasts, enabling direct assessment of prediction accuracy over time.

The primary objective of this framework is to evaluate whether the Digital Shadow can produce accurate and progressively improving predictions as more execution information becomes available. Validation is achieved when the observed final makespan lies within, or in close proximity to the simulated confidence intervals, and when the forecast exhibits stable convergence toward the realized value across successive forecast points.

Alternative formulations were evaluated during calibration, revealing that certain parameterizations can significantly degrade forecasting stability. In particular, formulations that overemphasize extreme observations were observed to produce erratic forecast trajectories and inconsistent convergence patterns. The percentile-based baseline formulation adopted in Section 4 provided the most stable and accurate forecasting behavior and was therefore retained for the final model.

Beyond validation, this framework provides insights into how uncertainty evolves during execution. Early forecasts are associated with wider confidence intervals due to the larger volume of remaining work and limited information, while later forecasts exhibit narrower intervals as the system approaches completion. The approach therefore provides not only an estimate of expected completion time, but also a dynamic characterization of uncertainty across the execution horizon.

6 Results and Discussion

The results of the Digital Shadow validation are presented through a sequence of rolling-horizon forecasts generated at discrete time points during execution. At each forecast point, the current system state is used to estimate both the remaining makespan and the projected completion time, enabling direct comparison between deterministic planning, stochastic forecasting, and observed execution outcomes. Figures 4 and 5 together provide a complementary view of forecast performance, capturing both the evolution of remaining work and the corresponding projected completion time at each decision point.

Figure 4 presents the evolution of the expected remaining makespan at each forecast point. Each marker represents a forecast generated at a specific time during execution. The figure compares three quantities: the deterministic remaining makespan using planned processing times, the Digital Shadow forecast, and the actual remaining makespan observed in the system.

The results show that the deterministic plan consistently underestimates the makespan, particularly in the early stages of execution.

For example, at the initial forecast point, the planned estimate suggests that 220 minutes are remaining to finish the daily work, whereas the actual remaining time is more than twice that value, closer to 500 minutes. In contrast, the Digital Shadow forecast provides a substantially more accurate and consistent estimate, closely tracking the actual remaining time across all forecast points. The associated confidence intervals further show that uncertainty is higher at the beginning of execution and decreases progressively as more information becomes available. At each forecast point, the Digital Shadow provides an estimate of the remaining makespan that is not directly observable from the deterministic plan, enabling more accurate anticipation of future delays.

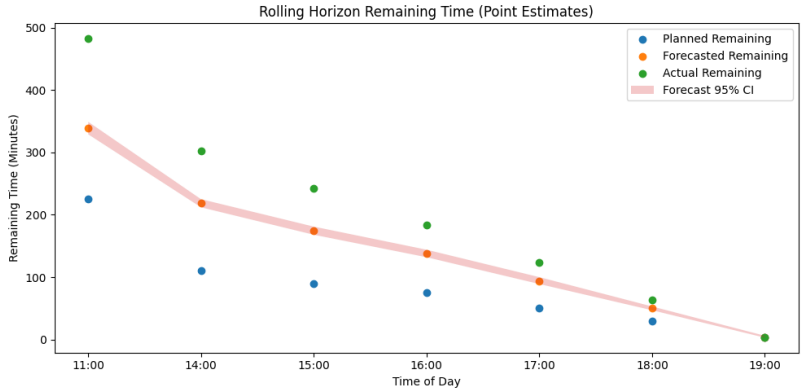


Figure 4. Rolling horizon Digital Shadow Forecast for the remaining processing time vs planned and actual remaining time

Figure 5 complements this analysis by showing the projected completion time at each forecast point. Each marker represents the estimated project finish time based on the information available at that specific moment. The deterministic planned projection remains systematically optimistic, consistently predicting completion significantly earlier than the observed outcome. Even at later stages of execution, when most tasks have already been completed, the planned projection continues to underestimate the remaining duration and suggests an unrealistically early finish time. In contrast, the Digital Shadow produces progressively improving estimates of the final makespan, with the forecast mean converging toward the actual completion time and the confidence interval narrowing as execution progresses. This demonstrates that, even when most of the work has been completed, the deterministic plan fails to capture residual uncertainty, while the Digital Shadow continues to provide meaningful predictive insight.

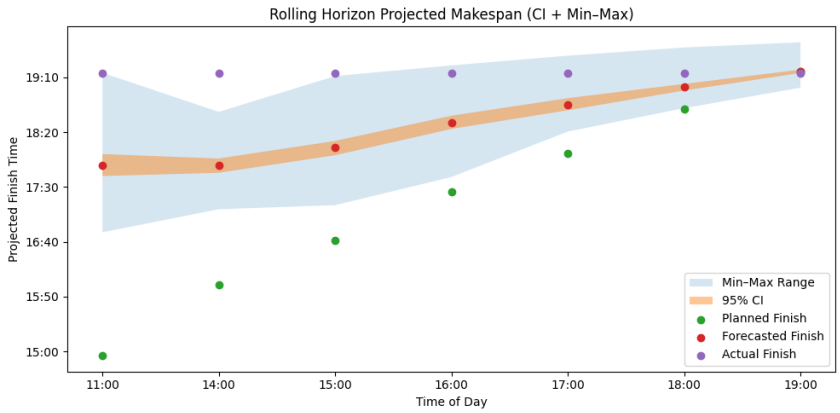


Figure 5. Rolling horizon Digital Shadow Forecast for the makespan vs planned projected makespan

The observed convergence behavior is consistent with the structure of the execution process. Early forecasts are based on limited information and are therefore subject to higher uncertainty. As execution progresses and delayed tasks become observable, the model incorporates this information, resulting in improved accuracy and reduced variance. This behavior is further influenced by operational practices in which workers dynamically reorder tasks to avoid disruptions, effectively postponing delayed tasks to later stages of execution. As these delays are revealed, the forecast adjusts accordingly, leading to the observed monotonic convergence toward the actual completion time.

For the final calibrated model, the Digital Shadow converges to within only a few minutes of the observed makespan, while maintaining stable and well-behaved confidence intervals across all forecast points. This demonstrates that the proposed calibration methodology successfully captures both baseline variability and disruption effects, enabling accurate and consistent forecasting performance.

From a methodological perspective, these results show that the Digital Shadow framework can effectively transform inaccurate deterministic planning inputs into reliable probabilistic forecasts. Rather than relying on planned task durations as direct predictors of execution performance, the approach uses empirical data to construct stochastic representations that reflect real system behavior. The rolling-horizon forecasting framework further strengthens this validation by evaluating prediction accuracy at multiple points in time, providing a more rigorous assessment than single-point comparisons at project completion.

From a Physical Internet perspective, these findings highlight the importance of uncertainty-aware validation for deployable manufacturing π-nodes. The ability to generate accurate, dynamically updated forecasts enables more informed decisions regarding facility deployment, resource planning, and operational risk. The Digital Shadow therefore serves as a practical mechanism for bridging the gap between planning and execution in distributed construction systems.

7 Conclusion

This paper presented a Digital Shadow enabled framework for execution level validation of deployable Rapid Plants within a Physical Internet inspired distributed construction network. The proposed approach integrates deterministic planning inputs with data driven stochastic calibration to generate probabilistic forecasts of system performance under uncertainty.

The results indicate that accurate and stable prediction of execution performance can be achieved through the proposed transformation of deterministic planning inputs into stochastic task representations. By transforming planned task durations through empirically derived stochastic multipliers, the Digital Shadow reproduces observed system behavior with high accuracy and stable convergence across successive forecast points.

The study also shows that dynamic forecasting provides a more rigorous validation framework than single-point comparisons, enabling evaluation of predictive performance throughout the execution process. This capability is particularly valuable in distributed construction environments characterized by variability and limited planning reliability. The calibration process was conducted through iterative experimentation, highlighting opportunities for future

work in simulation optimization and automated parameter tuning. Improvements in data availability and quality will further enhance the robustness of the proposed methodology.

The results presented in this study are based on a single pilot case and a limited dataset, and further validation across multiple facilities and larger datasets is required to generalize the findings. This work proposes a systematic methodology for transforming inaccurate deterministic planning inputs into reliable stochastic execution forecasts, providing a practical mechanism for uncertainty-aware validation of manufacturing π-nodes in Distributed Construction systems.

References

- Ballot, E., Montreuil, B. & Thivierge, C. (2013). Functional Design of Physical Internet Facilities: A Road-Rail Hub, *CIRRELT*.
- Bao, W., McGinnis, L., Campos, M., Liu, Z., Maurice, J., Montreuil, B., & Babalou, S. (2024). A modular and flexible design of hyperconnected assembly factory. Proceedings of the International Physical Internet Conference (IPIC), Savannah, USA. https://doi.org/10.35090/gatech/11066
- Campos, M., McGinnis, L., & Montreuil, B. (2023). Modeling and simulation of an agile assembly center in a Physical Internet-inspired manufacturing system. Proceedings of IPIC 2023, Athens, Greece. https://doi.org/10.35090/gatech/8031
- Cimino, C., Negri, E. & Fumagalli, L (2019). A review of digital twin applications in manufacturing. Computers in Industry, 113, 103130. https://doi.org/10.1016/j.compind.2019.103130
- Ding, H., Zhong, R. Y., Li, M., & Huang, G. Q. (2025). Out-of-order synchronization for off-site fit-out in smart prefabrication yard. *Automation in Construction, 177*, 106355. https://doi.org/10.1016/j.autcon.2025.106355
- Goyal, M., Cook, J., Kim, N., & Montreuil, B. (2016). Hyperconnected city logistics for furniture and large appliance industry: Simulation-based exploratory investigation. Proceedings of IPIC 2016. Available at: https://ipic2016.pi.events/sites/default/files/Workshop_TA1_Hyperconnected_City_Logistics_for_Furniture_Large_Appliance_Industry.pdf
- Grieves, M., Vickers, J. (2017). Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems. In: Kahlen, J., Flumerfelt, S., Alves, A. (eds) Transdisciplinary Perspectives on Complex Systems. Springer, Cham. https://doi.org/10.1007/978-3-319-38756-7_4
- Hakimi, D., Montreuil, B., Sarraj, R., Ballot, E., & Pan, S. (2012). Simulating a Physical Internet-enabled mobility web: The case of mass distribution in France. 9ème Conférence Internationale de Modélisation, Optimisation et SIMulation (MOSIM 2012). Available at: https://www.researchgate.net/publication/265426687_Simulating_a_physical_internet_enabled_mobility_web_the_case_of_mass_distribution_in_France
- Javaid, M., Haleem, A. & Suman, R. (2023). Digital Twin applications toward Industry 4.0: A Review, Cognitive Robotics, 3, 71-92. https://doi.org/10.1016/j.cogr.2023.04.003
- Jiang, Y., Liu, X., Wang, Z., Li, M., Zhong, R. Y., & Huang, G. Q. (2023). Blockchain-enabled digital twin collaboration platform for fit-out operations in modular integrated construction. *Automation in Construction, 148*, 104747. https://doi.org/10.1016/j.autcon.2023.104747

- Kaboudvand, S., & Montreuil, B. (2024). Simulation-Based Assessment of Hyperconnected Megacity Parcel Logistics. Logistics, 8(3), 66. https://doi.org/10.3390/logistics8030066
- Kim, N., Montreuil, B., Klibi, W., & Kholgade, N. (2021). Hyperconnected urban fulfillment and delivery. Transportation Research Part E, 145, 102104. https://doi.org/10.1016/j.tre.2020.102104
- Kritzinger, W., Karner, M., Traar, G., Henjes, J. & Sihn, W. (2018). Digital Twin in manufacturing: A categorical literature review and classification. IFAC-PapersOnLine, 51 (11), 1016-1022. https://doi.org/10.1016/j.ifacol.2018.08.474
- Legato, P., Matteucci, M., & Mazza, R. M. (2026). Towards a digital twin solution for manual order-picking operations in a wholesale distribution center. *Simulation Modelling Practice and Theory, 147*, 103246. https://doi.org/10.1016/j.simpat.2025.103246
- Montreuil, B. (2011). Toward a Physical Internet: meeting the global logistics sustainability grand challenge, *Logistics Research*, v3, 71-87. https://doi.org/10.1007/s12159-011-0045-x
- Montreuil, B., Cloutier, C., Labarthe, O., & Loubier, J. (2015). Holistic modeling, simulation and visualization of demand and supply chains. International Journal of Business Performance and Supply Chain Modelling, 7(1), 53-70. https://doi.org/10.1504/IJBPSCM.2015.068133
- Negri, E., Berardi, S., Fumagalli, L., & Macchi, M. (2020). MES-integrated digital twin frameworks. *Journal of Manufacturing Systems, 56*, 58–71. https://doi.org/10.1016/j.jmsy.2020.05.007
- Overbeck, L., Graves, S. C., & Lanza, G. (2024). Development and analysis of digital twins of production systems. *International Journal of Production Research, 62*(10), 3544–3558. https://doi.org/10.1080/00207543.2023.2242525
- Shaikh, S. J., Kim, N., Rababah, M., Montreuil, B., & Smith, J. S. (2023). Modular containerization of parcel logistics networks: Simulation-based impact assessment. Proceedings of IPIC 2023. Available at: https://repository.gatech.edu/server/api/core/bitstreams/e1a2f82c-23f3-4b62-a346-0ac544b1c1bb/content
- Sharma, A., Kosasih, E., Zhang, J., Brintrup, A., & Calinescu, A. (2022). Digital twins: State of the art theory and practice, challenges, and open research questions. Journal of Industrial Information Integration, 30, 100383. https://doi.org/10.1016/j.jii.2022.100383
- Sun, Y., Fu, Z., & Zhao, N. (2025). A modular multi-granularity simulation modelling method for manufacturing systems oriented towards digital twins. *International Journal of Production Research*. https://doi.org/10.1080/00207543.2025.2564280
- Tao, F., Zhang, M., & Nee, A. Y. C. (2019). Chapter 1 – Background and concept of Digital Twin. In F. Tao, M. Zhang, & A. Y. C. Nee (Eds.), *Digital Twin Driven Smart Manufacturing* (pp. 3–28). Academic Press. https://doi.org/10.1016/B978-0-12-817630-6.00001-1
- Uhlemann, T. H. J., Lehmann, C., & Steinhilper, R. (2017). The digital twin: Realizing the cyber-physical production system for Industry 4.0. Procedia CIRP, 61, 335-340. https://doi.org/10.1016/j.procir.2016.11.152
- Zhai, Y., Chen, K., Zhou, J., Cao, J., Lyu, Z., Jin, X, Shen, G., Lu, W., & Huang, G. Q. (2019). An Internet of Things-enabled BIM platform for modular integrated construction: A case study in Hong Kong. *Advanced Engineering Informatics, 42*, 100997. https://doi.org/10.1016/j.aei.2019.100997

- Zhang, Z., Guan, Z., Gong, Y., Luo, D., & Yue, L. (2022). Improved multi-fidelity simulation-based optimisation: Application in a digital twin shop floor. *International Journal of Production Research, 60*(3), 1016–1035. https://doi.org/10.1080/00207543.2020.1849846

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE inlecom

Operationalizing Book & Claim Through a Digital Chain of Custody for Freight Decarbonization in the Physical Internet

Harris Niavis[1], Ettore Gualandi[2], Andrea Condotta[2] and Ioanna Fergadiotou[1]
1. Inlecom Innovation, Greece
2. Gruber Logistics, Italy

Abstract: *Freight transportation decarbonization is hindered by high costs, fragmented supply chains and limited scalability of low-emission transportation services (LETS). The Book & Claim methodology enables the decoupling of environmental attributes from physical transport flows, but existing implementations face challenges related to trust, traceability and linkage to real-world activity. This paper proposes a digital platform that operationalizes Book & Claim through a verifiable chain of custody aligned with Physical Internet principles. The approach introduces a three-tier verification model combining Proof of Sustainability (PoS), Proof of Use (PoU), and Proof of Integrity (PoI), supported by a structured data model and a layered architecture leveraging zkEVM-based execution and blockchain registries. The platform enables scalable and auditable allocation of emissions reductions across stakeholders. This concept will be implemented and validated in the IKIGAI project through real-world freight decarbonization use cases.*

Keywords: *Physical Internet, book & claim, chain of custody, freight decarbonization, blockchain, smart contracts, emissions accounting.*

Physical Internet (PI) Roadmap Fitness*: ☒ PI Networks, ☒ System of Logistics Networks, ☒ Horizontal Supply Chain Alignment, ☒ Logistics/Commercial Data Platform.*

Targeted Delivery Mode-s*: ☒ Paper*

1 Introduction

Freight transportation is a significant and growing contributor to GreenHouse Gas (GHG) emissions (IEA, 2024), with decarbonization efforts hindered by high costs, limited availability of Low-Emission Transportation Services (LETS) (SFC, 2023), and the fragmented, multi-actor nature of logistics supply chains. Carriers, Logistics Service Providers (LSPs), and shippers operate across complex networks with limited coordination, making it difficult to access low-emission solutions and consistently attribute emissions reductions to specific transport activities. Approaches based on physical segregation of sustainable fuels or technologies face scalability limitations, as they require direct linkage between fuel use and individual shipments, which is often impractical in dynamic freight systems.

Addressing these challenges requires systemic changes in how logistics networks are designed, coordinated, and governed. Within the vision of the Physical Internet (PI) (Montreuil, 2011; Ballot et al., 2021), freight systems evolve into hyperconnected, open and standardized networks that enable seamless collaboration while improving efficiency, resilience and sustainability (Pan et al., 2025). However, current systems remain fragmented, with limited interoperability and insufficient mechanisms to allocate and verify emissions reductions across supply chains.

The Book & Claim methodology addresses these limitations by decoupling environmental attributes from physical transport flows, enabling emissions reductions to be allocated independently of where low-emission solutions are deployed. As a market-based measure (ITF, 2026), it supports scalable adoption of LETS but depends on robust digital infrastructures to ensure trust, traceability, and integrity. Current implementations lack interoperable and verifiable systems capable of linking certified sustainability attributes with real-world transport activity.

This paper addresses this gap by proposing a digital chain of custody platform for freight decarbonization. The approach combines a structured data model, a layered system architecture, and governance mechanisms to support the lifecycle of Book & Claim Units (BCUs). It introduces a three-tier verification model - Proof of Sustainability (PoS), Proof of Use (PoU) and Proof of Compliance (PoI) - to align certified inputs, real-world operations and digitally enforced transactions. The solution interoperates with established standards, including ISO 14083 and the GLEC Framework, and supports scalable coordination across logistics networks in line with PI principles.

The remainder of this paper is organized as follows. Section 2 presents the background and conceptual foundations. Section 3 introduces the proposed platform. Section 4 discusses industry impact and use cases, Section 5 outlines challenges, Section 6 reviews related work, and Section 7 concludes.

2 Background and Conceptual Foundations

2.1 Freight Decarbonization and Market-Based Measures

Freight transportation is a major contributor to global greenhouse gas (GHG) emissions (IEA, 2024), with continued growth expected without intervention. Decarbonization is therefore a central objective within climate policy and the PI, where sustainability complements efficiency and resilience.

LETS reduce emissions through alternative fuels, energy sources, or operational efficiencies. Sustainable fuels such as Sustainable Aviation Fuel (SAF), Hydrotreated Vegetable Oil (HVO), and Bio-LNG are key near-term solutions, particularly in hard-to-abate sectors, but their deployment is constrained by high costs, limited availability, and uneven geographic distribution (SFC, 2025).

These challenges are reinforced by structural characteristics of freight markets: high-cost premiums, limited access to sustainable fuels and fragmented supply chains involving multiple actors. This fragmentation makes it difficult to allocate emissions reductions to specific shipments or stakeholders.

Within the PI paradigm, these limitations highlight the need for interoperable and scalable coordination mechanisms. Market-based measures, including Book & Claim, enable the separation and reallocation of environmental attributes independently of physical transport flows. These mechanisms are particularly relevant in PI environments, where decoupling flows enables flexible allocation across distributed logistics networks.

2.2 Book & Claim Methodology

The Book & Claim methodology decouples environmental attributes from physical transport flows, allowing emissions reductions from low-emission fuels or technologies to be attributed independently of their physical use.

It relies on chain-of-custody models that define how sustainability attributes are tracked:

- **Physical segregation**, ensuring direct traceability through strict separation
- **Mass balance**, allowing mixing while maintaining proportional accounting
- **Book & Claim**, fully decoupling physical flows from environmental attributes through registries

While physical segregation and mass balance provide strong traceability, they are difficult to scale in complex logistics systems. Book & Claim offers greater flexibility and scalability, making it particularly suited to PI-enabled networks.

LETS can be generated directly (through physical fuel use) or indirectly (through allocation of certified attributes), enabling participation by actors without direct access to low-emission solutions.

However, integrity challenges remain, including ensuring additionality, preventing double counting and enabling traceability. Weaknesses in verification can undermine trust and enable fraudulent practices (ISCC, 2025).

A critical requirement is linking environmental attributes to real-world transport activity. This requires PoU, such as fuel invoices, fuel card transactions and Transport Management System (TMS) data. Without such linkage, systems risk relying on unverified declarations, increasing fraud risks. Robust Book & Claim systems must therefore ensure verifiable linkage between sustainability claims and real-world transport activity.

2.3 Digital Infrastructure and Chain of Custody Systems

Scaling Book & Claim requires digital infrastructure capable of supporting secure, transparent and interoperable transactions. Central to this is the registry, which records issuance, transfer and retirement of environmental attributes.

Registries ensure integrity by tracking ownership, enforcing rules, and supporting auditing. In the PI context, they must enable interoperability across distributed logistics networks.

Decentralized technologies (Niavis et al., 2023; Kouhizadeh et al., 2021; Saberi et al., 2018) enhance trust by enabling immutable records and automated rule enforcement through smart contracts. Digital identity frameworks, such as SSI and Verifiable Credentials, support authentication and regulatory compliance (eIDAS, 2024), while verification processes ensure data validity.

However, current systems remain fragmented across standards, platforms, and jurisdictions, limiting scalability. Inconsistent data models, limited interoperability, and insufficient verification mechanisms can lead to inefficiencies and risks of fraud and double counting.

Addressing these challenges requires interoperable and standardized digital infrastructures aligned with PI principles. The platform proposed in this paper builds on these foundations by integrating blockchain-based registries, smart contracts and trusted identities to enable a verifiable and scalable Book & Claim system.

3 Book & Claim Platform

3.1 System Overview and Methodology

The proposed platform operationalizes the Book & Claim methodology through a digital chain of custody that enables the issuance, transfer, attribution and retirement of LETS attributes across freight transportation networks. Within the context of the PI, the system facilitates coordination among distributed actors while ensuring interoperability, privacy, traceability and

scalability. The methodology defines the roles and interactions of four primary actors: i) low-emission solution providers, ii) carriers, iii) LSPs and iv) shippers, supported by v) an independent verification body.

The process begins with **solution providers**, who register sustainability characteristics of fuels or technologies in recognized registries through certified Proofs of Sustainability (PoS), which document the origin, production pathway and emissions characteristics of low-emission fuels or technologies in accordance with recognized certification schemes (e.g., RSB[1], ISCC). These PoS records form the basis for the issuance of BCUs, ensuring that all environmental attributes entering the system are verified and compliant with established sustainability criteria.

Carriers generate transport activity data (e.g., tonne-kilometers), and link these activities to verified proofs of physical activity (PoU)**,** including fuel invoices and fuel card transactions associated with specific vehicles. These records confirm that sustainable fuel has been physically purchased and used to perform real transport services, ensuring that BCUs are grounded in actual logistics operations rather than abstract allocations. The transport activities are linked to emissions profiles and create BCUs that represent the environmental attributes of low-emission transport operations. These BCUs are subsequently transferred to **LSPs**, who aggregate and attribute them across multiple shipments and customers. Finally, **shippers** or LSPs claim the attributed BCUs and report the corresponding emissions reductions within their inventories. Carriers and/or LSPs may also claim attributed BCUs to support their own internal decarbonization and emissions reporting objectives.

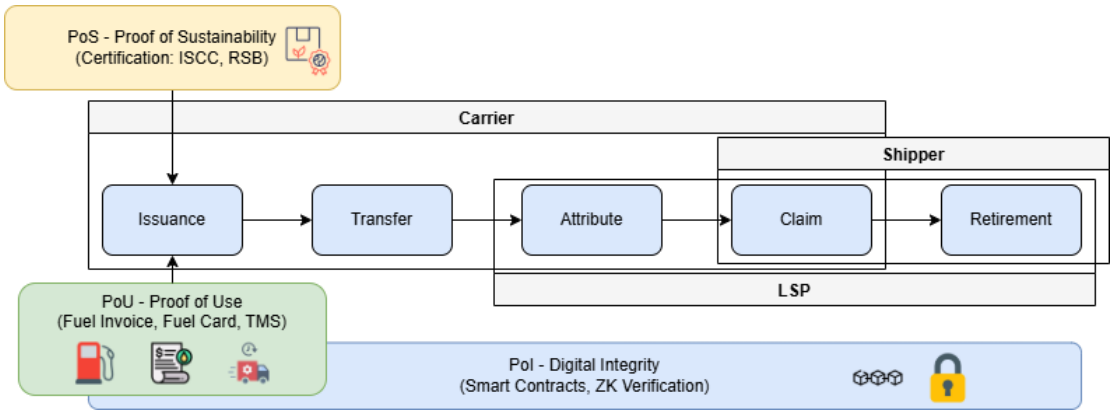


Figure 1: Three-tier Verification Model

Figure 1 illustrates the three-tier verification model underpinning the system, which is supported by the **verification body** who operates across all stages of the process, either by manually validating transactions or by setting up through the platform automated rule-based compliance checks, ensuring adherence to eligibility criteria and reinforcing trust among participants. PoS certifies the environmental attributes of fuels, PoU validates their use in real transport operations through physical evidence such as fuel invoices and TMS data, and PoI ensures that all digital transactions across the lifecycle of BCUs adhere to system rules, preventing double counting and ensuring auditability.

3.2 Data Model

The platform is underpinned by a structured data model that captures the entities, attributes, and relationships required to represent LETS and their lifecycle within a digital chain of custody.

[1] RSB, https://rsb.org/registry/

At the core of the model is the LETS unit, also referred to as a BCU, which encapsulates the environmental attributes associated with a specific low-emission transport activity or fuel input, as shown in Listing 1.

```
{

   "bcu_id": "BCU-000123",
   "fuel": { "type": "HVO", "liters": 1000 },
   "emissions": { "reduction_%": 84 },
   "pos": [{ "pos_id": "POS-ISCC-987654", "scheme": "ISCC EU", "feedstock":
"UCO"}],
   "PoU": { "invoice": "INV-987654", "tkm": 25000 },
   "PoI": { "stage": "issuance", "verified": true },
   "owner": "Carrier_Gruber",
   "status": "active"

}
```

Listing 1: Example of a Book & Claim Unit (BCU) representation in machine-readable format.

Each BCU is linked to one or more PoS, which provides verifiable metadata on the underlying fuel or technology, including certification scheme, feedstock origin and lifecycle emissions. These attributes ensure alignment with established emissions accounting frameworks, including ISO 14083 and the GLEC[2] Framework.

Transport activities are represented through standardized metrics such as tonne-kilometers (tkm), which provide the basis for linking physical transport operations to emissions profiles.

The lifecycle of a BCU follows a sequence of state transitions:

1. Issuance - creation of BCUs by carriers based on verified inputs, including PoS and PoU
2. Transfer - ownership transfer between actors (e.g., carrier to LSP or LSP to shipper)
3. Attribution - association of BCUs to specific shipments or customers
4. Claim - formal use of attributed BCUs for emissions reporting or decarbonization claims by an organization.
5. Retirement - final consumption of the BCU to prevent reuse

The lifecycle of a BCU is supported by two complementary verification artifacts: PoS and PoI. PoS is introduced at the issuance stage, remaining attached to the BCU throughout its lifecycle, while PoI is generated dynamically through smart contracts at each transaction stage to verify that all operations comply with system rules, including eligibility criteria, allocation constraints and double-counting prevention. While PoS ensures the integrity of environmental attributes, PoI ensures the correctness of their use, together forming a robust digital chain of custody to enable auditing and reporting, aligned with certification frameworks such as RSB.

In addition to PoS and PoI, the data model incorporates PoU as a core element linking digital records to physical operations. PoU includes structured data derived from fuel invoices (e.g., purchase date, fuel type, quantity, supplier) and operational data (e.g., vehicle identifiers, route information, tonne-kilometers). This ensures that each BCU is anchored to verifiable real-world transport activity.

[2] GLEC Framework, https://www.smartfreightcentre.org/en/our-programs/emissions-accounting/global-logistics-emissions-council/calculate-report-glec-framework/

3.3 System Architecture

The system architecture, shown in Figure 2, is designed as a layered digital infrastructure that enables trust and supports performance, privacy and scalability in a multi-actor environment. It integrates application interfaces, identity management, an off-chain execution layer for transaction pre-processing and batching and a Verifiable Data Registry (VDR).

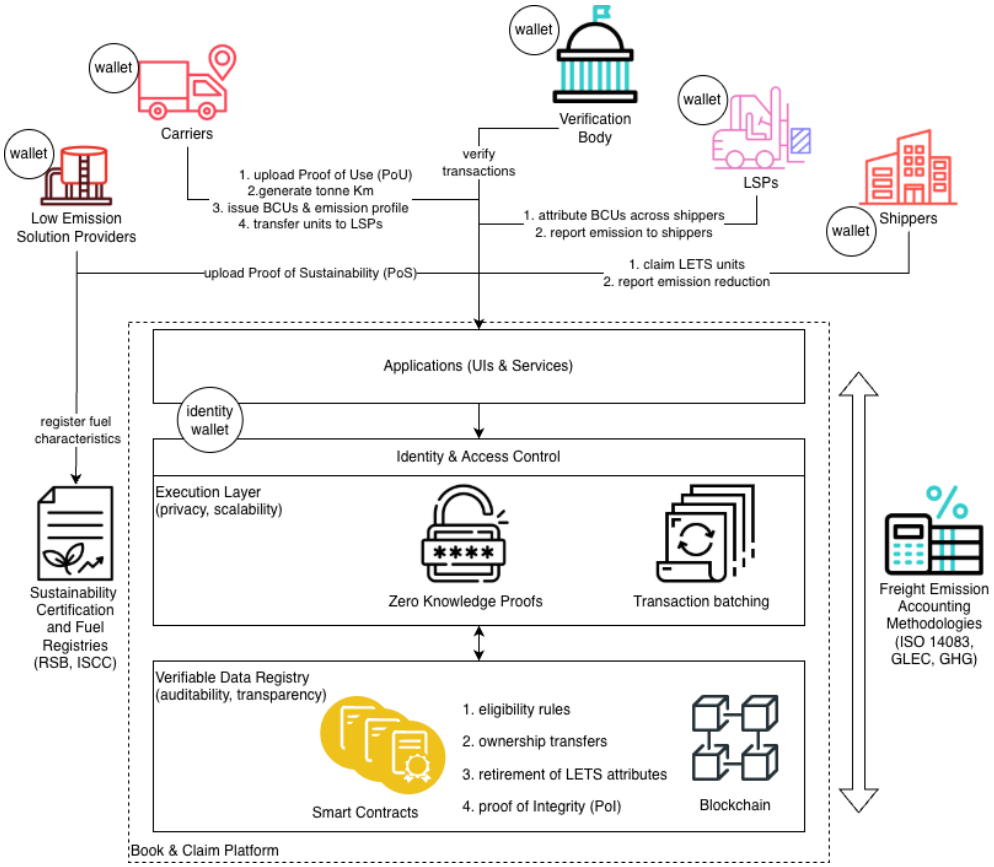


Figure 2: Book & Claim Platform

At the top, a web-based layer provides role-specific **applications and services** for shippers, carriers, LSPs and solution providers. Through this layer, users can onboard assets, issue and acquire BCUs, monitor transactions and generate emissions reports. It also enables verification bodies to perform audits, validate transactions and setup up dynamic compliance checks. This layer acts as the primary interaction point within the PI ecosystem, enabling seamless coordination across actors. An **identity and access control** component supports authentication and authorization through decentralized identity wallets. These components enable secure user onboarding, credential lifecycle management and compliance with regulatory frameworks such as eIDAS. Identity mechanisms ensure that only authorized actors can access the platform to issue, transfer, or claim BCUs, thereby strengthening trust and accountability.

Beneath the application layer, the **execution** layer enhances scalability and confidentiality while ensuring privacy-by-design. It is based on a L2 blockchain network that uses Zero-Knowledge Proofs (ZKPs) to enable verification of transactions without revealing sensitive data, while transaction batching improves throughput and reduces operational costs. These capabilities are essential for supporting large-scale adoption within the PI, where high transaction volumes and data sensitivity are expected.

At the core of the system lies a **blockchain-based VDR**, which serves as the immutable ledger for all LETS-related transactions. Smart contracts govern the behavior of the registry by enforcing eligibility rules, ownership transfers, allocation constraints and retirement conditions. Smart contracts further generate and validate PoI for each transaction, ensuring that all operations conform to predefined rules and accounting standards.

On the side, the platform is aligned with established freight emissions accounting frameworks, including ISO 14083, the GLEC Framework, and the GHG Protocol, ensuring compatibility with standardized emissions reporting. External sustainability registries (RSB, ISCC) provide PoS, while fuel providers and refueling infrastructure play an active role in supplying verified transaction data that supports PoU. They are integrated into the architecture to provide certified data on fuel characteristics and emissions factors.

Overall, the architecture reflects key PI principles, including modularity, interoperability and openness, while leveraging decentralized technologies to enable a trusted and scalable digital chain of custody.

3.4 Software Components

The platform is implemented through a set of modular software components that collectively support the functionality of the digital chain of custody.

The **web dashboard** provides a user-friendly interface for interacting with the platform, integrating data visualization, reporting tools and transaction management functionalities. It enables users to monitor LETS balances, track transaction histories and generate emissions reports aligned with standardized frameworks. Dedicated interfaces allow users to manage the lifecycle of BCUs, inspect transactions and verify associated proofs.

Digital identity wallets enable secure interaction between users and the platform. They manage cryptographic keys, store verifiable credentials and support authentication and authorization processes. These components are critical for enabling trusted digital identities and ensuring compliance with regulatory requirements such as eIDAS[3].

The execution layer is built on a private Layer 2 blockchain network that perform computationally intensive operations off the main chain before submitting cryptographic proofs to the public Layer 1 blockchain. In this work, the layer leverages open-source zkEVM technologies such as Linea[4], which batch transactions and generate ZKPs to ensure correctness while preserving data privacy. Standardized interfaces (e.g., RPC endpoints) enable seamless interaction with the L1 smart contracts and interoperability with EVM-compatible networks.

Smart contracts encode the rules governing BCU issuance, transfer, allocation, claim, and retirement. They enforce eligibility criteria, validate transaction inputs, and ensure correct ownership transfers. In addition, they prevent double counting by guaranteeing that each BCU can only be claimed and retired once. Smart contracts also verify the ZKPs generated by the execution layer and produce PoI, ensuring that all transactions adhere to system rules and maintaining the integrity of the chain of custody.

4 Industry Impact and Use Cases

The proposed solution enables scalable deployment of Book & Claim across freight transportation networks, addressing key barriers related to cost, access and fragmentation while aligning with PI principles of interoperability and network-wide optimization.

[3] eIDAS, https://digital-strategy.ec.europa.eu/en/policies/discover-eidas
[4] Linea, https://linea.build/

For **shippers**, the system provides a credible and auditable mechanism to report Scope 3 emissions reductions (GHG Protocol, 2011), even in the absence of direct access to low-emission transport solutions. By enabling the allocation of verified BCUs, shippers can decarbonize their logistics footprint while supporting market demand for sustainable fuels.

For **LSPs**, the platform facilitates aggregation and allocation of LETS across multiple shipments and customers, enabling more efficient coordination in fragmented supply chains. It also provides a structured mechanism to ensure traceability and compliance across subcontracted carriers, addressing a key limitation in current Book & Claim implementations.

For **carriers**, the system creates new revenue streams by enabling the monetization of investments in low-emission fuels and technologies. Through integration with operational systems such as fuel card platforms and transport management systems, carriers can generate BCUs directly from verified physical activity, reducing administrative burden and increasing trust.

For **solution providers and fuel suppliers**, the platform enhances visibility and traceability of sustainable fuel usage, enabling more efficient market participation and supporting the scaling of low-emission fuel supply chains. Their role becomes central in providing certified data (PoS) and enabling the linkage between upstream certification and downstream usage.

A representative use case is the deployment of HVO in road freight. A carrier such as Gruber Logistics, refuels at an approved station using a fleet-linked fuel card, generating a verified fuel transaction (PoU). This transaction is combined with certified sustainability data (PoS) to issue BCUs, which are then allocated through LSPs to shippers seeking emissions reductions. The process ensures that environmental benefits are both physically verified and digitally traceable.

5 Discussion and Challenges

Despite its potential, several challenges must be addressed for large-scale adoption. First, ensuring reliable physical activity verification across fragmented supply chains remains difficult, particularly when LSPs rely on subcontracted carriers with varying levels of digital maturity. Current industry practices often limit Book & Claim implementation to controlled fleets, highlighting the need for standardized and automated mechanisms, such as fuel card integration and TMS connectivity to enable broader participation.

Second, interoperability across registries and stakeholders represents a critical barrier. The coexistence of multiple certification schemes, fuel registries and digital platforms creates fragmentation and limits scalability. Addressing this “upstream partner problem” requires either interoperable registry architectures or coordinated governance frameworks, as well as a more active role for fuel providers in ensuring consistent and verifiable data exchange.

Third, data quality and auditability remain essential concerns. While digital technologies such as blockchain enhance traceability, the integrity of the system ultimately depends on the reliability of input data, including fuel transactions and operational records. The role of independent verification bodies remains critical to ensure auditability across actors and to validate both digital transactions and underlying physical activity.

The proposed approach addresses these challenges through the integration of a three-tier verification model - combining Proof of Sustainability (PoS), Proof of Use (PoU) and Proof of Integrity (PoI) - but further work is required to validate scalability, standardization and cross-platform interoperability in real-world deployments.

6 Related Work

Existing Book & Claim implementations have emerged across different transport sectors, particularly in aviation and fuels. Certification schemes and registries developed by organizations such as the Roundtable on Sustainable Biomaterials (RSB) and the International Sustainability and Carbon Certification provide mechanisms for issuing and tracking sustainability attributes of fuels, forming the basis for Book & Claim systems. Similarly, Sustainable Aviation Fuel (SAF) registries[5] enable the allocation of emissions reductions across aviation supply chains, demonstrating the feasibility of decoupling physical fuel use from environmental claims.

The Smart Freight Centre (SFC) has developed the Market-Based Measures (MBM) Framework (SFC, 2023), which provides guidance for implementing Book & Claim in freight transport, including principles for accounting, allocation and reporting of emissions reductions. This framework builds upon established standards such as the GLEC Framework and the GHG Protocol, contributing to the harmonization of logistics emissions accounting.

Recent research has also explored the application of blockchain technologies in supply chains to enhance transparency, traceability, and trust. Studies highlight the potential of distributed ledgers and smart contracts to support decentralized data sharing and automated compliance in sustainability applications. However, many implementations remain conceptual or limited to traceability use cases, without fully addressing the requirements of Book & Claim systems in complex, multi-actor logistics networks.

Despite these advances, a gap remains in integrating certification systems, physical activity verification and digital transaction mechanisms into a unified, interoperable infrastructure. Existing approaches often lack robust linkage between sustainability attributes and real-world transport activity, as well as mechanisms to ensure compliance across fragmented supply chains. This paper addresses this gap by proposing a digital chain of custody platform that combines PoS, PoU, and PoI within a layered architecture aligned with PI principles, enabling scalable and trustworthy Book & Claim implementation in freight transportation.

7 Conclusions and Future Work

This paper presented a conceptual framework and system architecture for operationalizing Book & Claim in freight transportation through a digital chain of custody aligned with PI principles. By integrating Proof of Sustainability, Proof of Use, and Proof of Compliance, the proposed approach addresses key challenges related to trust, traceability, and scalability in the allocation of low-emission transportation services across fragmented supply chains.

The work contributes a structured data model, a layered system architecture, and a set of interoperable software components that together enable auditable, privacy-preserving, and standards-aligned emissions accounting and transaction management. In doing so, it bridges the gap between certification systems, real-world transport activity, and digital enforcement mechanisms.

This paper is intentionally positioned as a concept paper, laying the foundations for future implementation. The proposed approach will be materialized within the IKIGAI project through real-world use cases in freight decarbonization, involving multiple stakeholders across the transportation value chain. Future work will focus on system deployment, validation of the three-tier verification model in operational settings, and assessment of scalability, interoperability, and user adoption.

[5] SAFc Registry, https://safcregistry.org/

References

- Ballot, E., Montreuil, B. and Zacharia, Z.G. (2021), Physical Internet: First results and next challenges. J Bus Logist, 42: 101-107. https://doi.org/10.1111/jbl.12268
- eIDAS 2.0, European Union (2024) 'Regulation (EU) 2024/1183 amending Regulation (EU) No 910/2014 as regards establishing the European Digital Identity Framework', Official Journal of the European Union, 30 April, L 2024/1183
- Greenhouse Gas Protocol (2011), Corporate Value Chain (Scope 3) Accounting and Reporting Standard. World Resources Institute (WRI) and World Business Council for Sustainable Development (WBCSD) https://ghgprotocol.org/standards/scope-3-standard
- IEA (2024), Breakthrough Agenda Report 2024, IEA, Paris https://www.iea.org/reports/breakthrough-agenda-report-2024
- ISCC (2025), Strengthening RED Certification for Waste- and Residue-Based Biofuels, https://www.iscc-system.org/wp-content/uploads/2025/06/ISCC_PostionPaper_20062025.pdf
- ITF (2026), Claiming the Future: Can Book and Claim Support Low-Emission Transport?, OECD Publishing, Paris, https://www.itf-oecd.org/sites/default/files/docs/book-and-claim-low-emission-transport.pdf
- Kouhizadeh M., Saberi S. and Sarkis J., Blockchain technology and the sustainable supply chain: Theoretically exploring adoption barriers, International Journal of Production Economics, Volume 231, 2021, 107831, ISSN 0925-5273, https://doi.org/10.1016/j.ijpe.2020.107831
- Montreuil B., Toward a Physical Internet: meeting the global logistics sustainability grand challenge. Logist. Res. 3, 71–87 (2011). https://doi.org/10.1007/s12159-011-0045-x
- Niavis H. and Zafeiropoulou M. A., A blockchain-based architecture and smart contracts for an interoperable Physical Internet, Transportation Research Procedia, Volume 72, 2023, Pages 2022-2029, ISSN 2352-1465, https://doi.org/10.1016/j.trpro.2023.11.684
- Pan S., Ballot E. and Huang Q G.. Operations with physical internet. Reference Module in Social Sciences, Elsevier, 2025, 978-0-443-15785-1. 10.1016/b978-0-443-28993-4.00171-2
- Saberi S., Kouhizadeh M., Sarkis J. and Shen L. (2019), Blockchain technology and its relationships to sustainable supply chain management, International Journal of Production Research, 57(7), pp. 2117–2135. doi: 10.1080/00207543.2018.1533261
- Smith D. and Lewis A., Smart Freight Centre (SFC), 2023, Voluntary Market Based Measures Framework for Logistics Emissions Accounting and Reporting, https://smart-freight-centre-media.s3.amazonaws.com/documents/SFC_MBM_FRAMEWORK_2023_Oct.pdf
- Smart Freight Centre (SFC), 2025, Empowering Shippers: How Book and Claim Plus a Clear Action Plan Can Unlock Heavy Road Freight Decarbonization, https://smart-freight-centre-media.s3.amazonaws.com/documents/Empowering_Shippers_White_Paper_December_2025_.pdf

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Strategic infrastructure design for a Hyperconnected Urban Delivery System integrating public transport.

Cécile DUPOUY[1,2], François CLAUTIAUX[2], Walid KLIBI[1] and Olivier LABARTHE[1]
1. The Centre of Excellence for Supply Chain Innovation & Transportation (CESIT), Kedge Business School, Bordeaux, France
2. Université de Bordeaux, Institut de Mathématiques de Bordeaux, UMR CNRS 5251, Centre Inria de l'Université de Bordeaux, France
Corresponding author: cecile.dupouy@kedgebs.com

Abstract: *This paper addresses the strategic design of a hyperconnected urban delivery system integrating public transport. We consider a three-tier urban delivery system in which freight circulates through the middle-mile using existing public transport services including tramway and bus lines, with containerized parcels handled at logistics hubs deployed at public transport stops. We propose a path-based mixed-integer programming model that jointly optimizes public transport lines selection and the logistics hubs deployment, while assigning freight demand to either public transport paths or private vehicles. Public transport capacity is assumed to be exogenous, and feasible paths are pre-generated to capture routing with transfers across the network. The model is applied to a real case study based on the city of Bordeaux, France. Results show that increasing the number of available lines raises the share of demand served by public transport from 47.3% to 100%, highlighting the role of network connectivity in enabling urban freight flow within public transportation.*

Keywords: *Physical Internet; Hyperconnected City Logistics; Freight Transportation; Public transportation integration; Network Design.*
Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☒ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*
Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

1 Introduction

The Physical Internet (PI) proposes a new way of organizing logistics systems by promoting open and interconnected networks. Its main objective is to improve efficiency and sustainability through the sharing of infrastructures, resources, and services between stakeholders. Instead of operating in separate and isolated freight networks, the PI encourages the use of modular and compatible systems that allow flows to circulate more efficiently *(Montreuil, 2011)*. These benefits are supported by earlier work on logistics pooling, which shows that sharing transport resources across supply chains can significantly reduce emissions by improving vehicle utilization and coordination between actors *(Pan et al., 2013)*.
In urban contexts, the PI vision is reflected in the concept of Hyperconnected City Logistics, which focuses on better coordination between freight transportation and passenger mobility *(Crainic et al., 2023)*. In this context, the integration of freight flows into urban public transport systems has gained increasing attention. Public transport networks already provide dense spatial coverage and operate on regular schedules. When unused capacity is available and freight operations remain compatible with passenger services, using part of this existing system to

transport goods can help reduce congestion, emissions, the use of dedicated freight vehicles, and the duplication of urban logistics infrastructures (*Savelsbergh et al. 2016)*.
Strategic network design in the PI context has received growing attention, from early conceptual frameworks for multi-tier urban hub networks *(Montreuil et al., 2018)* to optimization-based contributions addressing hub location in large-scale hyperconnected logistics networks *(Kulkarni et al., 2022)*. Recently, *Liu et al. (2025*) propose a methodology for the design and capacity management of hyperconnected urban logistics networks, combining multi-tier hub network design with dynamic capacity deployment. However, these works focus on dedicated freight infrastructure and do not capture the specific constraints that arise when freight is integrated into existing public transport systems. In such systems, the design problem is fundamentally different: it requires jointly deciding which public transport lines are used for freight and where hubs should be deployed within the network, under passenger service constraints. While hub location problems have been extensively studied, they are mostly addressed in large-scale or dedicated logistics networks, and do not account for the structure of urban public transport networks. In particular, the role of transfer locations in shaping feasible routing options remains largely overlooked, despite its strong impact on system performance *(Azcuy et al., 2021)*.
In this paper, we address this gap with the following contributions:

- We propose a path-based mixed-integer programming model for the strategic design of a hyperconnected urban delivery system integrating public transport, jointly optimizing line selection and hub deployment under capacity constraints.
- We explicitly model transfer-enabled routing, capturing how the location and type of hubs at interchange stops shape the feasibility and cost of public transport paths.
- We quantify, on a real case study based on the Bordeaux network, how network connectivity and hub deployment patterns drive the ability of the system to absorb urban freight flows.

2 Problem setting

We consider an urban delivery system structured around three interdependent tiers. The urban first-mile connects local logistics facilities to access points located near the public transport network. The urban middle-mile corresponds to the movement of consolidated freight flows within the city using public transport services. The urban last-mile ensures final delivery to customers using light vehicles. These three tiers are connected, and their coordination is necessary to ensure an efficient urban delivery process (*Delle Donne et al., 2025*). The middle-mile plays a central role, as it enables freight to circulate across the city by relying on existing mobility infrastructures. The three-tier urban delivery system is represented in *Figure 1*.

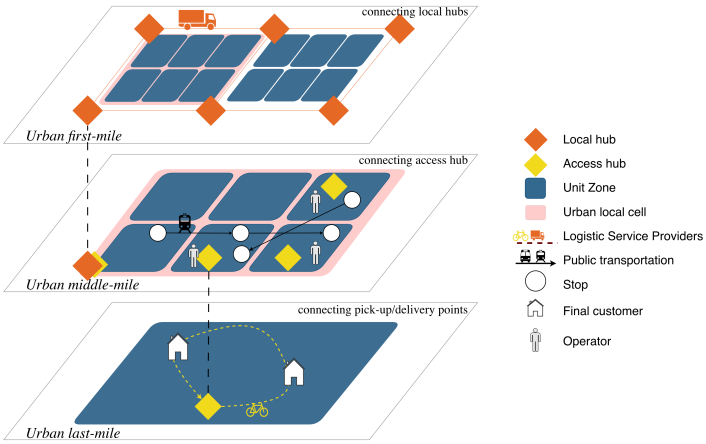


Figure 1. Three-tier urban delivery network (Dupouy et al. 2025).

To operate within public transport systems, parcels are grouped into standardized containers. Containerization is a core principle of the PI, as it enables freight flows to circulate across networks, transport modes, and multiple stakeholders using compatible and modular units. Consolidating parcels into containers reduces handling effort and improves efficiency in parcel logistics *(Kaboudvand, 2023).* We assume homogeneous container sizes, such that freight volumes and line capacities are expressed in a common unit. This standardization is essential to ensure interoperability across the different layers of the system.

This paper focuses on the design of the urban middle-mile infrastructure, where freight is integrated into the existing public transport network. The objective is to determine which lines to activate for freight, and where to deploy logistics hubs in the network. Freight can be routed through direct paths using a single line, or via paths involving transfers between lines. These operations are only possible at selected stops equipped with a hub where containers can be loaded, unloaded or transferred (*Labarthe et al., 2024*). We consider two types of hubs: access hubs, designed for handling smaller volumes, and local hubs, which handle larger volumes and act as main interfaces with first- and last-mile operations. In hyperconnected urban logistics, hubs are facilities that temporarily hold parcels and enable transfers between different modes and tiers of the network (*Faugère et al., 2022*). To structure these operations, the urban area is divided into unit zones. Each zone contains one or more public transport stops and may host logistics hubs. We consider a demand as a volume of parcel with an origin and destination zone. In the context of public transportation integration, freight operations must remain compatible with passenger services *(Li et al.,2021)*. If a demand pair cannot be served via public transport, either because no feasible public transport path exists, or because line or hub capacity is insufficient, freight is assigned to a private vehicle. This fallback mode guarantees full demand coverage across the urban area but induces a significantly higher unit transportation cost. Given the multi-stakeholder nature of such systems, an orchestrator coordinates the use of shared infrastructures and public transport capacity to ensure consistent operation. The orchestrator may be a dedicated entity, a logistics service provider, a city authority, or a public transport operator (*Kwon et al., 2023).*

3 Strategic decisions:

In this work, strategic design decisions are taken at the level of the urban middle-mile, where freight is integrated into public transport services. However, these decisions do not only affect middle-mile operations. By defining how freight can circulate within the public transport network, they also impact the organization of the urban first-mile and the urban last-mile. As a result, strategic choices made at the middle-mile level condition the overall performance and coordination of the three-tier urban delivery system. The strategic design of a hyperconnected urban delivery system integrating public transport relies on two main decisions. The first concerns the selection of public transport lines to be used for freight transportation. The second concerns the infrastructure design: the location and the size of logistics hubs deployed.

3.1 Line selection

The selection of public transport lines determines the connectivity of the urban middle-mile. Integrating a limited number of tramway or bus lines restricts the spatial coverage of the system and reduces routing and consolidation opportunities. As additional lines are integrated, the connectivity of the network increases. In graph terms, this expands the set of reachable origin-destination pairs and increases the number of feasible paths, as well as transfer opportunities between lines. Containers can change lines, parcels can be routed through alternative paths, and a larger share of demand can be served using public transport. However, increasing the number

of lines also raises coordination requirements, as transfers between lines require handling capacity at hubs.



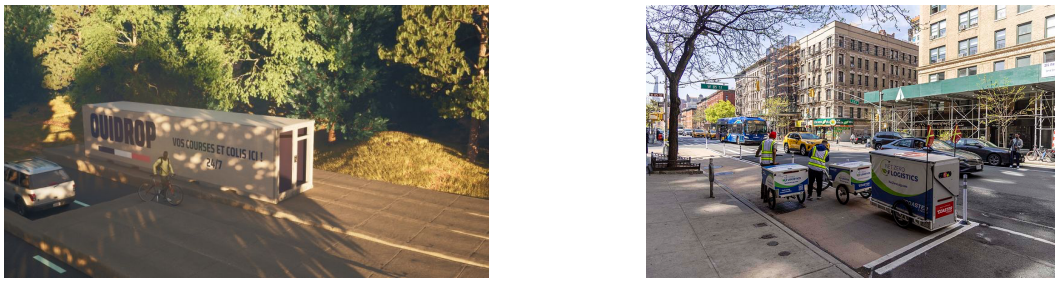
2.a) Full public transportation network *2.b) Subset of selected lines*

Figure 2. lines selection decision - Bordeaux example

Figure 2 illustrates the line selection decision on the Bordeaux (France) public transport network. While the full network (*Figure 2.a)*) offers extensive spatial coverage, only a subset of lines is selected for freight transportation (*Figure 2.b*)), defining the corridors along which containerized parcels can circulate. Increasing the number of selected lines expands network connectivity and allows a larger proportion of demand to be served by public transport.

3.2 Infrastructure design

The second strategic decision concerns the configuration of infrastructures within the public transport network. This decision includes two components: the location of hubs and the type of logistics hub deployed at each location. While some studies also optimize the spatial cover of zones *(Huang et al., 2018)*, we do not consider this aspect here. Zones are assumed to have a fixed coverage, and the strategic decisions of this paper are limited to hub location and infrastructure configuration. Each stop in the network is considered as a potential candidate for hub deployment. Freight operations are only allowed at stops where a hub is established. At these locations, containers can be loaded onto public transport vehicles, unloaded for last-mile delivery, or transferred between modes. Conversely, if no hub is deployed at a stop, it cannot support any freight operation, even if it belongs to an active line. Hub location directly determines which demand pairs can be served by public transport, as both origin and destination stop must be equipped with a hub. Moreover, hubs located at well-connected stops can act as transfer points, enabling flows to switch between lines and expanding routing possibilities.
In addition to their location, hubs may differ in terms of infrastructure configuration. Two hub types are considered: small hubs, which are less costly but have limited handling capacity, and large hubs, which allow higher consolidation opportunities at a higher cost. *Figure 3* illustrates the two hub types. *Figure 3.a)* shows a large hub (local hub) with dedicated storage space, enabling higher freight throughput and serving as an interface between the middle-mile and first-mile operations. *Figure 3.b)* shows a small access hub deployed directly in the urban environment, designed for limited volumes and last-mile parcel handling (New York City).



3.a) Ouidrop.fr *3.b) Micro-hub in New-York. Credit: NYC DOT*

Figure 3. Example of zone configuration

Hub deployment decisions are tightly coupled with line selection: a hub located at a stop not served by any selected line brings no operational value. Conversely, concentrating hubs at highly connected stops increases routing flexibility and reduces reliance on private vehicle deliveries. The joint optimization of these decisions defines a coupled network design and demand assignment problem

4 Mathematical model

We formulate the strategic infrastructure design problem as a Mixed-Integer Linear Program (MILP). The formulation follows a path-based approach modelling, which has been shown to be effective for capturing routing decisions with transfers in network design (*Zhu et al., 2025*). The model is defined over the following sets, parameters and variables:

Sets	$\mathcal{S}$	Set of stops
	$\mathcal{L}$	Set of public transport lines
	$\mathcal{L}(s)$	Set of lines passing through stop s
	$\mathcal{D}$	Set of demand pairs (origin, destination)
	$\mathcal{P}_d$	Set of public transport paths for demand pair d
	$\mathcal{L}(p)$	Set of lines used in path p
	$\mathcal{A}(p)$	Set of handling stops included in path p (entry, exit and transfer)
Parameters	v_d	Volume for demand d (in parcels)
	(o_d, e_d)	Origin and destination stop for demand d
	f_ℓ	Cost of activating line ℓ
	$(g^{\mathrm{B}}, g^{\mathrm{S}})$	Cost of opening a big and small hub in stop
	$(H^{\mathrm{B}}, H^{\mathrm{S}})$	Handling capacity of a big and small hub (in parcels)
	$c_{\mathrm{d}}^{\mathrm{PT}}$	Unit transportation cost with public transportation for demand d
	$c_{\mathrm{d}}^{\mathrm{PV}}$	Unit transportation cost with private vehicle for demand d
	cap_ℓ	Freight capacity available on line l (in parcels)
Variables	x_ℓ	1 if line ℓ is selected
	y_s^{B}	1 a large hub is opened at stop s
	y_s^{S}	1 a small hub is opened at stop s
	$z_{d,p}$	1 if demand d is served by public transport using path p
	w_d	1 if demand d is served by private vehicle

For each demand pair, the set of feasible public transport paths is generated by enumerating all possible paths between its origin and destination stops in the selected public transport network. These paths may be direct or may involve transfers between lines at connected stops. For each path we identify $\mathcal{A}(p)$ as the set of handling stops, including origin, destination, and transfer stops. The decision variables jointly capture infrastructure and demand assignment decisions. Binary variables x_ℓ, $y_s^{\mathrm{B}}$and y_s^{S} define the infrastructure design decisions, while variables $z_{d,p}$ and w_d represent demand assignment decisions. The model links these decisions by ensuring that a demand can only use a path if the corresponding lines are selected and the required hubs are activated at origin, destination, and transfer stops.

The model is written as follows:

$$\min \quad \sum_{\ell \in \mathcal{L}} f_\ell \cdot x_\ell + \sum_{s \in \mathcal{S}} (g^{\mathrm{B}} \cdot y_s^{\mathrm{B}} + g^{S} \cdot y_s^{\mathrm{S}}) + \sum_{d \in \mathcal{D}} v_d \cdot \left({c_d}^{\mathrm{PV}} \cdot w_d + \sum_{p \in \mathcal{P}_d} {c_d}^{\mathrm{PT}} \cdot z_{d,p} \right) \tag{1}$$

$$\sum_{p \in \mathcal{P}_d} z_{d,p} + w_d = 1 \qquad \forall d \in \mathcal{D} \tag{2}$$

$$\sum_{p \in \mathcal{P}_d} z_{d,p} \leq y^B_{o_d} + y^S_{o_d} \qquad \forall d \in \mathcal{D} \quad (3)$$

$$\sum_{p \in \mathcal{P}_d} z_{d,p} \leq y^B_{e_d} + y^S_{e_d} \qquad \forall d \in \mathcal{D} \quad (4)$$

$$y^{\mathrm{B}}_s + y^{\mathrm{S}}_s \leq 1 \qquad \forall s \in \mathcal{S} \quad (5)$$

$$z_{d,p} \leq x_\ell \qquad \forall d \in \mathcal{D}, \forall p \in \mathcal{P}_d, \forall \ell \in \mathcal{L}(p) \quad (6)$$

$$z_{d,p} \leq y^{\mathrm{B}}_s + y^{\mathrm{S}}_s \qquad \forall d \in \mathcal{D}, \forall p \in \mathcal{P}_d, \forall s \in \mathcal{A}(p) \quad (7)$$

$$\sum_{d \in \mathcal{D}} v_d \cdot \sum_{p \in \mathcal{P}_d : \ell \in \mathcal{L}(p)} z_{d,p} \leq cap_\ell \cdot x_\ell \qquad \forall \ell \in \mathcal{L} \quad (8)$$

$$\sum_{d \in \mathcal{D}} v_d \cdot \sum_{p \in \mathcal{P}_d : s \in \mathcal{A}(p)} z_{d,p} \leq H^{\mathrm{B}} \cdot y^{\mathrm{B}}_s + H^{\mathrm{S}} \cdot y^{\mathrm{S}}_s \qquad \forall s \in \mathcal{S} \quad (9)$$

$$x_\ell, y^{\mathrm{B}}_s, y^{\mathrm{S}}_s, z_{d,p}, w_d \in \{0,1\} \qquad \forall \ell \in \mathcal{L}, \forall s \in \mathcal{S}, \forall d \in \mathcal{D}, \forall p \in \mathcal{P}_d \quad (10)$$

Objective function (1) minimizes the total system cost, comprising three components: the fixed costs of selecting public transport lines, the cost of opening hubs (small or large) at stops, and the variable transportation costs incurred by serving demand pairs either by private vehicles or public transport. *Constraint (2)* ensures that each demand is served by exactly one option: either public transport path or private vehicle. *Constraints (3*) and *(4)* link demand assignment to hub deployment. A demand can be served by public transport only if hubs are activated at both its origin *(3)* and destination stops *(4)*, allowing freight to enter and exit the network. *Constraint (5)* ensures that at most one hub (small or large) can be opened at each stop. *Constraints (6)* and *(7)* ensure the feasibility of the selected path. *Constraint (6)* requires that all lines used by a path are activated. *Constraint (7)* ensures that if a demand path uses a transfer stop, a hub must be activated at that stop to enable the line change operation. *Constraints (8)* and *(9)* are capacity limitations. *Constraint (8)* limits the total volume of parcels routed through each line by its available transport capacity. The transport capacity is expressed as the product of the number of containers per passage, the line frequency and the container capacity. *Constraint* (*9*) restricts the handling capacity at hubs by accounting for all handling operations performed along a path, including entry, exit, and transfer operations. Finally, *Constraint (10)* defines the domain of all decision variables as binary.

5 Results

To evaluate the proposed model, we built a case study based on the city of Bordeaux, France using public transport data obtained from the open data platform provided by local authorities *(L'Atelier Opendata).* This dataset includes detailed information on stops, lines, and service schedules. Demand pairs are constructed from the spatial distribution of origins and destinations in the study area, while parcel volumes are synthetically generated to match a total demand of 30,000 parcels. For each demand pair, all feasible public transport paths between origin and destination stops are enumerated, including both direct paths and paths with transfers. We limit the number of transfers per path at two. The model is implemented and solved using IBM ILOG CPLEX 22.11. The maximum solving time is set to 3,600 seconds, with a target optimality gap of 0.0001.

Instances	Lines opened	Hubs opened		Demand pairs reached	Volume delivered by TP	Gap (%)	Time (s)
		Small	Large				
2 lines	2	14	1	162	47.3%	OPT	0.09
4 lines	4	33	4	762	60.3%	OPT	3.85
6 lines	6	44	5	1082	79.1%	OPT	96.4
10 lines	10	57	5	1532	96.6%	1.00	3600
15 lines	13	56	6	1676	100%	0.78	3600
20 lines	19	69	5	1902	100%	1.85	3600

Table 1: Model performance and solution.

Table 1 reports the results obtained for different network configurations. For the smaller instances (2, 4, and 6 lines), the solver reaches optimal solutions within the time limit. For the largest instances, the 1hour time limit is reached with a remaining optimality gap below 2%. From a system perspective, the results highlight two distinct effects of increasing the number of available lines. First, network coverage expands: the number of demands pairs whose origin and destination are both reachable via selected lines grows from 162 with 2 lines to 1,902 with 20 lines (column Demand pairs reached in *Table 1*). Second, among these reachable pairs, the volume of parcel effectively routed via public transport increases from 47.3% to 100%, reflecting the growing routing and consolidation opportunities as more transfer points become available. At the same time, the number of hubs increases with network size, reflecting the need for additional handling capacity to support higher volumes and more transfer opportunities. The results highlight a structural trade-off between infrastructure investment and transportation costs. Increasing network connectivity reduces reliance on private vehicles but requires deploying more hubs to enable transfers and maintain feasibility.

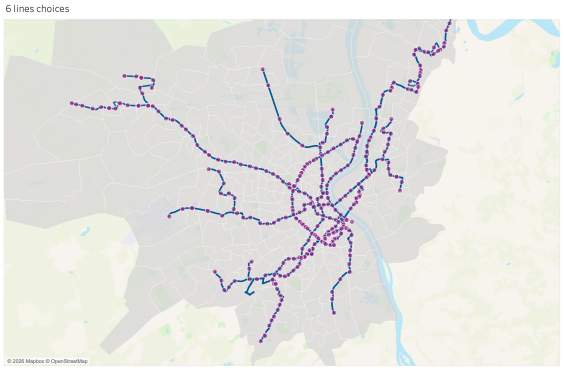
While these results show that full demand coverage can be achieved when the network is sufficiently connected, they are obtained for a fixed demand level (30,000 parcels). From a managerial perspective, an important question is to what extent such a mixed passenger-freight system can absorb higher demand volumes while maintaining a high level of service. To explore this, we consider an additional analysis based on the 20-line configuration, where demand is increased. The aim is to identify the capacity limits of the system and determine at which point public transport alone is no longer sufficient to absorb all flows, leading to a higher use of private vehicles.

Demand volume	Lines opened	Hubs opened		Demand pairs reached	Demand delivered by TP	Gap (%)	Time (s)
		Small	Large				
30 000	19	69	5	1902	100%	1.85	3600
40 000	19	65	7	1902	100%	4.19	3600
50 000	19	65	13	1902	98.2%	3.65	3600
60 000	19	68	13	1902	95.5%	2.19	3600
70 000	20	63	14	1902	92.8%	2.77	3600
100 000	20	75	19	1902	77.6%	2.42	3600

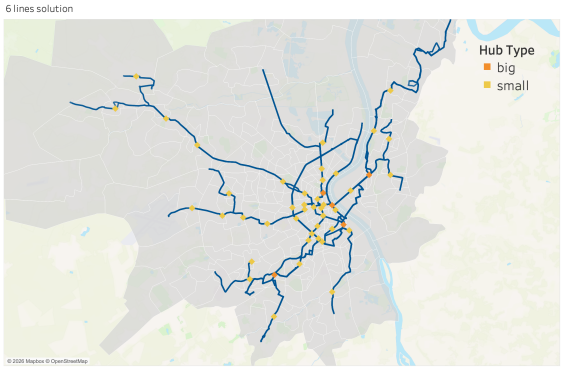
Table 2: Demand increases with the 20 lines scenario.

Table 2 reports the results of this sensitivity analysis on demand volumes using the 20-line configuration. As demand grows, hub deployment adapts: the number of large hubs increases from 5 to 19 while the number of selected lines remains almost stable up to 70,000 parcels. This indicates that the model first reinforces handling capacity at key transfer locations before extending the network. However, this adjustment is not sufficient to fully absorb higher demand levels, and part of the flow is shifted to private vehicles. The share of demand served by public transport decreases gradually from 100% to 77.6% at 100,000 parcels. At higher demand levels, the capacity is mainly constrained by line frequency and capacity. It is worth noting that even at 100,000 parcels, this still represents more than 70,000 parcels transported via public transport

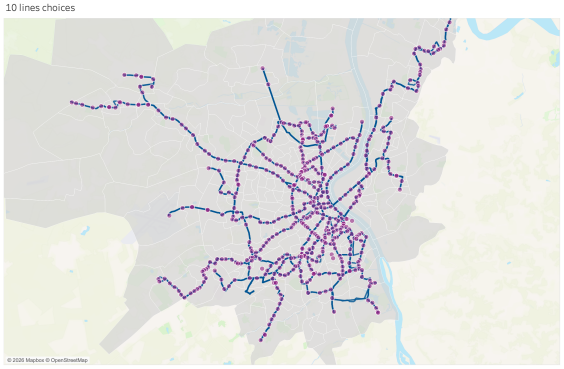
rather than private vehicles over a single day, which highlights the significant environmental potential of such systems.



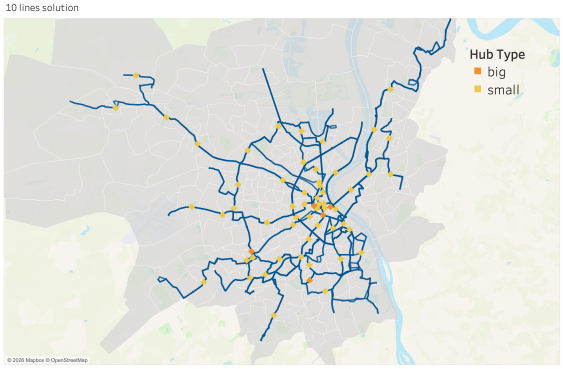
4.a) 6 lines choices



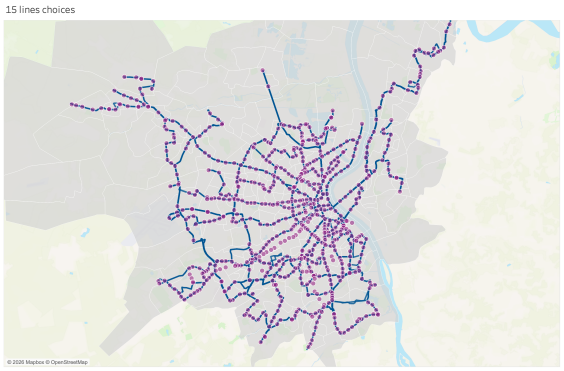
4.b) 6 lines solution



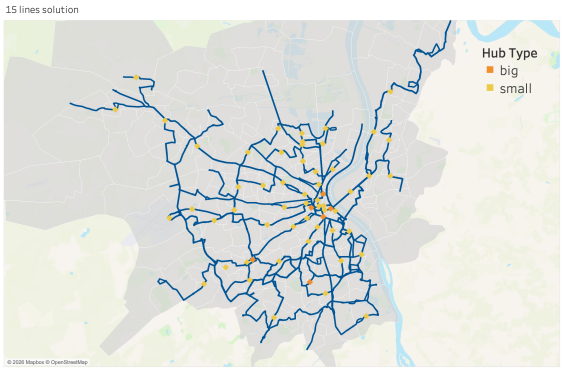
4.c) 10 lines choices



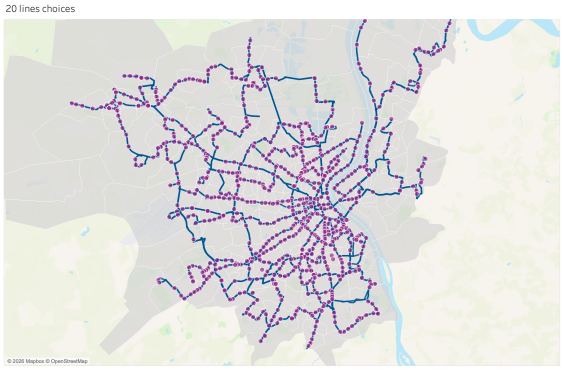
4.d) 10 lines solution



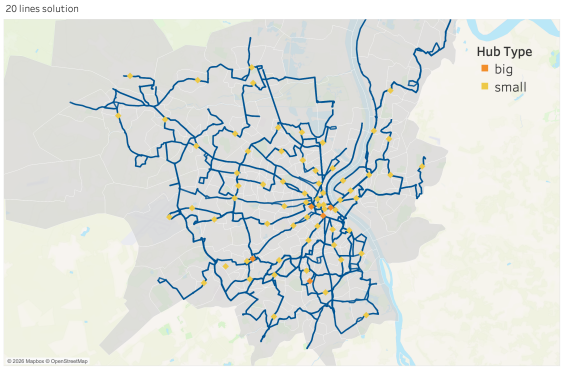
4.e) 15 lines choices



4.f) 15 lines solution



4.g) 20 lines choices



4.h) 20 lines solution

Figure 4. Hub location results for the 6-, 10-, 15-, and 20-line scenarios

Figure 4 illustrates the resulting network configurations for the 6, 10, 15 and 20-line scenarios. In each pair, the left figures (*a, c, e, g*) show the candidate network with all stops eligible for hub deployment and all lines available for selection, while the right figures (*b, d, f, h*) show the optimal solution returned by the model, with selected lines and activated hubs by size.
The spatial distribution of hubs provides additional insights. Hubs concentrate at well-connected stops, which act as key transfer points and support multiple lines. These locations maximize routing flexibility, enabling flows to be consolidated and redirected across the network. This structure is reflected in the deployment of large hubs, which are systematically located at the most connected stops. For instance, in the 15-line scenario, large hubs appear at nodes served by 3 to 6 lines. This confirms that the model prioritizes these locations to handle higher transfer volumes. In contrast, small hubs are more dispersed along the network corridors. Their role is mainly to ensure access to the system by covering origin and destination points, rather than supporting transfers. This results in a functional structure: large hubs support consolidation and transfers, while small hubs ensure spatial coverage.

Conclusion

This paper addresses the strategic design of a hyperconnected urban delivery system integrating public transport. We propose a path-based mixed-integer programming formulation that jointly optimizes line selection and hub deployment decisions, while explicitly accounting for routing with transfers and capacity constraints at both the line and hub levels.

The results obtained on a real case study based on the Bordeaux public transport network highlight the importance of connectivity in enabling the use of public transport for freight. Increasing the number of available lines from 2 to 15 raises the share of demand served by public transport from 47.3% to 100%. The model also reveals a structured hub deployment with large hubs located at highly connected stops where transfers are frequent, while small hubs are distributed along the network to ensure access to the system. From a managerial perspective, this highlights the importance of highly connected stops as priority locations for investment and that a limited set of well-chosen lines may be sufficient to achieve high coverage. The sensitivity analysis on demand shows that the system can absorb increasing volume before capacity constraints become limiting. Even at 100,000 parcels when capacity saturation occurs, 77.6% of demand is still served by public transport, representing more than 70,000 parcels transported via the public network rather than private vehicles. Mixed systems have the potential to reduce the use of dedicated freight vehicles in urban areas.

This work has several limitations that open directions for future research. Demand is assumed deterministic, and passenger-freight interactions are captured through aggregate capacity constraints. Future work could incorporate demand uncertainty and better represent interactions with passenger flows. In addition, feasible paths are generated through exhaustive enumeration, which may become challenging for larger networks and could be replaced by more scalable path-generation approaches. Finally, the model focuses on strategic decisions and does not explicitly account for tactical and operational planning. Extending the framework to include container routing, service scheduling, and real-time demand management would provide a more complete decision-support tool for hyperconnected urban delivery systems.

References

- Azcuy, I., Agatz, N., & Giesen, R. (2021). *Designing integrated urban delivery systems using public transport.* Transportation Research Part E: Logistics and Transportation Review, 156, 102525.
- Crainic, T. G., Klibi, W., & Montreuil, B. (2023). *Hyperconnected city logistics: a conceptual framework.* In Handbook on city logistics and urban freight (pp. 398-421). Edward Elgar Publishing.
- Delle Donne, D., Santini, A., & Archetti, C. (2025). *Integrating public transport in sustainable last-mile delivery: Column generation approaches*. European Journal of Operational Research.
- Dupouy, C., Clautiaux, F., Klibi, W., & Labarthe, O. (2025, June). *Multi-stakeholder Insights on Integrating Public Transport for Urban Freight Deliveries.* In IPIC 2025-11th International Physical Internet Conference. Georgia Institute of Technology.
- Faugère, L., Klibi, W., White III, C., & Montreuil, B. (2022). *Dynamic pooled capacity deployment for urban parcel logistics.* European Journal of Operational Research, 303(2), 650-667.
- Huang, Y., Savelsbergh, M., & Zhao, L. (2018). Designing logistics systems for home delivery in densely populated urban areas. *Transportation Research Part B: Methodological*, *115*, 95-125.
- Kulkarni, O., Dahan, M., & Montreuil, B. (2022). *Resilient hyperconnected parcel delivery network design under disruption risks.* International Journal of Production Economics, 251, 108499.
- Kwon, S., Montreuil, B., Dahan, M., and Klibi, W. (2023). *Hyperconnected logistic service networks: Bidding-based design framework*. In IPIC 2023-9th International Physical Internet Conference. Georgia Institute of Technology.
- Labarthe, O., Ahmadi, G., Klibi, W., Deschamps, J. C., & Montreuil, B. (2024*). A sustainable on-demand urban delivery service enabled by synchromodality and synergy in passenger and freight mobility.* Transportation Research Part C: Emerging Technologies, 161, 104544.
- Li, Z., Shalaby, A., Roorda, M. J., & Mao, B. (2021). *Urban rail service design for collaborative passenger and freight transport.* Transportation Research Part E: Logistics and Transportation Review, 147, 102205.
- Liu, X., Muthukrishnan, P., & Montreuil, B. (2025). *Network Design and Capacity Management in Hyperconnected Urban Logistic Networks*. In IPIC 2025-11th International Physical Internet Conference. Georgia Institute of Technology.
- L'atelier opendata Bordeaux: https://opendata.bordeaux-metropole.fr/explore/dataset/offres-de-services-bus-tramway-gtfs/export/
- Montreuil, B. (2011*). Toward a physical internet: Meeting the global logistics sustainability grand challenge*. Logistics Research, 3, 71–87.
- Montreuil, B., Buckley, S., Faugere, L., Khir, R., & Derhami, S. (2018). *Urban parcel logistics hub and network design: The impact of modularity and hyperconnectivity.*
- Pan, S., Ballot, E., & Fontane, F. (2013). *The reduction of greenhouse gas emissions from freight transport by pooling supply chains.* International journal of production economics, 143(1), 86-94.
- Savelsbergh, M., & Van Woensel, T. (2016). *50th anniversary invited article-city logistics: Challenges and opportunities.* Transportation science, 50(2), 579-590.
- Zhu, R., Ding, J. Y., Song, S., Ye, H., & You, K. (2025). *Express transport network design with transshipments*. Omega, 103397.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

TNO innovation for life

Design Principles for Human-Centered and Resilient Autonomous Logistics Systems

Frank Phillipson[1,2] and Ruben Fransen[1]

1. TNO, Den Haag, The Netherlands
2. Maastricht University, Maastricht, The Netherlands

Corresponding author: frank.phillipson@tno.nl

Abstract: *Global logistics systems are undergoing rapid transformation, driven by advances in digitalization, artificial intelligence, robotics, and networked sensing. These developments enable increasingly autonomous and self-organizing logistics architectures, yet the design of such systems is often approached narrowly as a technical optimization challenge. This paper proposes a conceptual framework of six design principles to guide the development of autonomous logistics systems that remain transparent, resilient, and socially responsible. The framework synthesizes insights from autonomous logistics research, socio-technical design, and critical theories of technological decentralization. Its relevance is demonstrated through analytical case applications to emerging Physical Internet infrastructures, including federated PI hubs and neutral logistics data spaces. The results highlight how autonomy can be operationalized in ways that support interoperability, multi-stakeholder governance, ecological stewardship, and long-term societal benefit. The framework offers a foundation for shaping future autonomous logistics ecosystems that are efficient and scalable while strengthening ethics, resilience, and sustainability.*

Keywords: *Physical Internet; autonomous logistics; self-organizing systems; socio-technical design; decentralization; sustainability; resilience; governance; logistics data spaces; federated PI hubs*

Physical Internet (PI) Roadmap Fitness*: System of Logistics Networks, Governance*

1 Introduction

1.1 Background: Toward Autonomous and Self-Organizing Logistics

Global logistics systems are undergoing a profound transformation, driven by advances in digitalization, artificial intelligence, robotics, and networked sensing (Gerrits et al., 2024). These technologies enable logistics infrastructures that can increasingly perceive, decide, and act without continuous human intervention. Such autonomous logistics systems, often described as self-organizing logistics (SOL), coordinate the movement, storage, and distribution of goods through decentralized, data-driven decision-making embedded in warehouses, vehicles, distribution networks, and information platforms. Instead of relying solely on centrally planned operations, autonomous logistics systems allow multiple actors and technical components to dynamically adjust routing, inventory management, scheduling, and resource allocation based on real-time information.

The concept of self-organizing logistics draws on developments in distributed computing, cyber-physical systems, and agent-based coordination. In these systems, the decision-making authority is (partially) delegated to autonomous agents, such as vehicles, robots, containers, or software agents that interact through digital infrastructures and collectively produce coordinated system behavior. Research in logistics and operations management has explored

such approaches using terms like autonomous control (Windt, 2010), agent-based logistics (Berndt, 2011), and self-organizing logistics networks (Daudi, 2020; Quak, 2018), highlighting the potential for improved flexibility, robustness, and responsiveness compared with traditional hierarchical planning systems. In dynamic environments characterized by demand uncertainty, supply chain disruptions, and increasing logistical complexity, decentralized coordination can allow systems to adapt in real-time while maintaining operational efficiency.

The growing relevance of autonomous logistics systems is closely tied to broader structural pressures on global supply chains. E-commerce expansion, demand volatility, and rising expectations for rapid delivery have intensified the need for logistics systems capable of continuous optimization and high responsiveness (Gajewska, 2018). At the same time, disruptions such as pandemics, geopolitical tensions, and climate-related events have revealed vulnerabilities in highly centralized and tightly coupled supply networks (Ivanov, 2020). Autonomous and self-organizing logistics architectures offer the potential to increase system resilience by distributing decision-making across multiple nodes and enabling adaptive responses to changing conditions (Windt, 2008). Robotics and automated vehicles can also reduce repetitive manual labor (Boysen, 2019), while AI-driven analytics enable more efficient routing, inventory allocation, and capacity utilization (Alii, 2025).

1.2 Existing Literature on Self-Organizing Logistics and Autonomous Systems

Research in self-organizing logistics has begun to address the architectural and coordination challenges associated with decentralized decision-making. Early work on the topic explored the application of principles from complex adaptive systems and swarm intelligence to logistics networks, emphasizing emergent coordination among autonomous agents (Svenson, 2004). Research has demonstrated that decentralized control can improve system robustness and responsiveness in environments characterized by uncertainty and rapid change (Shahmoradi-Moghadama, 2021). Within this literature, concepts such as autonomous control in logistics, agent-based transport systems, and cyber-physical logistics networks have provided theoretical and computational foundations for distributed coordination (Windt, 2010). These approaches typically rely on local decision rules, information sharing mechanisms, and adaptive feedback loops that enable system-level order to emerge from the interactions of multiple autonomous components, as illustrated in synchromodal transportation frameworks (De Juncker, 2017).

Parallel research streams address the design and governance of autonomous systems more broadly. Work in human–computer interaction and socio-technical systems design emphasizes transparency, interpretability, and human oversight in algorithmic decision systems (Shneiderman, 2022). Studies of distributed infrastructures and platform ecosystems highlight the risks of centralization and technological lock-in when digital infrastructures become dominated by a small number of actors (Srnicek, 2017). Meanwhile, research in sustainable logistics and circular supply chains (McKinnon, 2018) stresses the need to align logistics infrastructures with ecological constraints, emphasizing resource efficiency, reverse logistics, and lifecycle thinking. Together, these bodies of literature suggest that the design of autonomous logistics systems must integrate technical performance with considerations of governance, sustainability, and human interaction.

The challenge of designing autonomous systems, outside logistics, that remain accountable, adaptable, and socially embedded has also been examined in broader institutional and philosophical discussions of technology, for example Gorz (1982), Illich (1973), and Bookchin (1982). Theories of convivial and decentralized technologies emphasize the importance of systems that remain understandable and controllable by the communities that use them. These

perspectives argue that technological infrastructures should enhance human agency and collective decision-making rather than concentrating power or obscuring decision processes. Applying such perspectives to logistics raises important questions about transparency in algorithmic routing systems, participatory governance of data infrastructures, and the distribution of control across supply networks.

1.3 Problem Statement

Given the increasing autonomy of logistics infrastructures, there is a growing need for design frameworks that move beyond purely technical optimization and instead address the systemic implications of autonomous coordination. Designing autonomous logistics systems requires careful consideration of how decision-making authority is distributed, how systems interact with human operators and institutions, and how operational efficiency can be balanced with social and ecological objectives. In other words, the design challenge is not only how to build autonomous logistics technologies, but how to embed them within resilient, transparent, and socially responsible logistics ecosystems.

1.4 Contribution & Approach

This paper addresses this challenge by proposing a set of design principles for autonomous logistics systems that integrate technical effectiveness with human-centered, context-sensitive, and socially embedded considerations. Building on the 'heuristics for autonomous institutional design' proposed by Meacham (2025) and drawing conceptual inspiration from critical theories of technology and decentralization, the framework translates these ideas into the logistics domain. The resulting principles aim to guide the design of autonomous logistics architectures that remain transparent, distributed, ecologically compatible, temporally sustainable, and democratically governed. By articulating these principles, the paper contributes to the emerging discussion on how autonomous logistics systems can be designed not only for efficiency and scalability but also for resilience, accountability, and long-term societal benefit. While related to human-centered automation and responsible AI frameworks, the contribution of this paper lies in translating these concerns into a coherent design framework specifically for autonomous and Physical Internet-oriented logistics infrastructures. In particular, the framework integrates governance, interoperability, ecological constraints, and distributed coordination into a single socio-technical perspective tailored to autonomous logistics ecosystems.

The remainder of this paper is organized as follows. Section 2 outlines the methodological approach, describing the theory-building process through which the design principles were developed. Section 3 presents the six design principles in detail. Section 4 demonstrates their applicability through two analytical case examples: federated Physical Internet hubs and neutral logistics data spaces. Section 5 concludes by reflecting on the broader implications of the framework for the development and governance of autonomous logistics systems, with particular attention to sustainability, resilience, and ethical alignment.

2 Methodology

This work employs a qualitative, theory-building research approach. We conducted a structured conceptual synthesis drawing on three bodies of literature: (1) autonomous and self-organizing logistics systems, (2) socio-technical and human-centered design of autonomous systems, and (3) critical theories of technological decentralization, as shown in the previous section. The six design principles were derived through conceptual translation and adaptation of the heuristics for autonomous institutional design proposed by Meacham (2025), which themselves build on critical and decentralized technology traditions associated with Gorz (1982), Illich (1973), and Bookchin (1982). Through abductive reasoning, these heuristics were iteratively interpreted in

relation to autonomous logistics and Physical Internet research, resulting in six logistics-oriented principles addressing governance, interoperability, resilience, sustainability, and distributed coordination. This method aims to derive a coherent conceptual design framework rather than an empirical performance model.

The proposed design principles are illustrated through analytical case application. Two emerging Physical Internet infrastructures, federated PI hubs and neutral logistics data spaces, serve as analytical illustration cases. These examples demonstrate how the principles manifest in real-world or near-term logistics architectures, thereby assessing their relevance, coherence, and practical applicability. Future validation will include expert interviews with logistics practitioners and structured comparison with existing autonomous logistics design frameworks.

3 Design principles

Autonomous logistics systems are rapidly transforming the way goods and services are transported, stored, and distributed, offering unprecedented efficiency, responsiveness, and scalability. Advances in robotics, artificial intelligence, and networked sensing enable autonomous decision-making across warehouses, delivery fleets, and inventory systems. However, the design of such systems often focuses narrowly on technical optimization, automation, or cost reduction, frequently overlooking critical social, environmental, and ethical dimensions. This work proposes a set of design principles for autonomous logistics systems that integrate technical effectiveness with human-centered, context-sensitive, and socially embedded considerations. Within Physical Internet contexts, these principles support interoperable, federated, and multi-stakeholder logistics infrastructures such as shared PI hubs, cross-carrier routing systems, and neutral data spaces.

The first principle, **appropriability**, emphasizes that autonomous systems should be understandable, transparent, and adaptable for those who interact with them. Operators, planners, and local communities must be able to comprehend system logic, influence outcomes, and adapt processes to local circumstances. Decision-making algorithms must be explainable, and interfaces configurable, allowing users to align system operations with real-world conditions, regulations, and cultural norms. Appropriability ensures that technology remains a tool under human guidance rather than an opaque, uncontrollable entity.

The second principle, **non-monopolistic design**, addresses the risks of centralization. Autonomous logistics systems should encourage diversity, collaboration, and interoperability rather than consolidating control in a single algorithm, server, or company. Distributed architectures enable multiple nodes or actors to contribute to decision-making, share data safely, and complement each other's operations. This principle mitigates vulnerabilities inherent in centralized systems, such as single points of failure, data monopolies, or unequal power dynamics, and promotes resilient and cooperative networks.

Scale-appropriateness, the third principle, highlights the importance of aligning decision-making with the operational and environmental scale of the system. Not all logistics challenges require global optimization; local delivery routing, inventory replenishment, and distribution planning may be most effective when handled at the scale where the decisions have direct impact. By matching system decisions to the scale of their effects, organizations can optimize efficiency while remaining responsive to local contexts, infrastructure constraints, and dynamic environmental conditions.

Ecological Compatibility is the fourth principle. Autonomous logistics systems must operate in alignment with ecological limits rather than intensifying extractive and linear throughput models. System architectures, routing algorithms, warehousing strategies, and infrastructure deployment should minimize material and energy consumption, prioritize renewable energy integration, and reduce emissions across the lifecycle. Logistics flows should be designed to enable circular material cycles, including reverse logistics, repair, remanufacturing, reuse, and recycling, rather than reinforcing one-way extraction-to-disposal chains.

The fifth principle, **liberatory temporality**, refers to designing logistics systems with temporal buffers, recovery capacity, maintenance cycles, and planning horizons that support long-term resilience rather than continuous real-time optimization. Logistics operations should avoid extreme 'just-in-time' pressure or continuous optimization cycles that overextend human and material resources. Incorporating temporal buffers, seasonal rhythms, and cycles for maintenance and learning allows systems to adapt to disruptions, absorb variability, and operate sustainably. By embedding time structures that account for recovery and reflection, logistics systems can achieve efficiency without compromising the health of workers, infrastructure, or ecosystems.

Finally, **democratic embeddedness** emphasizes the integration of autonomous logistics systems within the social and institutional contexts in which they operate. Stakeholders, including drivers, warehouse staff, customers, and local authorities, should participate in the design, governance, and evaluation of systems. Participatory mechanisms such as cooperative data ownership, shared algorithmic decision-making, and feedback-driven adaptation ensure that systems reflect social values as well as operational objectives. Democratic embeddedness transforms logistics systems from isolated technological artifacts into accountable, responsive components of broader socio-technical ecosystems.

These six principles together offer a framework for designing autonomous logistics systems that balance technical performance with social responsibility, resilience, and adaptability. By foregrounding transparency, collaboration, context-sensitivity, sustainable temporal rhythms, and participatory governance, the framework addresses both the opportunities and the ethical challenges presented by automation and self-organization in logistics. The adoption of these design principles can guide engineers, planners, policymakers, and organizations in creating logistics infrastructures that serve operational goals while respecting human, environmental, and community values. This approach envisions autonomous logistics not as a replacement for human judgment but as a set of tools that enhance human capability, collective decision-making, and societal well-being.

Table 1: Operationalization of the six design principles in Physical Internet infrastructures.

Principle	PI Hubs	Data Spaces
Appropriability	Explainable routing interfaces	Traceable data usage
Non-monopolistic design	Open hub protocols	Federated data sovereignty
Scale-appropriateness	Multi-level coordination	Distributed computation
Ecological compatibility	Emissions-aware routing	Sustainability metrics
Liberatory temporality	Buffer windows	Shared planning horizons
Democratic embeddedness	Cooperative governance	Participatory data rules

4 Examples

Two emerging Physical Internet infrastructures illustrate how the proposed design principles can be operationalized in practice. Federated PI hubs represent shared logistics facilities where multiple carriers coordinate container flows through standardized modular infrastructure rather than proprietary networks. In parallel, neutral logistics data spaces provide federated digital infrastructures that allow companies to exchange operational data while retaining control over their information. Both approaches embody decentralized coordination, interoperability, and collaborative governance, demonstrating how autonomous logistics architectures can be designed to support transparency, resilience, and multi-stakeholder participation. Table 1 provides an overview of how the six design principles translate into concrete design requirements for two emerging Physical Internet infrastructures: federated PI hubs and neutral logistics data spaces. The table summarizes the operational implications of the framework before the principles are discussed in more detail for each case.

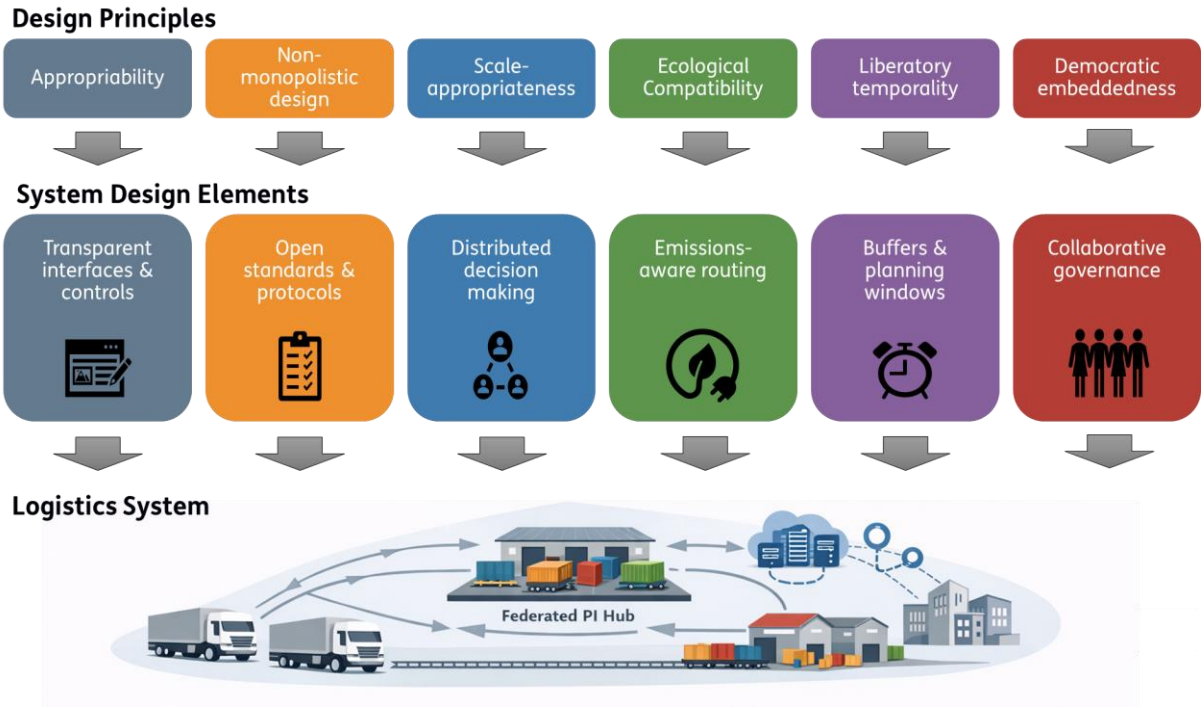


Figure 1: Design framework for autonomous logistics systems. The six design principles (top) inform system-level design choices (middle), which shape the behavior of distributed logistics infrastructures (bottom). The framework illustrates how normative principles translate into concrete architectural and operational decisions in self-organizing logistics systems.

4.1 Federated Physical Internet hubs

Federated Physical Internet (PI) hubs should be designed not merely as shared infrastructure, but as distributed coordination environments that operationalize autonomy, interoperability, and multi-actor governance. Rather than optimizing throughput within a single operator's network, the design challenge is to enable cross-organizational coordination without central control, while maintaining transparency, adaptability, and accountability.

Building on the Physical Internet vision proposed by Montreuil (2011), this requires rethinking both the technical architecture (e.g., routing logic, interfaces, standards) and the institutional design (e.g., governance, access rules, control rights) of logistics hubs.

Design requirements derived from the principles

Appropriability: Design for inspectability and intervention
Hub control systems should expose decision logic (e.g., routing, prioritization, allocation) through interpretable models and operator-facing interfaces. Local actors must be able to:

- inspect why specific routing or consolidation decisions are made
- override or adjust decision parameters in response to contextual conditions

This implies avoiding fully opaque optimization models in favor of explainable, parameterized decision systems.

Non-monopolistic design: Design for open access and protocol-based coordination
Federated hubs should be built on open technical standards and shared protocols rather than proprietary systems. Concretely:

- access rules should be non-discriminatory and standardized
- routing and handling interfaces should be interoperable across carriers
- no single actor should control core coordination algorithms

This shifts the hub from a platform model to a shared infrastructure model.

Scale-appropriateness: Design for multi-level decision autonomy
Control architectures should explicitly allocate decision-making across levels:

- local: docking, handling, short-range routing
- regional: consolidation, flow balancing
- network: long-haul coordination

This requires modular decision layers rather than a single optimization engine, enabling local responsiveness without losing system coherence.

Ecological compatibility: Design environmental objectives into core logic
Environmental constraints should be embedded directly into routing and scheduling mechanisms:

- include emissions, energy use, and modal options as decision variables
- enable consolidation across carriers to reduce empty miles
- integrate with low-carbon transport modes

This ensures sustainability is not an afterthought but a first-class optimization objective.

Liberatory temporality: Design for temporal flexibility, not continuous optimization
Instead of continuous real-time optimization, hubs should incorporate:

- configurable consolidation windows
- buffering mechanisms
- scheduled downtime and maintenance cycles

This requires temporal design parameters that balance efficiency with system resilience and human workload.

Democratic embeddedness: Design governance as part of the system architecture
Governance should be explicitly designed, not assumed. This includes:

- multi-stakeholder governance structures (e.g., cooperatives, public–private partnerships)
- transparent rule-setting for prioritization and access
- mechanisms for stakeholder feedback and adaptation

In effect, governance becomes a designed layer of the system, not an external constraint.

This approach is visualized in Figure 1.

4.2 Neutral / shared logistics data spaces

Neutral logistics data spaces should be designed as federated coordination infrastructures that enable autonomous decision-making while preserving data sovereignty and preventing platform centralization. The key challenge is to enable data-driven coordination without creating new asymmetries of control. Architectures such as IDS and Gaia-X (Daudi, 2024; Windolph, 2024) provide a starting point, but their effectiveness depends on how they are implemented in practice.

Design requirements derived from the principles

Appropriability: Design for data transparency and algorithmic traceability
Data spaces should provide:

- clear visibility[1] into what data is shared, with whom, and for what purpose
- traceability of which decisions were based on which data (e.g., routing outcomes)
- user-configurable access and usage policies

This requires auditable data flows and explainable decision pipelines, not just secure data exchange.

Non-monopolistic design: Design for federated data sovereignty
The architecture should ensure:

- data remains with the originating actor unless explicitly shared
- no central entity aggregates or controls all data
- interoperability is achieved through standards, not centralization

This implies federated identity, access control, and data exchange mechanisms.

Scale-appropriateness: Design for distributed computation
Instead of central data processing:

- computations (e.g., routing, forecasting) should occur where data resides
- only necessary insights or signals are shared across the network
- privacy-preserving coordination techniques (e.g., secure multi-party computation) can be used to enable joint optimization across actors without requiring raw data sharing

This supports context-sensitive decision-making and reduces dependency on central infrastructure.

Ecological compatibility: Design for sustainability-aware data use
Data-sharing mechanisms should explicitly support:

- emissions-aware routing
- visibility into capacity utilization
- coordination of reverse and circular flows

This requires integrating environmental metrics into shared data models and APIs.

Liberatory temporality: Design shared temporal frameworks
Data spaces should enable coordination not only in space but in time:

[1] Data spaces should provide controlled transparency, ensuring visibility into data-sharing relationships, access rights, and usage purposes, while preserving confidentiality of the underlying data through mechanisms such as usage control, encryption, and secure computation.

- shared planning horizons
- visibility into future capacity and demand
- mechanisms for smoothing peaks and distributing workloads

This moves systems away from reactive, short-term optimization toward coordinated temporal planning.

Democratic embeddedness: Design participatory data governance
Governance mechanisms should include:

- collective rule-setting for data access and usage
- representation of different stakeholder groups
- mechanisms for dispute resolution and accountability

This transforms the data space into a commons-like infrastructure, rather than a neutral technical layer.

5 Conclusion

The increasing autonomy of logistics infrastructures presents both significant opportunities and substantial socio-technical challenges. This paper argued that the design of autonomous logistics systems cannot be approached solely as a technical optimization problem but must instead be grounded in principles that support transparency, resilience, sustainability, and participatory governance. Building on abductive, theory-driven synthesis and conceptual inspiration from both logistics research and critical theories of technological decentralization, we developed six design principles, appropriability, non-monopolistic design, scale-appropriateness, ecological compatibility, liberatory temporality, and democratic embeddedness. Together, these principles provide a coherent framework for shaping autonomous logistics architectures that remain accountable and aligned with societal values.

The framework is illustrated through analytical application to federated PI hubs and neutral logistics data spaces. Specifically, the examples of federated PI hubs and neutral logistics data spaces, demonstrates that these principles are not merely theoretical but can be meaningfully operationalized in Physical Internet–oriented infrastructures. These cases illustrate how autonomously coordinated logistics can be designed to encourage interoperability, distributed decision-making, ecological stewardship, and multi-stakeholder governance. Future empirical validation will expand this work through expert interviews, comparative analysis with existing design frameworks, and applications in emerging PI deployments.

The potential impact of this work could be significant. As autonomous logistics capabilities expand, the sustainability, ethics, and resilience implications of automation will increasingly shape both industry practice and policy. The framework presented in this paper offers guidance for organizations and policymakers seeking to ensure that autonomous logistics systems enhance, not erode, human agency, environmental responsibility, and democratic oversight. In the context of accelerating Physical Internet adoption, these principles can support the development of logistics ecosystems that are not only efficient and scalable but also socially robust, transparent, and future-ready.

References

- Alii, H. M., Ugli, A. A. G., Mannonov, A., & Lalnunthari, L. (2025). Artificial intelligence and iot in logistics management for the intelligent supply chain management for the future era. In 2025 International Conference on Computational Innovations and Engineering Sustainability (ICCIES) (pp. 1-6). IEEE.

- Berndt, J. O. (2011): Self-organizing logistics process control: an agent-based approach. In International Conference on Agents and Artificial Intelligence (pp. 397-412). Springer Berlin Heidelberg.
- Bookchin, Murray (1982): The ecology of freedom.
- Boysen, N., De Koster, R., & Weidinger, F. (2019). Warehousing in the e-commerce era: A survey. European Journal of Operational Research, 277(2), 396-411.
- Daudi, M., & Thoben, K. D. (2020). Self-organizing logistics networks for Less-Than-Truckload. Procedia Manufacturing, 52, 101-106.
- De Juncker, M. A., Huizing, D., del Vecchyo, M. O., Phillipson, F., & Sangers, A. (2017). Framework of synchromodal transportation problems. In International Conference on Computational Logistics (pp. 383-403). Cham: Springer International Publishing.
- Diallo, M. A., Li, Z., Sun, J., Zhang, X., & Fu, D. (2024). Optimizing logistics operations: An International Data Space (IDS) inspired multi-sided data space marketplace for enhanced data sovereignty in the logistics industry. In Proceedings of the 3rd International Conference on Big Data Economy and Digital Management (BDEDM 2024)
- Gajewska, T., & Zimon, D. (2018). Study of the logistics factors that Influence the development of e-commerce services in the customer's opinion. Archives of Transport, 45(1), 25-34.
- Gerrits, B., van Heeswijk, W., & Mes, M. (2024): Towards self-organizing logistics in transportation: a literature review and typology. International Transactions in Operational Research, 31(3), 1309-1374.
- Gorz, André (1982): Farewell to the working class: an essay on post-industrial socialism. Pluto Press.
- Illich, Ivan, and Anne Lang (1973): Tools for conviviality.
- Ivanov, Dmitry. "Predicting the impacts of epidemic outbreaks on global supply chains: A simulation-based analysis on the coronavirus outbreak (COVID-19/SARS-CoV-2) case." Transportation Research Part E: Logistics and Transportation Review 136 (2020): 101922.
- McKinnon, A. (2018). Decarbonizing logistics: Distributing goods in a low carbon world. Kogan Page Publishers.
- Meacham, Darian (2025): Inaugural Speech: Institution Work and Technology, Maastricht University.
- Montreuil, B. (2011). Toward a Physical Internet: meeting the global logistics sustainability grand challenge. Logistics Research, 3(2), 71-87.
- Quak, H., van Kempen, E., & Hopman, M. (2018). Moving towards practical implementation of self-organizing logistics–making small steps in realizing the PI vision by raising awareness. In International Physical Internet Congress (pp. 106-119).
- Shahmoradi-Moghadama, Hani, and Jörn Schönberger. "A robust decentralized decision-making approach for mobile supply chains under uncertainty." Logistics Research 14.1 (2021): 1-15.
- Shneiderman, B. (2022). Human-centered AI. Oxford University Press.
- Srnicek, N. (2017). Platform capitalism. John Wiley & Sons.
- Svenson, P., Mårtenson, C., Sidenbladh, H., & Malm, M. (2004). Swarm intelligence for logistics: background. Swedish Defence Research Agency.
- Windolph, J., Pahl, M., & Kaup, S. (2024). Gaia-X as an Enabler to Shape Interconnected Logistics in Europe. IPIC 2024
- Windt, K., Becker, T., Jeken, O., & Gelessus, A. (2010). A classification pattern for autonomous control methods in logistics. Logistics Research, 2(2), 109-120.
- Windt, Katja, Felix Böse, and Thorsten Philipp. "Autonomy in production logistics: Identification, characterisation and application." Robotics and Computer-Integrated Manufacturing 24.4 (2008): 572-578.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Converged Digital–Physical Identity Resolution Architecture for Enabling Traceability in the Physical Internet

Wenqing Lei[1], Yong-Hong Kuo[1], George Q. Huang[2] and Ming Li[2]

1. Department of Data and Systems Engineering, The University of Hong Kong, Hong Kong SAR
2. Department of Industrial and Systems Engineering, The Hong Kong Polytechnic University, Hong Kong SAR

Corresponding author: ming.li@polyu.edu.hk

Keywords: *Physical Internet; Cyber–Physical Internet; PI Networks; Digital–Physical Synchronization; Address Resolution; Traceability; Interoperability.*

Conference Domain Fitness: *This contribution directly addresses the advancement of PI Networks, Logistics Data Platforms, and System-of-Logistics-Networks integration as highlighted in the IPIC 2026 call for contributions. By proposing a standardized digital–physical identity resolution architecture, the work supports interoperable logistics nodes, synchronized cyber–physical execution, and scalable cross-network coordination within emerging Physical Internet infrastructures. The framework contributes to foundational standardization efforts required for large-scale deployment and operational integration of the Physical Internet.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☒ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☒ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

Research Contribution Abstract

The Physical Internet (PI) envisions a globally interconnected logistics system built upon standardized modular containers and interoperable nodes. However, in emerging Cyber-Physical Internet（CPI) environments, logistics nodes must function as dual-layer entities: digital routing units on the network layer and physical operational hubs on the physical layer. This dual role creates a critical requirement for a synchronization mechanism that can map logical addresses to physical execution. To bridge this gap, this research formalizes three critical inquiries:

- How can traditional Internet addressing mechanisms be adapted to accommodate the unique dual layer demands of the CPI?
- How can resolution protocols be designed to maintain performance and scalability in dynamic, high-mobility logistics environments?

- How can we ensure the seamless migration of these identities as assets move across different network domains?

By addressing these questions, this work proposes a Converged Digital–Physical Identity Resolution Architecture that links Physical Internet Protocol (PIP) addresses with Physical Media Access Control (PMAC) identifiers and geographic coordinates. This approach directly aligns with Cyber-Physical Systems (CPS) principles by ensuring that cyber-layer routing decisions are instantly resolvable into physical-layer actions.

The framework establishes a structured relationship between logical and physical identifiers to ensure end-to-end traceability.

- **Physical Internet Protocol (PIP)**: As the logical identifier, the PIP address inherits the traditional IP structure, supporting both static and dynamic allocation through specialized DHCP services designed for CPI local area networks. This ensures compatibility with existing Internet infrastructures while enabling precise device tracking.

- **Physical Media Access Control (PMAC)**: To complement traditional ID-based standards like GS1, which focus primarily on asset identification, our architecture introduces two PMAC formats that bridge the gap between identity, real-time location, and network routing. PMAC-Fixed serves as a unique identifier for logistics node areas, such as warehouses or ports, by encoding geographic coordinates into transmission-friendly formats. In contrast, PMAC-Mobile ensures identity continuity for dynamic assets like trucks and shipment units. This integration allows the system to resolve 'Where' (Physical Location) directly from 'Who' (Network ID) within the protocol layer.

Algorithmic Encoding for Geographic Precision: A key contribution of this architecture is the transformation of geographic contour data into a standardized PMAC-Fixed address. The encoding process is realized through a two-stage algorithmic approach:

1. **Point Selection**: Utilizing a specialized algorithm (e.g., based on building outlines from OpenStreetMap API), the system selects a sorted list of coordinates (typically eight points) that represent the geometric shape of the logistics node.

2. **Normalization and Integration**: These coordinates are rounded to six decimal places, normalized to eliminate negative values, and combined with a Plus Code derived from the centroid of the building. This results in a 26-byte unique identifier that encapsulates both identity and precise spatial reference.

To facilitate real-time resolution, we developed the CPI-ARP protocol, which operates over the User Datagram Protocol (UDP) to minimize connection overhead. The detailed packet field specifications are illustrated in Table 1. The protocol maps PIP addresses PMAC identifiers through a request-response workflow:

- Request Phase: A device broadcasts a CPI-ARP request containing the target's PIP to retrieve its corresponding physical location.

- Response Phase: The target device or a centralized CPI-ARP server replies with the PMAC address, allowing the requester to update its local resolution cache.

Table 1

CPI-ARP Packet Structure

Field Name	Size(Bytes)	Description
Physical Type (PTYPE)	2	Specifies the type of physical address, e.g., Fixed is 1.
Network Type (NTYPE)	2	Identifies the network address type, e.g., the network address is 0x0800.
Physical Address Length (PLEN)	1	Indicates the length of the physical address (in bytes), e.g., the PMAC address is 26 bytes.
Network Address Length (NLEN)	1	Indicates the length of the physical address (in bytes), e.g., the PIP address is 4 bytes.
Operation Code (OPER)	2	Specifies the type of operation, 1 for request, 2 for reply.
Sender Physical Address (SPA)	26	The sender's physical address.
Sender Network Address (SNA)	4	The sender's network address.
Target Physical Address (TPA)	26	The physical address of the target (empty in request).
Target Network Address (TNA)	4	The target network address.

To evaluate the feasibility and scalability of the proposed architecture, we conducted demonstrative experiments in a controlled, isolated virtual network environment. This setup utilized Python-based scripts to simulate high-concurrency requests across two primary logistics scenarios:

- Scenario 1 (Intracity): Communication between logistics nodes within Hong Kong (Hong Kong Island to Kowloon), testing the system in a low-latency, high-stability urban environment.

- Scenario 2 (Cross-Border): Communication between nodes in Hong Kong and Guangzhou, China, representing the complexities of long-distance routing and varied network conditions within the Greater Bay Area.

The performance of the CPI-ARP protocol was assessed using four key Quality of Service (QoS) metrics: Latency, Jitter, Success Rate, and Packet Loss Rate.

- Network Stability: The Intracity (HK) network demonstrated superior stability, maintaining a low packet loss rate of approximately 0.5% even under high concurrent loads (50 threads). In contrast, the Cross-Border (CB) network exhibited a sharp increase in latency and a packet loss rate exceeding 6% as the system reached its processing limits.

- Scalability: While both scenarios maintained operational stability, the HK network showed better scalability with gradual latency increases, whereas the CB network encountered performance bottlenecks due to complex routing and longer transmission distances. The empirical relationship between throughput and latency under varying load conditions is illustrated in Figure 1.

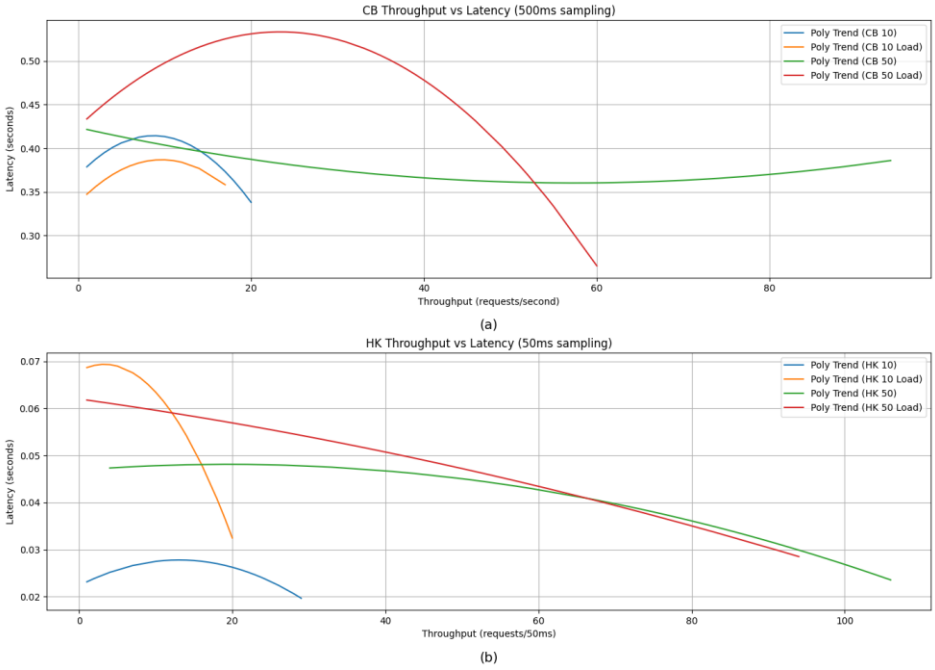


Figure 1. Plot of Throughput versus latency (polynomial fitting under different thread numbers and load conditions)

Recognizing the inherent vulnerabilities of traditional protocols, we evaluated a secure configuration of CPI-ARP employing AES-GCM encryption and Pre-Shared Key (PSK) authentication.

Attack Mitigation: The security mechanisms successfully reduced the success rate of spoofing, Man-in-the-Middle (MITM), and replay attacks from 100% to 0%.

Overhead Costs: Implementing these protections introduced a measurable computational overhead, increasing average latency by approximately 29.4% for low-load (10 threads) and 48.2% for high-load (50 threads) scenarios. This quantified trade-off confirms that while security measures are indispensable for resilient PI networks, they require careful optimization in resource-constrained environments.

This research introduces a converged digital–physical identity resolution architecture for the CPI environment, providing a standardized mechanism to link network-layer identities (PIP) with precise physical locations (PMAC). Our experimental results confirm that the proposed CPI-ARP protocol maintains high success rates and manageable latency in both intracity and cross-border logistics scenarios. The dual-layer addressing scheme effectively bridges the gap between cyber-layer routing and physical-layer execution, laying the groundwork for a more traceable and interoperable Physical Internet.

However, as the CPI is in its nascent stages, several areas remain for further exploration:

- Dynamic Security Mechanisms: Future work will focus on implementing dynamic key exchange and more robust resilience against Distributed Denial of Service (DDoS) and Sybil attacks in large-scale deployments.

- AI-Driven Address Management: Integrating artificial intelligence (AI) and machine learning (ML) to optimize address allocation and routing based on task priority and device mobility.

- Real-World Validation: Transitioning from virtual simulations to large-scale pilot projects in global logistics networks to assess performance under real-world noise and interference.

References

Ballot, E. (2019). The Physical Internet. In H. Zijm, M. Klumpp, A. Regattieri, & S. Heragu (Eds.), *Operations, Logistics and Supply Chain Management* (pp. 719–734). Springer International Publishing. https://doi.org/10.1007/978-3-319-92447-2_31

Hercog, D. (2020). ARP Protocol. In D. Hercog, *Communication Protocols* (pp. 321–322). Springer International Publishing. https://doi.org/10.1007/978-3-030-50405-2_19

Lei, W., Li, M., Kuo, Y.-H., & Huang, G. Q. (2025). Converged Address Resolution Protocol for Traceability and Visibility in Cyber–Physical Internet. *IEEE Internet of Things Journal*, *12*(14), 27322–27338. https://doi.org/10.1109/JIOT.2025.3562803

Montreuil, B., Ballot, E., & Fontane, F. (2012). An Open Logistics Interconnection model for the Physical Internet. *IFAC Proceedings Volumes*, *45*(6), 327–332. https://doi.org/10.3182/20120523-3-RO-2023.00385

Tran-Dang, H., Krommenacker, N., Charpentier, P., & Kim, D.-S. (2020). Toward the Internet of Things for Physical Internet: Perspectives and Challenges. *IEEE Internet of Things Journal*, *7*(6), 4711–4736. https://doi.org/10.1109/JIOT.2020.2971736

Wu, H., Huang, L., Li, M., & Huang, G. Q. (2024). Cyber-Physical Internet (CPI)-enabled logistics infrastructure integration framework in the greater bay area. *Advanced Engineering Informatics*, *60*, 102551. https://doi.org/10.1016/j.aei.2024.102551

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE
Universidad de La Sabana

Assessing the Role of Firm Size Distribution and Hub Localization in Collaborative Supply Chain Design under the PI Paradigm

Luz H. Mancera-Méndez[1], William J. Guerrero[1], and Jairo R. Montoya-Torres[1]
School of Engineering, Universidad de La Sabana, Campus del Puente del común, Km. 7, Autopista Norte de Chía, Cundinamarca, Colombia
Corresponding author: william.guerrero1@unisabana.edu.co

Abstract: *This research addresses Supply Chain Network Design (SCND) within the Physical Internet framework, focusing on resource pooling and the shared utilization of transport assets and PI-hubs. We propose a Mixed-Integer Linear Programming (MILP) model to optimize strategic flow allocation and facility positioning. A core contribution is the integration of an "individual rationality" constraint, ensuring a "win-win" scenario where operational costs for all participants remain lower than or equal to traditional non-collaborative baselines. Validated through diverse computational experiments involving SMEs and large enterprises and different hub locations, the results demonstrate significant improvements in both individual cost reduction and hub capacity utilization.*

Keywords: *Supply Chain Network Design, Network of Networks, Asymmetric Horizontal Collaboration, Physical Internet*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☒ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☐ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

1 Introduction

Supply chain management is one of the most significant factors in a country's economic development, as it reduces costs, expands market access, boosts employment, encourages investment, and increases business productivity (Calatayud & Montes, 2021). It is not a coincidence that higher-income countries score highest on the Logistics Performance Index.

This has boosted research aimed at designing and operating increasingly efficient networks. In this context, the Physical Internet (PI) emerges as an alternative for this sustainable design by proposing a global logistics system based on the interconnection of logistics networks, enabling better use of resources at all levels and a lower environmental impact (Montreuil et al., 2012). Studies show the potential of the PI to improve supply chain performance in terms of costs (Sarraj et al., 2014); (Kim et al., 2021), environmental impact (Sarraj et al., 2014), (Ariztegui-Beltrán et al., 2024) and responsiveness (Kulkarni et al., 2022); (Yang et al., 2017); however, the studies reviewed have examined the benefits of PI when supply chains consist of firms of similar sizes, but given the heterogeneity of firms within a supply network, it is worthwhile to study the benefits that firms of various sizes may derive, as well as the potential effect of distribution center location. In this regard, it is important to determine whether there is a

relationship between company size and the cost savings they achieve; whether the competitive advantages of companies within the network are shared—for example, due to the location of their distribution centers or lower operating costs—and whether synergies are generated by considering the transformation of traditional distribution centers into π-Hubs where multiple companies can share operations.

With the aim of contributing to this effort, this article presents assessing the performance of a supply chain designed under the Physical Internet paradigm, considering small, medium, and large companies that share their hubs, in line with the principle of open nodes proposed by the PI as the basis for so-called π-Hubs, where horizontal interconnectivity enables the consolidation of multi-companies flows. These hubs can be located near production facilities or near consumption centers. Likewise, it is considered the shared use of vehicles is permitted, seeking to optimize transport and storage capacity in collaborative competition environments.

This article is organized as follows: The second section presents the literature review. The third section gives an overview of the mixed-integer mathematical programming model. The fourth section presents the experimental setting and the discussion of results. The last section concludes the paper with key findings and potential avenues for future research.

2 Literature Review

Although the study of supply chains has been ongoing for decades, problems persist, such as underutilization of vehicles, low use of multimodal transport, poor working conditions for drivers, inefficiencies in product distribution, inefficient use of production and storage facilities, poor coordination within distribution networks, and low network security and resilience, among others (Montreuil, 2011). This is partly due to the existence of independent, heterogeneous logistics networks dedicated to a single company or a small group of collaborating companies (Yang et al., 2017), which sometimes makes the decision-making process slow and fragmented; therefore, a design based on PI could help overcome these situations.

According to the literature review conducted, early research on this topic focused on the proposal of standardized containers, known as π-containers, and their routing. In this regard, (Meller et al., 2012) proposed mathematical models that integrate the use of standardized containers into network design, while in 2014 (Sarraj et al., 2014) evaluated the implementation of asynchronous shipments in smart modular containers of different sizes using multimodal transport. In 2016, (Sallez et al., 2016) presented a proposal outlining the characteristics and functionalities these containers should possess, while more recently, (Hayek et al., 2023) developed a π-container for the timber supply chain in the European Union.

Multimodal transportation was subsequently incorporated into network design, with studies that allow for the shared use of resources and various transportation options (Furtado & Frayret, 2015),utilize railways (Walha et al., 2016), and include waterway transportation (Ariztegui-Beltrán et al., 2024).

Regarding network design that considers multiple objectives, there are studies that take into account economic aspects and network resilience (Tordecilla et al., 2025) as well as environmental aspects (Ji et al., 2023).

Previous studies conclude that designing a network based on PI principles reduces network costs while improving its resilience and its social and environmental indicators. These benefits have drawn the attention of the public sector, leading the European Union and Japan to develop

roadmaps for their implementation. Specifically, the European Union has established five implementation stages, with the initial stages focusing on integrating nodes and connecting previously separate networks (Alliance for Logistics Innovation through Collaboration in Europe, 2019). This can serve as a starting point for leveraging existing infrastructure and building confidence in the operation of this integrated model. However, when connecting independent networks, the characteristics of their stakeholders must be considered. The studies identified in the literature review analyze the overall performance of the network, focusing primarily on economic and environmental dimensions, and in some cases social benefits; however, no study has been found that analyzes the benefits for companies of different sizes of participating in a network designed according to IP principles. This paper aims to propose a protocol based on a mathematical model to fill this gap.

3 Mathematical model

This study applies a traditional supply chain design model aimed at minimizing total cost of a system composed by multiple companies in eq. (1). The total cost per company is defined in eq. (2). Also, we incorporated a constraint that ensures that none of the participating firms incurs higher costs in eq. (3). To this end, the operating cost for each firm was estimated by solving a model under a non-collaborative scenario, in which each firm uses its own resources. This cost was used as the maximum allowable cost per firm, given that any increase in this cost could discourage firms from participating in a π-network, even more when considering the organizational effort involved in sharing resources with other firms.

$$Min\ f(n) = \sum_{n} CtotalE_n \qquad \forall\, n \quad (1)$$

$$CtotalE_n = CPRE_n + CA_n + CCFCD_n + CV1E_n + CV2E_n + CV4E_n + CDE_n \qquad \forall\, n \quad (2)$$

$$CtotalE_n \leq CE_n \qquad \forall\, n \ (3)$$

$\boldsymbol{CtotalE_n}$	Total cost of the company *n*
$\boldsymbol{CPRE_n}$	Total cost of production of the company *n*
$\boldsymbol{CA_n}$	Total storage costs for the company *n*
$\boldsymbol{CCFCD_n}$	Fixed cost for the use of the company's hubs *n*
$\boldsymbol{CV1E_n}$	Cost of trips made by the company *n* from its production plants to its hubs
$\boldsymbol{CV2E_n}$	Cost of trips made by company *n* from its production plants to the hubs of other companies
$\boldsymbol{CV4E_n}$	Cost of trips made by company *n* from hubs to consumption centers
$\boldsymbol{CDE_n}$	Cost per kilometers traveled on all trips made by the company *n*
$\boldsymbol{CE_n}$	Individual cost of the company *n*

The proposed model was implemented for a theoretical case study based on a distribution chain in Colombia. The characteristics of the supply networks in this country are in line with the objective of this study since there are large companies, which coexist with micro, small and medium-sized enterprises that represent 92%, 5.9% and 1.6% of the total companies in the country, respectively (Sebastián González Patiño et al., 2024). Additionally, in this type of network, sharing economies of scale can be beneficial, especially for small and medium-sized enterprises since according to the country's recent national logistics survey, the percentage of

the logistics cost of small companies is 24.3%, doubling that of large companies, which report a cost of 12.1%, which generated an increase of 5 points in the country's logistics cost between 2020 and 2022(Departamento Nacional de Planeación - DNP-, 2022). Finally, the market in Colombia is distributed in the main cities, which are geographically dispersed and have dissimilar demands due to their population.

The chain under consideration consists of three companies, each with two production plants and two distribution centers, as well as 18 consumption centers. The production plants were randomly located in the cities with the highest industrial presence in the country, while for the consumption centers, 5 large cities and 13 medium-sized cities were considered, which determined the demand for each one. The model is based on the following assumptions:

- Total network demand follows a uniform distribution between 18,000 and 32,400 and is distributed proportionally among consumption centers based on their size.
- Demand is met by each firm in proportion to its production capacity.
- The costs for using distribution centers are calculated based on the costs of the company that owns them and are borne by the company that owns the cargo.
- The fleet is shared for transportation between distribution centers and consumption centers. Since a single vehicle can carry cargo from multiple companies, this cost is distributed proportionally among all companies.

To analyze changes in the supply chain configuration under different distribution center locations and different company capacities, the models were solved considering the following scenarios:

Table 1: Description of Scenarios

Scenario	Hub's location	Companies' production and storage capacity	Probability distribution of production capacity by plant
S1	Near the production facilities	Symmetric	Uniform (4,000, 5,000)
S2	Near the production facilities	Asymmetric	Large plants: Uniform (6,000, 7,000); Medium plants: Uniform (4,000, 5,000); Small plants: Uniform (1,000, 2,000)
S3	Near the main consumption centers	Symmetric	Uniform (4,000, 5,000)
S4	Near the main consumption centers	Asymmetric	Large plants: Uniform (6,000, 7,000); Medium plants: Uniform (4,000, 5,000); Small plants: Uniform (1,000, 2,000)

Five random data sets were generated for each scenario, in accordance with the distributions listed in Table 1. In order to analyze the network's behavior, these scenarios were solved using three mathematical models: when each firm uses its own resources (T-M); under a collaborative framework aimed at minimizing total cost (C-M); and, finally, under the PI design, which ensures that no firm experiences an increase in its costs (PI-M). The results are presented below.

4 Computational Results

4.1 Overall Results

Table 2 shows the results of the average total network cost when each firm uses its resources individually (column 2). The collaborative model is analyzed in columns 3–5, and the PI model

in columns 6–8. For the collaborative model (C-M) and the Physical Internet Model, the average cost reduction for each company is presented (Avg. Reduction column), the maximum cost increase and reduction is also presented (Max. increase/ reduction respectively).

Table 2: Average variation in total network cost

Scenario	Avg Total Cost	C-M			PI-M		
		Avg. reduction	Max. increase	Max. reduction	Avg. reduction	Max. increase	Max. reduction
S1	$ 29,316	-2.6%	4.0%	-15.2%	-2.5%	0.0%	-9.2%
S2	$ 32,682	-3.2%	1.2%	-34.4%	-3.2%	0.0%	-28.6%
S3	$ 29,909	-2.5%	4.8%	-17.8%	-2.4%	0.0%	-10.3%
S4	$ 36,921	-3.1%	0.9%	-20.2%	-3.1%	0.0%	-18.5%

Similar reductions in total cost for both the collaborative model and the PI model; however, upon analyzing the collaborative model, an increase in cost is observed for some companies in all scenarios, which is offset by reductions for others, such that the overall result is lower than what would be achieved if each company used its own infrastructure. Table 2 shows that for this model, the maximum increase observed was 4.8% for one company in Scenario 3. On the other hand, although the reductions in total cost under the PI model are smaller than those under the collaborative model, there is no evidence of increased costs for any of the companies.

Related to the use of hubs, Table 3 shows the results of storage capacity utilization when each firm uses its resources individually (column 2), when the network is designed under a collaborative scheme (column 3), and when the network is designed using PI (columns 4).

Table 3: Average hub occupancy

Scenario	T-M	C-M	PI-M
S1	74,8%	89,5%	89,7%
S2	86,6%	97,3%	95,8%
S3	76,2%	88,0%	88,0%
S4	84,9%	95,1%	94,9%

The results show that when each company uses its resources individually, hubs operate at an average capacity of 80%; however, under the PI model, this percentage rises to 92.4%. A similar situation occurs with the collaborative model.

Finally, regarding CO_2 emissions, table 4 shows that the computational results are not conclusive, as both models show increases and decreases across the different scenarios.

Table 4: Average variation in total emissions

Scenario	C-M	PI-M
S1	5,9%	1,5%
S2	-1,0%	0,3%
S3	-7,3%	2,7%
S4	0,8%	-8,5%

4.2 Analysis of Company Size and Location

Related to total cost reduction, as shown in Table 5, in the asymmetric scenarios (S2 and S4), a greater reduction in costs is observed in both the collaborative model and the PI model, exceeding 3% in some cases. In these scenarios, small firms (C3) benefit the most, with reductions ranging from 15,51% to 18,63% for the collaborative model and from 13,80% to 16,04% for the PI model. Scenarios S1 and S3, as well as S2 and S4, show similar cost reductions, suggesting that the location of hubs does not significantly affect total costs, whether they are located near production plants (S1 and S2) or near consumption centers (S3 and S4).

Table 5: Average variation in total cost per company

Company	S1		S3		S2		S4	
	PI-M	C-M	PI-M	C-M	PI-M	C-M	PI-M	C-M
C1	-1,61%	1,11%	-0,44%	2,22%	0,00%	0,58%	-0,04%	0,41%
C2	0,00%	1,23%	-0,98%	-0,36%	-0,02%	0,50%	0,00%	0,39%
C3	-6,78%	-11,95%	-6,28%	-10,30%	-16,04%	-18,63%	-13,80%	-15,51%
Average Total Cost	-2,52%	-2,58%	-2,37%	-2,51%	-3,18%	-3,24%	-3,06%	-3,14%

Regard to hub utilization, the table 6 shows that the traditional design shows utilization rates below 70% in all scenarios for some of the companies, whereas both the network designed under the collaborative scheme and the PI model show utilization rates close to or above 70% in most cases, with the exception of one small company in instance 2 of scenario 4.

Table 6: Average hub occupation per company

Company	S1			S3			S2			S4		
	T-M	PI-M	C-M	T-M	PI-M	C-M	T-M	PI-M	C-M	T-M	PI-M	C-M
C1	83%	87%	92%	83%	88%	86%	91%	98%	98%	90%	97%	97%
C2	82%	92%	85%	80%	86%	88%	91%	98%	98%	91%	97%	97%

Company	S1			S3			S2			S4		
	T-M	PI-M	C-M	T-M	PI-M	C-M	T-M	PI-M	C-M	T-M	PI-M	C-M
C3	60%	90%	93%	67%	92%	90%	77%	74%	75%	73%	86%	88%
Average all companies	75%	90%	90%	77%	89%	88%	86%	90%	90%	85%	93%	94%
Min all companies	58%	79%	79%	56%	78%	80%	69%	86%	82%	58%	57%	60%
Max all companies	83%	99%	100 %	100%	100%	98%	91%	100%	100%	91%	100%	100%

This capacity can be used both to store the company's own products and to store products from other companies in the network. Table 7 shows the percentage of capacity given to other companies. It can be seen that in the asymmetric scenarios (S2 and S4), small companies (C3) allocate a small percentage of their capacity, in contrast to large and medium-sized companies.

Table 7: Average variation in storage capacity given to other companies

Company	S1		S3		S2		S4	
	PI-M	C-M	PI-M	C-M	PI-M	C-M	PI-M	C-M
C1	32%	27%	44%	43%	32%	33%	22%	30%
C2	34%	18%	35%	20%	39%	25%	37%	25%
C3	17%	15%	30%	34%	4%	10%	15%	22%

5 Conclusions

The results show that there are differences in network performance when networks consist of firms with varying capacities, revealing an advantage for small firms. This cost advantage stems from the ability to leverage the resources of other firms with lower costs. This has a potential impact on logistics costs in countries such as Colombia, where micro, small and medium-sized enterprises make up the business ecosystem to a greater extent.

Similarly, sharing distribution center capacity leads to the discontinuation of network distribution centers in all scenarios; particularly in asymmetric scenarios, distribution centers belonging to small firms were no longer used due to their higher costs. This generates lower operating costs due to an average use of more than 89% of the hubs' capacity, while allowing the division of costs between the different companies in the network, benefiting not only small companies.

Regarding the location of distribution centers, the results show no difference in cost between locating them near production plants or near consumption centers.

On the other hand, as expected, the collaborative design model results in lower total network costs than the PI model; however, some firms experience cost increases, particularly in asymmetric scenarios, where this increase is observed among large and medium-sized firms,

this is because this model does not take into account the proposed win-win restriction. Finally, the results regarding emissions are inconclusive.

The findings from this research provide indications of potential benefits for companies of various sizes; however, further experimentation is required, both in terms of the number of instances and network characteristics—for example, by considering more realistic networks with a greater number of nodes or allowing new nodes to be opened, among other approaches. Additionally, other indicators can be evaluated, such as resilience and maximum travel time, the latter of which is related to driver well-being.

Finally, future research should address the multidimensional challenges of operationalizing the Physical Internet, specifically focusing on the friction points of network interworking. This includes analyzing governance policies for shared flows and the deployment of standardized π-containers that facilitate collaborative transport and storage modularity, thereby overcoming the barriers to horizontal resource integration.

6 References

Alliance for Logistics Innovation through Collaboration in Europe. (2019). *Alice-Zero-Emissions-Logistics-2050-Roadmap-WEB*. https://www.etp-logistics.eu/wp-content/uploads/2019/12/Alice-Zero-Emissions-Logistics-2050-Roadmap-WEB.pdf

Ariztegui-Beltrán, O. H., Cortés-Murcia, D. L., Guerrero Rueda, W., & Péton, O. (2024). *PI-Service network design for fluvial cargo transportation*. In International Physical Internet Conference (pp. 273-282).

Calatayud, A., & Montes, L. (2021). *Logística en América Latina y el Caribe: oportunidades, desafíos y líneas de acción*.

Departamento Nacional de Planeación - DNP-. (2022). *Coordinación editorial-Oficina Asesora de Comunicaciones del DNP*. www.dnp.gov.co

Furtado, P., & Frayret, J. M. (2015). Proposal sustainability assessment of resource sharing in intermodal freight transport with agent-based simulation. *IFAC-PapersOnLine*, *28*(3), 436–441. https://doi.org/10.1016/j.ifacol.2015.06.120

Hayek, M., Mahringer, G., Segner, R., Landschützer, C., & Sihn, W. (2023). Development of a Physical Internet container for an optimized wood supply chain. *Transportation Research Procedia*, *72*, 1950–1957. https://doi.org/10.1016/j.trpro.2023.11.675

Ji, S., Zhao, P., & Ji, T. (2023). A Hybrid Optimization Method for Sustainable and Flexible Design of Supply–Production–Distribution Network in the Physical Internet. *Sustainability (Switzerland)*, *15*(7). https://doi.org/10.3390/su15076327

Kim, N., Montreuil, B., Klibi, W., & Kholgade, N. (2021). Hyperconnected urban fulfillment and delivery. *Transportation Research Part E: Logistics and Transportation Review*, *145*. https://doi.org/10.1016/j.tre.2020.102104

Kulkarni, O., Dahan, M., & Montreuil, B. (2022). Resilient Hyperconnected Parcel Delivery Network Design Under Disruption Risks. *International Journal of Production Economics*, *251*. https://doi.org/10.1016/j.ijpe.2022.108499

Meller, R. D., Lin, Y. H., & Ellis, K. P. (2012). The impact of standardized metric Physical Internet containers on the shipping volume of manufacturers. *IFAC Proceedings Volumes (IFAC-PapersOnline)*, *14*(PART 1), 364–371. https://doi.org/10.3182/20120523-3-RO-2023.00282

Montreuil, B. (2011). Toward a Physical Internet: meeting the global logistics sustainability grand challenge. *Logistics Research*, *3*(2), 71–87. https://doi.org/10.1007/s12159-011-0045-x

Montreuil, B., Ballot, E., & Fontane, F. (2012). An open logistics interconnection model for the physical internet. *IFAC Proceedings Volumes (IFAC-PapersOnline)*, *45*(6 PART 1), 327–332. https://doi.org/10.3182/20120523-3-RO-2023.00385

Sallez, Y., Pan, S., Montreuil, B., Berger, T., & Ballot, E. (2016). On the activeness of intelligent Physical Internet containers. *Computers in Industry*, *81*, 96–104. https://doi.org/10.1016/j.compind.2015.12.006

Sarraj, R., Ballot, E., Pan, S., Hakimi, D., & Montreuil, B. (2014). Interconnected logistic networks and protocols: Simulation-based efficiency assessment. *International Journal of Production Research*, *52*(11), 3185–3208. https://doi.org/10.1080/00207543.2013.865853

Sebastián González Patiño, J., Claudia, M., & Valenzuela, L. (2024). *Una mirada a las mipymes en Colombia*.

Tordecilla, R. D., Montoya-Torres, J. R., & Guerrero, W. J. (2025). Resilient design of hyperconnected multiactor Physical Internet supply chain networks. *International Transactions in Operational Research*, *32*(6), 3528–3564. https://doi.org/10.1111/itor.13615

Walha, F., Bekrar, A., Chaabane, S., & Loukil, T. M. (2016). A rail-road PI-hub allocation problem: Active and reactive approaches. *Computers in Industry*, *81*, 138–151. https://doi.org/10.1016/j.compind.2016.04.007

Yang, Y., Pan, S., & Ballot, E. (2017). Mitigating supply chain disruptions through interconnected logistics services in the Physical Internet. *International Journal of Production Research*, *55*(14), 3970–3983. https://doi.org/10.1080/00207543.2016.1223379

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE
KEDGE BUSINESS SCHOOL

Resource Pooling for Multiple Stakeholders in Hyperconnected Warehouse Networks under Demand Uncertainty

Yaxin PANG[1], Shenle PAN[2], Olivier LABARTHE[1], and Walid KLIBI[1]
1. The Centre of Excellence in Supply Chain and Transportation (CESIT), Kedge Business School, France
2. Mines Paris, PSL University, Centre for management science (CGS), i3 UMR9217 CNRS, Paris, France

Corresponding author: yaxin.pang@kedgebs.com

Abstract: *The Physical Internet envisions open, hyperconnected logistics where resources flow across organizational boundaries. This paper addresses an underexplored element: resource pooling among warehouses in dense clusters. We propose a framework where multiple independent stakeholders form a coalition, governed by a neutral orchestrator, to share different types of resources. Under demand uncertainty, the coalition decides how many existing resources to reserve and how many new resources to purchase. A two-stage stochastic program ensures each stakeholder benefits. The research advances the Physical Internet by operationalizing hyperconnected warehouse networks through a practical, multi-stakeholder sharing mechanism.*

Keywords: *Physical Internet, Hyperconnected warehouse network, Two-stage stochastic programming, Demand uncertainty, Logistics collaboration*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☒ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☒ *PI Networks,* ☐ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

1 Introduction

The rapid expansion of e-commerce has transformed warehouses from passive storage facilities into dynamic fulfillment centers (Boysen, N., and de Koster, R. 2024). Orders arrive in unpredictable waves. Same-day and next-day delivery expectations have compressed response times from days to hours. Moreover, customers expect near-perfect reliability. Behind the scenes, warehouse operators face a difficult balancing challenge. They must maintain enough capacity to handle peak demand, yet they cannot afford to let expensive resources sit idle during slow periods.

Three pressures have made this balancing challenge increasingly difficult. First, labor shortages have become structural. According to the U.S. Bureau of Labor Statistics (2026), approximately 180,000 workers in the transportation, warehousing, and utilities sector quit their jobs in February 2025 alone. Warehousing is physically demanding work, and competition from other sectors has made it hard to attract and retain pickers and packers. Even when workers are available, absenteeism rates are high. Second, automation costs remain high. Autonomous

mobile robots (AMRs) and automated wrapping machines offer productivity gains, but their upfront capital cost puts them out of reach for many operators, especially smaller ones (Baker and Canessa, 2019; Fragapane et al., 2021). Third, demand volatility has intensified. Promotions, seasonal peaks, and unexpected events create demand patterns that are difficult to predict and even harder to match with fixed capacity.

The Physical Internet (PI) offers a way forward. The Physical Internet envisions a global logistics system where goods are transported, stored, and handled through standardized, shared infrastructure (Montreuil 2011). In this vision, warehouses are not isolated facilities but hyperconnected nodes in a network. Resources such as workers, robots, equipment, can flow between nodes in response to uncertain demand crossing multiple stakeholders as needed.

Precisely, a practical step toward this vision is the hyperconnected warehouse network (Pang et al. 2025). Figure 1 illustrates a hyperconnected warehouse network in Physical Internet and Physical Intranet. In areas near major cities, ports, and transportation hubs, dozens of warehouses now operate within a few square kilometers. In such areas, each warehouse typically operates independently. One facility may struggle to meet demand while its neighbor, with a different demand pattern, has residual workers and residual robots. This inefficiency persists despite physical proximity. But the physical proximity creates an opportunity that has been unexploited: the possibility of pooling resources across the warehouses even for different logistics service providers.

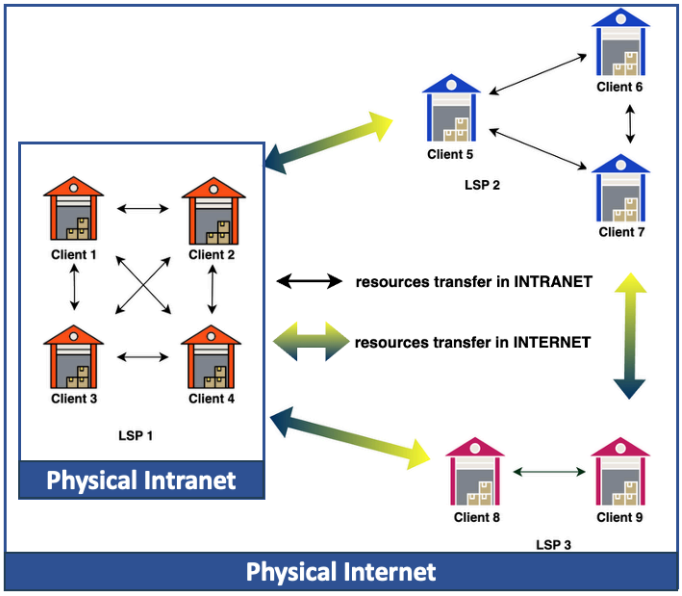


Figure 1: Hyperconnected Warehouse Network (example of three stakeholders), adapted from Pang et al. (2025).

In the hyperconnected warehouse network, multiple warehouses in the same geographic area agree to pool certain resources. When demand surges at one warehouse, resources from the pool or from other warehouses can be deployed there temporarily. When demand is low, resources remain in the pool, available to others. This arrangement offers potential benefits for all participants. First, demand uncertainty makes capacity planning notoriously difficult; our two-stage sharing mechanism allows warehouses to commit capacity strategically before demand is known and adjust operationally afterwards. Second, better capacity management through pooling reduces the need for each warehouse to over-invest for rare peaks, lowering total system cost. Third, labor shortages in logistics can be mitigated by sharing a flexible pool of workers across multiple warehouses, smoothing peaks and reducing idle time. Fourth, the slow integration of robotics in warehousing can be accelerated because a shared pool allows

multiple stakeholders to collectively invest in automation, and spreading the benefits. However, realizing these benefits requires careful planning. How large should the shared pool be? How should resources be allocated when demand is known? And how can we ensure that every stakeholder benefits enough to participate willingly?

This paper bridges the gap between the deterministic, single-stakeholder model and the full vision of a hyperconnected warehouse network under the PI. We make following contributions. On the one hand, while existing PI research has focused on transportation and single-stakeholder warehouse intranets, this paper addresses multi-stakeholder resource pooling in dense warehouse clusters. Unlike prior work that assumes deterministic demand, we introduce a neutral orchestrator and a two-stage sharing mechanism with reserved resources. On the other hand, our work further moves the PI vision from intra-firm coordination to truly horizontal, inter-firm collaboration under uncertainty.

The remainder of this paper is organized as follows. Section 2 reviews the relevant literature. Section 3 elaborates the proposed coalition mechanism. Section 4 explains the problem.

2 Literature review

Under PI paradigm, logistics resources such as vehicles, containers, warehouses, and handling equipment become interconnected and accessible across different organizations. Over the past decade, PI has attracted growing attention from both academia and industry, with research focusing on collaborative transportation, shared freight networks, and synchronodal mobility (Ballot, 2018; Chabot et al., 2018; Labarthe et al., 2024). A key theme in this literature is the potential for resource pooling to improve utilization, reduce waste, and enhance system resilience (Wu et al., 2025).

Within the PI framework, warehouses are conceptualized as hyperconnected nodes. However, most existing PI resource pooling research has focused on transportation rather than warehousing. As critical nodes for inventory management and order fulfillment, warehouses remain relatively underexplored in the transition toward hyperconnected logistics. Resource pooling in warehouse settings has been studied mainly through two lenses: inventory pooling and collaborative resource sharing.

Inventory pooling is a well-established practice where multiple entities share a common stock to reduce total inventory holding costs while maintaining service levels (Alptekinoğlu et al., 2013). Research has demonstrated that inventory pooling can mitigate demand uncertainty, reduce disruption risks, and improve sustainability outcomes, though the benefits depend on factors such as demand correlation and information quality (Yilmaz, 2025).

Collaborative resource sharing within warehouses has received increasing attention. Jamili et al. (2022) investigated the impact of sharing labor and dock doors among multiple clients operating in a single warehouse, developing optimization models that quantify the resulting improvements in order tardiness. Extending this work, Jamili et al. (2025) examined how different resource-sharing configurations affect warehouse design, showing that the optimal strategy depends on client characteristics such as demand rate and variability.

Beyond single-warehouse settings, Mathlouthi et al. (2015) proposed the concept of "pooled warehouses," where multiple storage facilities share capacity to maximize accepted requests while maintaining product quality. Using agent-based simulation, they compared different storage allocation strategies among actors. However, their work did not consider operational resources such as labor or equipment, nor did it address demand uncertainty.

The most relevant work is Pang et al. (2025), which introduced the hyperconnected warehouse network (HWN) under the PI framework. In their model, geographically proximate warehouses operated by the same logistics service provider pool three critical resources: human labor (with location preferences), wrapping machines, and autonomous mobile robots. They formulate a mixed-integer linear programming model that minimizes total costs, including resource utilization, transfers, and workforce preference mismatches. Their illustrative example demonstrates that resource pooling can reduce total costs by approximately 10% compared to dedicated warehouse operations.

This work extends the work of Pang et al. (2025). On the one hand, we assume demand is uncertain instead of deterministic. On the other hand, we extend to multiple independent stakeholders which introduces participation constraints that each stakeholder must see a clear benefit to join the pool.

To summarize, existing literature on warehouse resource pooling has focused on inventory rather than operational resources (labor, robots, equipment) and assumed deterministic demand. Research on shared mobility (e.g., bike-sharing, car-sharing) involves similar two-stage resource allocation problems but typically assumes identical resources and does not account for efficiency differences across users or the multi-resource nature of warehouse operations.

Thus, a model that combines multiple stakeholders, multiple operational resources, and demand uncertainty remains absent from the literature. Our research fills this gap by extending the hyperconnected warehouse network concept to a multi-stakeholder, uncertain environment.

3 Hyperconnected warehouse network

3.1 The network

Traditional warehouse networks operate independently. Each logistics service provider or retailer manages its own warehouses, its own workforce and its own equipment. Even when warehouses stand side by side in an industrial park, they rarely share resources. Occasional collaboration exists, for example, sharing storage space or swapping inventory for contingency, but systematic, dynamic sharing of operational resources such as workers, robots, and wrapping machines remains exceptional. The PI vision changes this perspective. PI calls for open, hyperconnected logistics where resources flow across company boundaries as seamlessly as data packets flow across the digital internet.

Beyond being a purely theoretical concept, geographic proximity is what makes collaboration feasible and economically attractive. In a dense warehouse area such as an industrial park, warehouses are located very close to one another. The distance between any two warehouses is at most a few kilometers, and most are within walking distance, typically between 100 meters and one kilometer. This proximity allows mobile resources to be shared across warehouses, with travel times measured in minutes rather than hours. Moreover, warehouses are connected by a network of roads, and walkways, further facilitating efficient movement.

3.2 The role of orchestrator

Figure 2 illustrates a simple coalition with three stakeholders, i.e., A, B and C, each operating warehouses A_1 and A2, B_1, and C_1. All communicate with a neutral orchestrator. The orchestrator facilitates the coalition contract (strategic, fixed for the horizon) and then, before demand is known, produces a resource reserving plan. After actual demand is revealed, it executes dynamic relocation of resources between warehouses (transfers, rentals, overtime).

Stakeholders do not interact directly. All coordination goes through the orchestrator, which acts as a single trusted decision-maker.

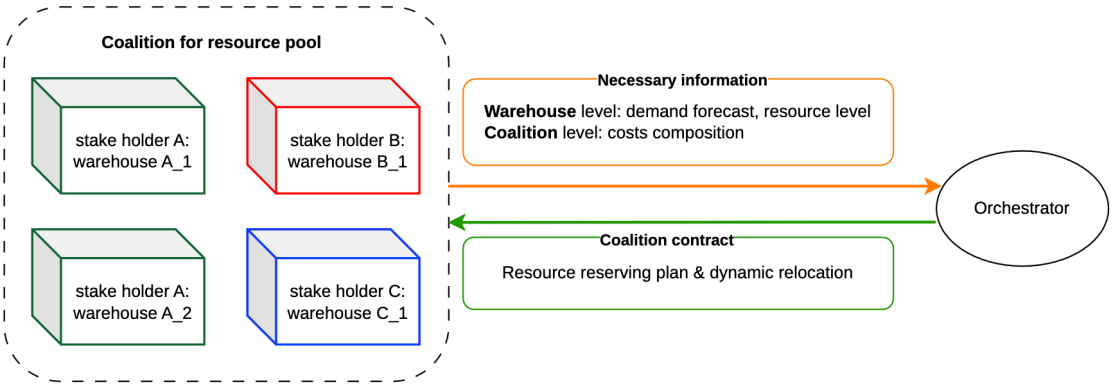


Figure 2: illustration of the orchestrator role and coalition structure.

Specially, we assume that the orchestrator knows each warehouse's demand forecasts, existing resource levels, and all cost parameters. It is benevolent which means it minimizes total coalition cost but not its own profit.

3.3 Sharing mechanism

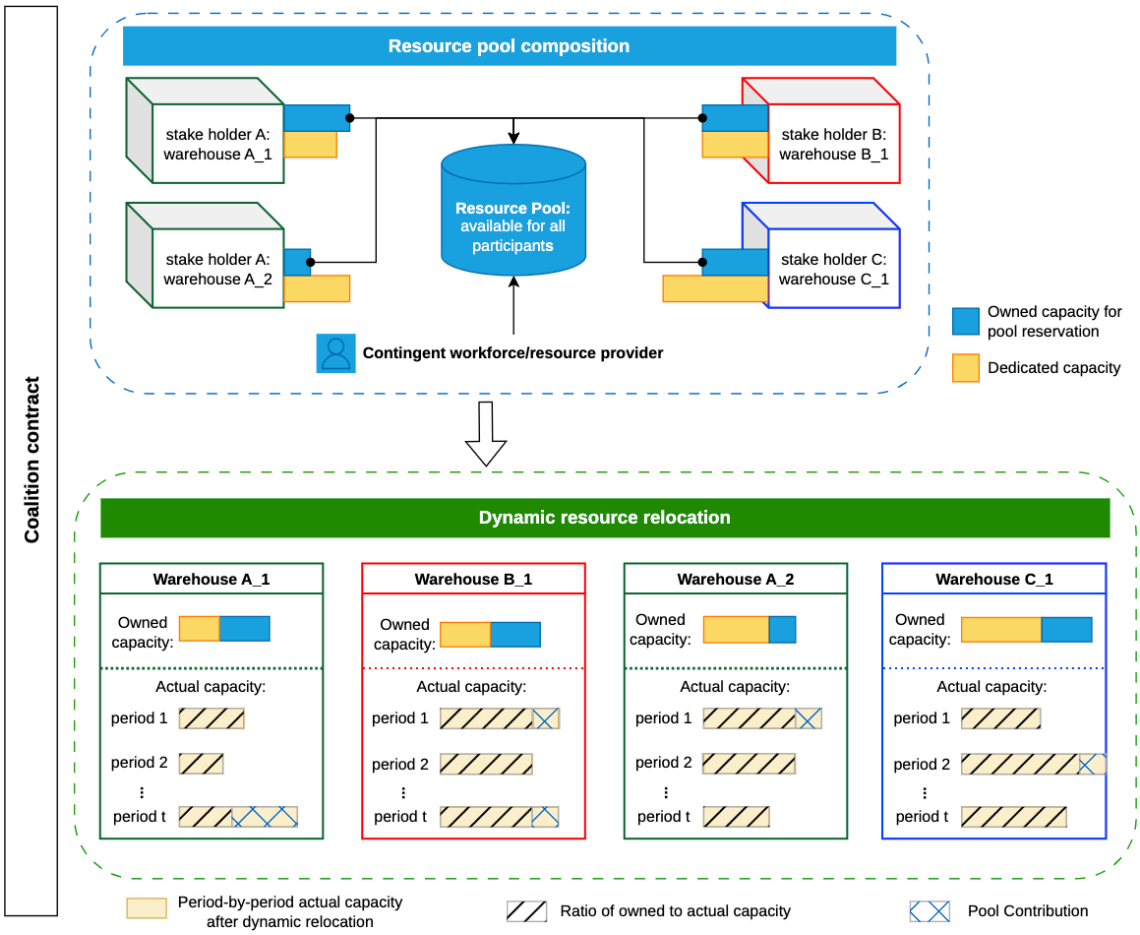


Figure 3: Coalition-based resource pooling and dynamic capacity sharing across hyperconnected warehouse networks.

Figure 3 illustrates a two-phase sharing mechanism within a coalition contract among multiple stakeholders who each operate one or more warehouses.

In *Resource Pool Composition*, each stakeholder contributes a portion of their owned warehouse capacity into a shared resource pool (shown as the central cylinder). Capacity is split into two types: dedicated capacity which stays exclusively with the warehouse owner, and owned capacity for pool reservation which is committed to the shared pool. A contingent workforce or resource provider can also feed into this pool, making it available to all coalition participants.

In *Dynamic Resource Relocation*, the pooled resources are redistributed across warehouses on a period-by-period basis according to actual demand. Each warehouse's actual capacity in a given period can differ from its owned capacity: it may grow (by drawing from the pool) or shrink (by contributing idle capacity back). The hatched bars show the ratio of owned-to-actual capacity, while the cross-hatched portions represent pool contributions.

Each stakeholder joining the coalition is independent for the scheduling inside their own warehouses. Each warehouse has its own demand distribution deriving from historical data. Each manage the cost for his own private resources which are resources owned exclusively outside the pool. When the resources in the pool and private resources are not enough for the realized demand, the stakeholders can use services of overtime labor, rental equipment to satisfy the demand with additional cost. If demand is unmet, there would be a penalty for the coalition for reduced service level. Stakeholders will only agree to participate in the pool if their net benefit from pooling is acceptable.

The coalition contract and the reserving plan are fixed for the entire coalition contract. The orchestrator does not revise the pool size or contribution levels after the horizon starts. Only transfers, rentals, and overtime are adjusted dynamically. The key idea is that instead of each stakeholder independently sizing their warehouse capacity for peak demand (which leads to underutilization), the coalition allows them to share surplus capacity dynamically over time which has the potential of reducing costs, improving utilization, and providing flexibility to handle demand fluctuations across the network.

4 Problem description

4.1 Resources for pooling

Warehouses differ in several dimensions that affect resource sharing. Some warehouses are primarily storage facilities, others are cross-docks or fulfillment centers. Therefore, some have requirements for pallet wrappers, while others do not. Besides, the warehouses also differ in traditional picking, or automated systems, which affects resource suitability and labor efficiency for workers assigned from other warehouses.

We consider three types of resources. These are representative of the resources most commonly used in warehouse operations:

Human labor. Workers are the most flexible resource but also the most constrained. A worker assigned to a warehouse that is not their home warehouse works less efficiently. We model this as an efficiency factor $e_{w,n} \in [0,1]$ for a worker from their home warehouse w when working at warehouse n. For a worker at their home warehouse, $e = 1$. For a neighboring warehouse, e could be 0.8. For a warehouse further away or with a different layout, e could be 0.5. Workers also have preferences for specific warehouses, based on commute convenience, workplace familiarity, or social ties. Assigning a worker to a non-preferred warehouse incurs a

compensation cost. This reflects real human resource considerations that are often overlooked in optimization models.

Autonomous mobile robots (AMRs). AMRs transport goods within and between warehouses. They can be shared across warehouses by allowing them to travel on designated pathways connecting facilities or transferred by vehicles. AMRs are homogeneous (same speed, battery capacity, payload). Transferring an AMR from one warehouse to another takes time and incurs a transfer cost (e.g., vehicle cost, distance-dependent variable cost, setup time).

Pallet wrapping machines. Wrapping machines are used to secure pallets before shipping. In a shared setting, wrapping machines can be placed in a shared yard accessible to multiple warehouses, or mobile units can be moved between warehouses. Each machine can wrap one pallet at a time, and wrapping takes a fixed amount of time.

4.2 Resources transfer

Not every resource can be used in every warehouse. For example, wrapping machines are only applicable at warehouses requiring them. AMRs are applicable at all warehouses with suitable floor surfaces. Labor is applicable at all warehouses (since workers can pick in any facility, but with different efficiency and preference violation cost).

Resources can be transferred between warehouses. Transferring a unit of resource from warehouse i to warehouse j incurs a cost depending on the distance between the two warehouses and resource type, which may include transportation, setup activities, and idle time. We assume that any transfer within the warehouse network can be completed within one period (e.g., one day), so we do not model transfer time explicitly.

For labor, transfers are also subject to a maximum travel distance. A worker's preference to the warehouse is based on the distance between the warehouse and their home warehouse. This reflects practical limits on commuting or walking between facilities.

4.3 Uncertainty modeling

The demand in each warehouse is uncertain for the planning horizon, while following a known distribution extracted from historical data. Different warehouses may have different demand distributions. We assume demand is independent across warehouses. This is realistic when warehouses serve different customers or product categories. Demand is also independent across periods, meaning that knowing demand in one period gives no information about demand in another period.

4.4 Two-stage tactical stochastic planning

The planning horizon is a strategic period of 3 months, which is a contract period between the stakeholders in the hyperconnected warehouse network. The horizon is divided into T operational periods. Each period could be a week (T = 12) or a day (T = 90).

Warehouses that join the coalition agree to contribute a portion of their existing resources to a shared pool, potentially purchase additional resources for the pool, share in the costs of pool maintenance and new purchases, and allow their resources to be transferred to other coalition members when needed. The coalition is formed at the beginning of the planning horizon and remains fixed for the entire contract period.

The shared pool for each resource type is composed of two parts: (1) Reserved resources from each warehouse: existing resources that warehouses agree to make available to the pool. These resources are physically located at the home warehouse but are available for transfer to other

warehouses when needed. (2) Newly purchased resources for the pool: additional units purchased specifically for the pool. These new resources are initially placed at a central pool location (e.g., a shared yard).

Private resources are not part of the optimization. These are considered given and are not decision variables. The orchestrator only optimizes over the shared pool. Precisely, the coalition collectively decides the total number of units of each type of resource that will be available in the shared pool. The pool consists of existing resources that warehouses contribute to the pool and new resources purchased specifically for the pool. Both are decisions to be made before any uncertainty (of demand) is realized. Maintenance cost for existing resources and purchase cost for new resources added to the pool will incur when forming this coalition.

Then with complete knowledge of demand and availability for all periods, the orchestrator makes allocation decisions sequentially over $t = 1 \ldots T$. Because the uncertainty distribution is already known, the entire horizon can be planned optimally at once for each uncertainty scenario. Then the orchestrator's decisions for each period t (scenario-dependent) includes: units of each type of resource transferred between warehouses, units of resource rented from a third party when pool resources are insufficient, overtime hours for workers at each warehouse used when regular pool labor is insufficient, and demand that cannot be fulfilled incurring a penalty cost.

Costs incurred in Stage 2 (per scenario ω, summed over all periods t): regular labor wages for pool workers assigned to warehouses, overtime labor wages, usage fees for shared AMRs and wrapping machines for resources actually used), rental costs for backup resources, inter-warehouse transfer costs per unit transferred, worker preference compensation when workers assigned to non-preferred warehouses, and unmet demand penalties for per unit of unmet demand.

To address the problem described above, we will adopt a two-stage stochastic programming framework. This choice is natural because the coalition faces a clear sequence of decisions: first, under uncertainty, it commits to pool composition (reserved resources and new purchases); then, after demand is revealed, it makes operational adjustments through transfers, rentals, overtime, and unmet demand. To ensure that every warehouse finds the coalition beneficial, we impose constraints on each participant's benefit relative to operating alone. The goal of the model is to minimize total costs including 1st stage resource pooling costs and expected total operational costs over the stochastic scenarios in the 2nd stage.

Conclusion

This paper proposes a hyperconnected warehouse network framework in which multiple independent stakeholders form a coalition to share operational resources. The framework centers on two key elements. First, a neutral orchestrator governs the coalition, collects information, and executes resource allocation. Second, a two-stage sharing mechanism separates strategic commitment (reserving existing resources and purchasing new ones) from operational relocation (transfers, rentals, overtime) after demand uncertainty is resolved. Future work can extend the model to include resource availability uncertainty, correlated demand, and fair cost allocation.

References

- Alptekinoğlu, A., Banerjee, A., Paul, A., & Jain, N. (2013). Inventory pooling to deliver differentiated service. *Manufacturing & Service Operations Management*, 15(1), 33-44.
- Baker, P., & Canessa, M. (2019). Warehouse automation in the age of e-commerce: Productivity versus affordability. *International Journal of Logistics Management*, 30(3), 821–842.

- Ballot, E. (2018). The physical internet. In Operations, Logistics and Supply Chain Management (pp. 719-734). Springer.
- Boysen, N., & De Koster, R. (2025). 50 years of warehousing research—An operations research perspective. European Journal of Operational Research, 320(3), 449-464.
- Chabot, T., Bouchard, F., Legault-Michaud, A., Renaud, J., & Coelho, L. C. (2018). Service level, cost and environmental optimization of collaborative transportation. Transportation Research Part E, 110, 1-14.
- Fragapane, G., de Koster, R., Sgarbossa, F., & Strandhagen, J. O. (2021). Planning and control of autonomous mobile robots for intralogistics: Literature review and research agenda. *European Journal of Operational Research*, 294(2), 405–426.
- Jamili, N., Van Den Berg, P. L., & De Koster, R. (2022). Quantifying the impact of sharing resources in a collaborative warehouse. European Journal of Operational Research, 302(2), 518-529.
- Jamili, N., Van Den Berg, P. L., & de Koster, R. (2025). The impact of resource sharing on the design of multi-client warehouses. IISE Transactions, 57(3), 336-350.
- Labarthe, O., Ahmadi, G., Klibi, W., Deschamps, J. C., & Montreuil, B. (2024). A sustainable on-demand urban delivery service enabled by synchronodality. Transportation Research Part C, 161, 104544.
- Mathlouthi, W., Saoud, N. B. B., & Shoui, S. (2015). Agent-based modeling and simulation of pooled warehouse intelligent management. In *Proceedings of the Conference on Summer Computer Simulation* (pp. 1-8).
- Montreuil, B. (2011). Towards a Physical Internet: Meeting the global logistics sustainability grand challenge. *Logistics Research*, 3(2-3), 71-87.
- Pang, Y., Klibi, W., & Labarthe, O. (2025). Resource Pooling in Hyperconnected Warehouse Networks.
- U.S. Bureau of Labor Statistics. (2026). *Job openings and labor turnover – February 2025*. https://www.bls.gov/news.release/pdf/jolts.pdf
- Wu, H., Li, M., Yu, C., Ouyang, Z., Lai, K. H., Zhao, Z., ... & Huang, G. Q. (2025). Towards cyber-physical internet: A systematic review, fundamental model and future perspectives. *Transportation Research Part E*, 197, 104051.
- Yilmaz, O. (2025). A literature review on inventory pooling with applications. *Sustainability*, 17(2), 797.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Cyber-Physical Internet-enabled Spatial-Temporal Traceability for Production Logistics Synchronization: An Agentic Reasoning Approach

Jinpeng Li[1, 2] , Zhiheng Zhao[1, 2, *] and George Q. Huang[1, 2]

1. Department of Industrial and Systems Engineering, The Hong Kong Polytechnic University, Hong Kong, China
2. Research Institute for Advanced Manufacturing, The Hong Kong Polytechnic University, Hong Kong, China

Corresponding author: Zhiheng Zhao (zhiheng.zhao@polyu.edu.hk)

Abstract: *In discrete manufacturing, the growing demand for personalized production has significantly increased process complexity and placed higher requirements on production logistics (PL). To achieve zero inventory and even zero warehousing, PL synchronization has become increasingly important, since it ensures that operators, tools, and materials are available at the required location within a specified time window, thereby enabling production, assembly, and transportation activities to start and finish as planned. In this context, spatial-temporal traceability is essential for tracking PL resources and determining their availability in time and space. Recently, the Cyber-Physical Internet (CPI) has provided a promising foundation for connecting physical and digital entities and supporting real-time information sharing. However, relevant traceability information is often fragmented and stored across various CPI routers, which makes its collection, integration, and analysis highly challenging. Meanwhile, although large language models (LLMs) have shown strong capabilities in information understanding and reasoning, they remain prone to hallucinations and may lack reliability in complex industrial decision-making scenarios. To address these challenges, this paper proposes an agentic reasoning approach for spatial-temporal traceability. The approach introduces multiple function-specific agents to handle tasks such as information retrieval, data fusion, and stakeholder-oriented query support. Through collaboration among these agents, it improves the accuracy, coherence, and interpretability of traceability reasoning over fragmented data, thereby enabling different stakeholders to obtain the spatial-temporal information they need through CPI more effectively.*

Keywords: *Cyber-Physical Internet; Large Language Model; Spatial-Temporal Analysis; Production Logistics; Industry 4.0.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☐ *PI Networks,* ☒ *System of Logistics Networks,* ☒ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

1 Introduction

Production logistics (PL) plays a critical role in discrete manufacturing systems by coordinating the planning, control, and execution of material flows within industrial facilities. In practice, PL accounts for approximately 95% of total production time and up to 20%-40% of overall operational costs, making it a dominant factor influencing manufacturing efficiency and

responsiveness (Li et al., 2025b). PL synchronization refers to the capability of ensuring that all required production resources, including materials, equipment, and human operators, are available at the right location within the appropriate time window, such that production, assembly, and transportation activities can be executed as planned without delay or conflict. In this context, spatial-temporal traceability plays a critical role in enabling effective PL synchronization. By continuously tracking the states of resources across both space and time, it allows the system to determine resource availability, anticipate potential conflicts, and support coordinated logistics decision-making under dynamic conditions.

Recent research has explored various advanced technologies to enhance PL synchronization, with particular attention to the evolution from the Physical Internet (PI) to the Cyber-Physical Internet (CPI). The term PI was first proposed in 2006, introducing a paradigm for designing large-scale and complex logistics and transportation networks by emphasizing standardized protocols, modular containers, and intelligent interfaces (Samadhiya et al., 2023). While PI provides a robust framework for the interconnection of physical resources, its effectiveness diminishes in highly dynamic and uncertain environments (Boysen et al., 2025). To address these limitations, the CPI has recently emerged as a next-generation paradigm that extends beyond the traditional PI (Wu et al., 2025). Inspired by the transparency and flexibility of instant information exchange, CPI envisions logistics systems that deliver goods with comparable efficiency. By integrating digital twin technologies, CPI facilitates real-time synchronization of both cyber and physical aspects of PL activities, thereby overcoming the static limitations of PI (Li et al., 2025a). Through continuous coordination across multiple nodes, including material states, task progress, and resource availability, CPI enables more adaptive resource configuration, cooperative negotiation, and distributed decision-making, ultimately providing a scalable and resilient architecture for PL systems.

Despite these advancements, several critical challenges remain unresolved. First, traceability information is inherently fragmented across multiple CPI routers and stored in their routing tables, making unified representation and integration difficult. Second, while spatial-temporal data contains rich implicit knowledge about resource interactions and system dynamics, the lack of effective reasoning mechanisms limits its utilization for subsequent resilient decision-making. Third, although LLMs offer promising capabilities for knowledge understanding and interactive reasoning, their integration into industrial environments remains challenging due to reliability concerns, as they may produce inconsistent or untrustworthy outputs when handling complex, domain-specific tasks (Meriem et al., 2026).

To address the above challenges, this paper proposes an agentic reasoning approach for spatial-temporal traceability in PL under the CPI framework. The core idea is to decompose complex reasoning tasks into multiple function-specific agents, each responsible for subtasks such as information retrieval, data fusion, and knowledge reasoning. Through collaboration among these agents, fragmented spatial-temporal information can be systematically integrated and transformed into actionable knowledge. This approach not only enhances the accuracy and interpretability of reasoning results, but also improves robustness by constraining the reasoning process and reducing the impact of LLM hallucinations.

2 Literature Review

2.1 PI and CPI

A paradigm shift in logistics systems has been driven by the concept of the Physical Internet (PI), which rethinks the organization of logistics operations through principles analogous to data communication networks (Montreuil, 2011). Rather than treating logistics activities as

isolated processes, PI envisions an open and interconnected system in which standardized protocols, modular containers, and interoperable interfaces enable seamless collaboration across different logistics networks (Pan et al., 2021). Under this vision, the movement of physical goods can be orchestrated in a manner similar to packet transmission in the Internet, thereby improving efficiency, scalability, and sustainability. Such a transformation not only reshapes how physical objects are handled, transported, and stored, but also opens new possibilities for redesigning existing business models (Peng et al., 2021).

Building upon this perspective, the Cyber-Physical Internet (CPI) extends the PI paradigm by incorporating cyber-physical integration and emphasizing spatial-temporal awareness. In contrast to PI, which primarily focuses on logistics interconnectivity, CPI highlights the importance of fine-grained visibility and traceability of resources across both space and time. This enables a deeper understanding of resource interactions and supports more informed and resilient decision-making in complex production environments. Recent studies have explored the application of CPI in various domains. For instance, CPI-based frameworks have been developed to improve coordination in construction supply chains through enhanced information sharing and incentive mechanisms (He et al., 2025). In logistics transportation, CPI has been utilized to model and optimize resource allocation in multi-stage manufacturing (Li et al., 2026). Moreover, CPI-inspired network architectures have been proposed to reorganize logistics infrastructures using routing protocols, enabling more efficient and cost-effective delivery of physical goods (Lu et al., 2026). These efforts demonstrate the potential of CPI to provide a unified foundation for intelligent and data-driven logistics management.

2.2 PL Synchronization

PL synchronization concerns the coordinated delivery of required resources to designated locations at appropriate times, ensuring the continuity of operations across interconnected production stages. This concept is closely aligned with the principles of Just-In-Time (JIT) production, which aims to reduce inventory levels by tightly coupling upstream and downstream processes. However, with the emergence of Industry 4.0 and Industry 5.0, production environments have become increasingly dynamic, heterogeneous, and demand-driven (Sun et al., 2026). Under such conditions, traditional JIT strategies, which rely on relatively stable and predictable workflows, are often insufficient to handle frequent disruptions, customized production requirements, and complex resource interactions.

To address these challenges, recent studies have shifted toward more flexible synchronization mechanisms that explicitly incorporate spatial and temporal constraints into logistics planning and control. In this context, PL synchronization is no longer limited to static scheduling, but evolves into a dynamic orchestration problem involving real-time information updates and adaptive decision-making. Existing approaches typically integrate optimization models with heuristic or metaheuristic strategies to manage precedence relationships, resource constraints, and time-dependent task requirements (Su et al., 2026). In parallel, digital twin-based frameworks have been introduced to maintain a virtual representation of the physical system, enabling continuous monitoring and feedback-driven adjustments. Similarly, ticket-driven coordination mechanisms have been explored to facilitate decentralized decision-making and improve responsiveness in distributed production settings (Zhao et al., 2024).

3 Methodology

To facilitate PL synchronization in CPI, an agentic reasoning approach is proposed to coordinate multiple specialized agents with distinct functional roles, as shown in Figure 1. Specifically, a perception agent collects and parses heterogeneous shop-floor information from

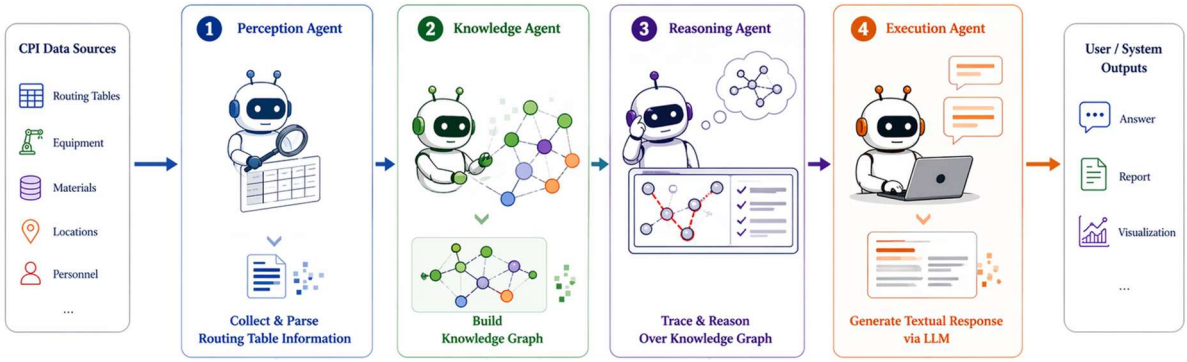


Figure 1: Agentic Reasoning Framework for PL Synchronization in CPI

CPI routing tables. A knowledge agent then organizes these structured records into a dynamic knowledge graph to represent transportation events. Subsequently, the reasoning agent performs traceability analysis of target objects and associated resources. Finally, an execution agent translates the reasoning outcomes into structured textual context and generates responses using a LLM.

3.1 Perception Agent for Information Aggregation

The perception agent is responsible for aggregating and structuring fragmented PL information. In the CPI framework, such information is natively maintained in the routing tables of CPI routers, which record resource allocation and transportation decisions across the shop floor. Each routing table entry encodes a logistics decision in a structured format, including the destination address, candidate next hop, resource attributes (e.g., type, specification, and quantity), and transportation mode. Compared with unstructured shop-floor data, these records inherently capture the flow of resources between locations, providing a concise representation of logistics interactions.

A key component underlying this representation is the CPI addressing scheme. Unlike conventional location descriptions based on Cartesian coordinates or GPS-style tuples, CPI assigns each physical location a hierarchical address. The high-order fields denote functional regions such as workshops or storage areas, while the low-order fields specify fine-grained spatial cells. As a result, address prefixes distinguish different operational areas, and complete addresses uniquely identify specific locations. This design embeds spatial and functional semantics into the addressing structure, enabling the system to differentiate between locations that are geographically close but operationally distinct, such as machining stations and buffer zones. By combining routing-table records with CPI addressing, the perception agent transforms heterogeneous shop-floor data into a unified and structured representation. This representation not only preserves the underlying logistics flow but also provides the necessary semantic and spatial context for subsequent knowledge construction and reasoning.

3.2 Knowledge Agent for Context Augmentation

Building upon the structured records produced by the perception agent, the knowledge agent augments them with semantic and temporal context by organizing the data into a dynamic knowledge graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. In this graph, nodes represent PL resources and sites, while edges encode relationships such as transfer, containment, and association. As illustrated in *Figure 2*, the entries of CPI routing tables are lifted from isolated records into a connected structure, where resource movements are explicitly modeled as directed links across different locations and processing stages.

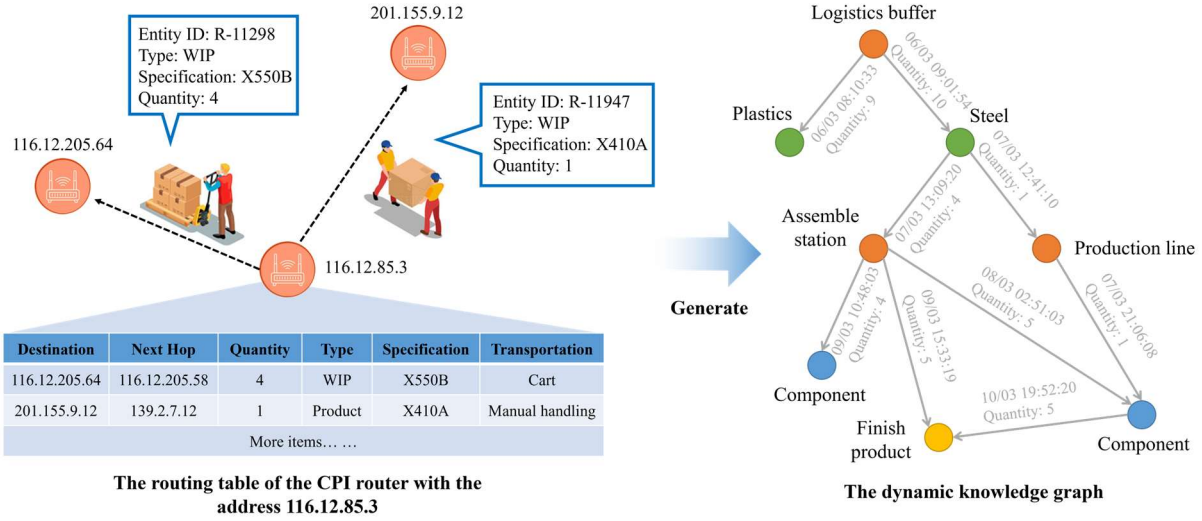


Figure 2: Construction of Dynamic Knowledge Graph

The knowledge graph integrates three complementary categories of information. Resource knowledge describes materials, work-in-progress items, components, and finished products. Allocation knowledge captures logistics events, including dispatching and transportation activities. Site knowledge provides the spatial context of operations, covering workshops, stations, buffers, and other functional areas. Each resource node is associated with attributes such as identifier, specification, status, quantity, and current location, enabling fine-grained traceability. Allocation edges further connect resources with timestamps, transportation modes, and source-destination pairs, thereby forming an explicit representation of the lifecycle of each logistics task. In parallel, site nodes contribute the spatial semantics necessary to interpret where these interactions occur and under what operational conditions. Through this organization, routing-table records are transformed from flat entries into a relational structure that supports reasoning over both interactions and context.

3.3 Reasoning Agent for Traceability Analysis

The reasoning agent aims to reconstruct the complete logistics trajectory of a target resource. The reasoning process consists of two stages, namely target identification and traceability analysis. Given a query Q describing the target resource, the query is first mapped into the semantic embedding space as

$$z_q = Embedding(Q), \tag{1}$$

where $Embedding(\cdot)$ denotes a text embedding model that transforms textual descriptions into dense semantic vectors, and z_q denotes the embedding of the query. For each node $v \in \mathcal{V}$ in the knowledge graph, its textual attributes x_v are also encoded as

$$z_v = Embedding(x_v). \tag{2}$$

The reasoning agent then computes the semantic similarity between the query and each node, and identifies the target node as

$$v^* = arg \max_{v \in \mathcal{V}} Sim(z_q, z_v), \tag{3}$$

where $Sim(\cdot,\cdot)$ denotes the similarity function. Starting from the identified target node v^*, the reasoning agent performs temporal traversal to recover its associated logistics history. Let $\mathcal{E}_t \subseteq$

$\mathcal{E}$ denote the set of temporal edges, where each edge $e(u, v)$ is associated with a timestamp $t(e)$ and an event extraction function $\phi(e)$. The traceability event set is defined as

$$\mathcal{F}(v^*) = \{\, \phi(e) \mid e(u,v) \in \mathcal{E}_t, u \text{ is reachable from } v^* \; via \; v \,\}. \tag{4}$$

During traversal, each newly discovered event is collected once, and the search is expanded along temporally connected nodes so as to preserve the causal order of logistics interactions.

Finally, all collected events are sorted by timestamp to produce the traceability timeline, which can be expressed as

$$\mathcal{T}(v^*) = \mathrm{Sort}(\mathcal{F}(v^*),\, t), \tag{5}$$

which describes the complete evolution of the target resource across different locations and operational stages. In this manner, the reasoning agent converts a high-level manager query into an ordered sequence of logistics events, thereby providing an interpretable basis for downstream analysis and response generation.

3.4 Execution Agent for Solution Generation

The execution agent focuses on transforming structured event sequences into textual context and generating interpretable responses. Specifically, each event $\phi(e)$ in $\mathcal{T}(v^*)$ is mapped into a natural language description through a textualization operator, and the resulting descriptions are concatenated, given by

$$\mathcal{X}_{req} = \mathcal{C}\mathrm{oncat}\Big(Textualize\big(\mathcal{T}(v^*)\big)\Big) \tag{6}$$

where $\mathcal{X}_{req}$ denotes the constructed textual context. This process preserves the temporal structure and key attributes of logistics events, allowing the downstream model to access both sequential dependencies and operational details.

The generated textual context is then provided to a pretrained LLM to produce task-specific outputs, represented as

$$p_\theta\big(Y \mid \mathcal{X}_{req}\big) = \prod_{i=1}^{L} p_\theta\big(y_i \mid y_{<i}, \mathcal{X}_{req}\big) \tag{7}$$

where θ denotes the model parameters, L is the output length, and $Y = (y_1, y_2, \ldots, y_L)$ represents the generated response sequence.

By grounding the generation process in the traceability timeline rather than the full knowledge graph, the execution agent ensures that the produced outputs are consistent with the observed logistics history. This design improves interpretability and reduces the risk of irrelevant or unsupported content generation. Consequently, the execution agent translates structured reasoning outcomes into coherent textual responses, completing the end-to-end pipeline for PL traceability and decision support.

4 Experiments

4.1 Experimental Setup

4.1.1 Dataset

The experiments are conducted on simulated data to emulate realistic shop-floor logistics scenarios under the CPI framework. Specifically, we construct a synthetic CPI environment

consisting of multiple workshops, stations, and storage areas, where each location is uniquely identified by a hierarchical CPI address. Based on this spatial structure, we generate 30 routing tables corresponding to different CPI routers deployed across the shop floor. Each routing table entry represents a logistics decision and contains the destination address, candidate next hop, resource category, specification, quantity demand, and transportation mode. The generation process follows a stochastic yet structured procedure: resource types and quantities are sampled from predefined distributions, while routing paths are constructed according to feasible transitions between functional areas. Temporal attributes are further assigned to each routing record to reflect realistic production sequences, including processing delays and transportation durations. In total, the dataset contains approximately 10,000 routing records.

All experiments are performed on a server equipped with an NVIDIA GeForce RTX 4090D GPU (24 GB memory) and an Intel Xeon Gold 6430 CPU (32 cores), running Ubuntu 20.04. The proposed framework is implemented in Python 3.8.10 using PyTorch 2.0.1. DeepSeek-R1-Distill-Qwen-14B, a pretrained open-source large language model, is adopted for textual reasoning and response generation.

4.1.2 Baseline Methods

To evaluate the effectiveness of the proposed reasoning framework, we compare it with three representative graph traversal methods for resource traceability:

(1) Breadth-First Search (BFS): BFS explores the graph in a level-wise manner, visiting all neighboring nodes before proceeding to the next depth. It provides comprehensive coverage of reachable nodes but may introduce redundant exploration in dense graphs.
(2) Depth-First Search (DFS): DFS traverses the graph by following each branch to its maximum depth before backtracking. It is suitable for discovering long dependency chains but may overlook alternative paths in early stages.
(3) Random Walk (RW): The random walk baseline selects the next node stochastically from the current node's neighbors, thereby representing unguided exploration without semantic or temporal constraints.

4.2 Performance Comparison

The time required by different algorithms to trace the same resource is presented in Figure 3. The results show that our proposed method achieves the shortest traceability analysis time, requiring only 68.3s, while RW, DFS, and BFS require 210.4s, 150.8s, and 130.1s, respectively. Compared with these three baseline methods, the proposed method reduces the analysis time by approximately 67.5%, 54.7%, and 47.5%. This improvement can be mainly attributed to the fact that traditional graph traversal methods perform traceability analysis in a relatively blind or exhaustive manner. RW depends heavily on random path selection, which may repeatedly visit low-value nodes and generate redundant search steps. DFS tends to explore long branches before backtracking, resulting in unnecessary traversal when the target-related records are distributed across multiple paths. BFS can ensure layer-wise expansion, but it still needs to scan a large number of neighboring nodes, leading to high computational overhead as the traceability network grows. In contrast, the proposed method incorporates more targeted traceability constraints and priority-guided search rules, allowing the analysis process to focus on records that are more strongly associated with material movement, storage location, and operation events. Therefore, irrelevant nodes and redundant paths can be effectively filtered out at an early stage, which reduces the search space and shortens the overall reasoning chain, thereby improving the efficiency of traceability analysis.

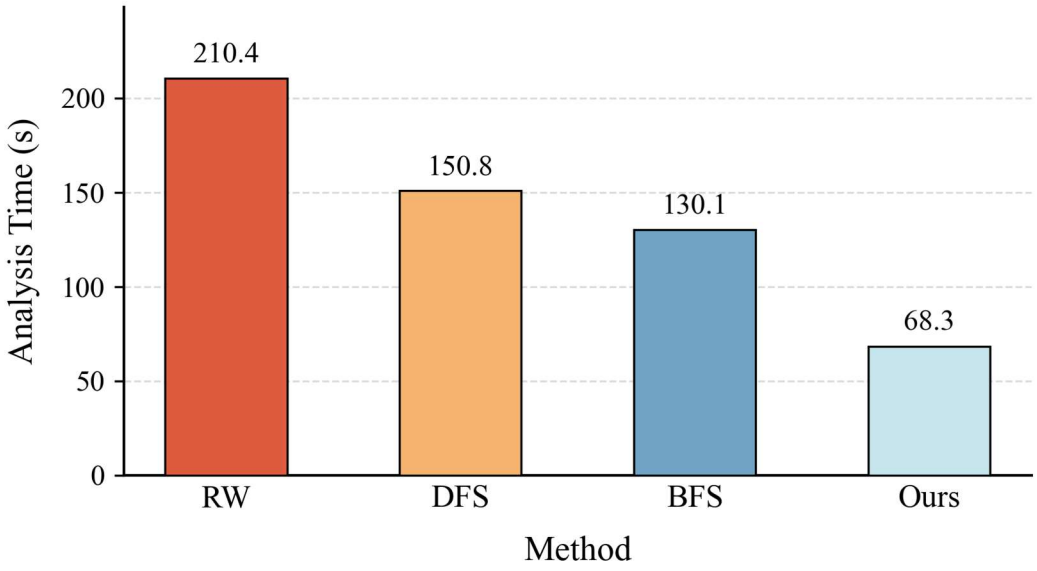


Figure 3: Computational Cost of Traceability Analysis Methods

4.3 Interaction Analysis

Figure 4 illustrates a representative interaction example driven by the proposed agentic reasoning approach. The example includes two typical tasks, namely status query and resource tracing, showing that the reasoning framework can support both forward and backward traceability requests through natural language interaction.

In the first query, the user requests confirmation of whether a batch of circuit boards is ready for dispatch. The response directly provides the current status, including quality inspection

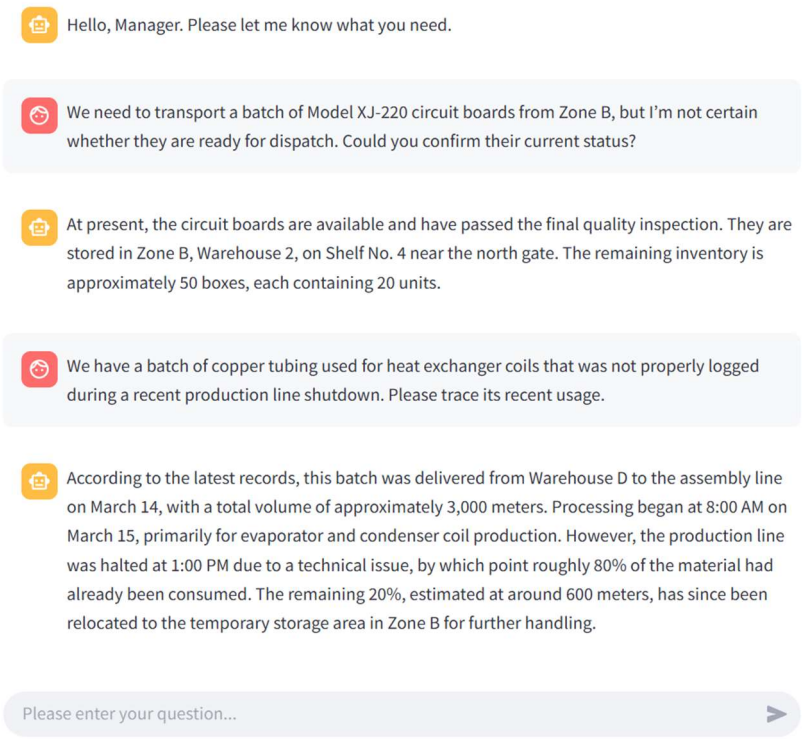


Figure 4: Interactive Traceability Query via Agentic Reasoning

results, storage location, and inventory details. This reflects that the reasoning process can effectively retrieve and integrate relevant information into a coherent result, enabling the user to obtain all necessary decision-related information in a single interaction.

In the second query, the user asks to trace the recent usage of a batch of copper tubing. The response reconstructs the material flow by integrating records distributed across different CPI routers, forming a complete trajectory that spans delivery, processing, interruption, and subsequent storage. Compared with the first query, this task requires correlating fragmented information sources rather than directly retrieving a single status, which highlights the capability of the reasoning approach to synthesize dispersed records and reveal the underlying evolution of the material flow.

5 Conclusion

This paper presents an agentic reasoning approach for traceability analysis within the CPI framework. The proposed approach enables unified reasoning over distributed data sources by coordinating multiple agents, allowing dispersed records to be effectively aggregated to support both status queries and resource tracing. In addition, it supports natural language interaction, thereby simplifying user access to traceability functions without relying on predefined query structures. Future work will focus on improving the scalability and robustness of the proposed approach. In particular, more efficient strategies for cross-router information retrieval and prioritization will be investigated to reduce overhead in large-scale settings.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (No. 52305557), Hong Kong Research Grants Council (No. 15203025, T32-707/22-N, R7036-22), Guangdong Basic and Applied Basic Research Foundation (No. 2024A1515011930), Research Institute for Advanced Manufacturing (RIAM) of The Hong Kong Polytechnic University (No. CDLU, CDLM), and the PAIR Seed Funding Scheme for Joint Research Projects with Leading Universities. The authors would also like to express their sincere gratitude to the CPI Research Group at The Hong Kong Polytechnic University for their valuable support and collaboration.

References

- Boysen, N., Briskorn, D., Montreuil, B., & Zey, L. (2025). The π-transportation problem: On the value of split transports for the Physical Internet concept. European Journal of Operational Research, 324(2), 629-643.
- He, S., Zhang, M., Wang, S., & Huang, G. Q. (2025). Channel structures and subscription strategies for AI-driven logistics data products. European Journal of Operational Research, 326(3), 597-614.
- Li, J., Fan, J., Zhao, Z., Wang, L., & Huang, G. Q. (2026). Cyber-Physical Computer in Action: Generative Recommendation via Spatial-Temporal Subgraph Reasoning for Production Logistics Orchestration. IEEE Transactions on Industrial Informatics.
- Li, J., Sun, M., Zhao, Z., & Huang, G. Q. (2025a). Cyber-Physical Internet Spatial-Temporal Reasoning for Production Logistics Resource Recommendation in Industry 4.0. Journal of Industrial Information Integration, 100935.
- Li, M., Qian, X., Li, M., Qu, T., & He, Z. (2025b). An order postponement strategy for multi-stage production-logistics synchronization towards zero-warehousing smart manufacturing. Journal of Manufacturing Systems, 82, 389-404.

- Lu, Z., Zhao, Z., & Huang, G. Q. (2026). STAR: Spatial-Temporal Attention Reasoning model for dynamic logistics network routing in Cyber-Physical Internet. Advanced Engineering Informatics, 69, 103830.
- Meriem, G. B., Pan, S., Klibi, W., & Selmi, M. H. (2026). LLM-Augmented Strategic Framing for PI-Oriented Supply Chain Networks. In Supply Chain and Logistics Intelligence in the Era of the Physical Internet: Bridging High-Tech and Low-Tech Solutions.
- Montreuil, B. (2011). Toward a Physical Internet: meeting the global logistics sustainability grand challenge. Logistics Research, 3(2), 71-87.
- Pan, S., Trentesaux, D., McFarlane, D., Montreuil, B., Ballot, E., & Huang, G. Q. (2021). Digital interoperability in logistics and supply chain management: state-of-the-art and research avenues towards Physical Internet. Computers in industry, 128, 103435.
- Peng, X., Ji, S., Thompson, R. G., & Zhang, L. (2021). Resilience planning for Physical Internet enabled hyperconnected production-inventory-distribution systems. Computers & Industrial Engineering, 158, 107413.
- Samadhiya, A., Agrawal, R., Kumar, A., & Garza-Reyes, J. A. (2023). Regenerating the logistics industry through the Physical Internet Paradigm: A systematic literature review and future research orchestration. Computers & industrial engineering, 178, 109150.
- Su, Z., Niu, Y., Li, J., Zhou, Y., Qiu, Y., & Zhao, Z. (2026). Decentralised schedule for mobile robots in Industry 4.0: an adaptive path planning strategy selection mechanism under dynamic environments. International Journal of Production Research, 64(7), 2616–2635.
- Sun, M., Li, J., Fan, J., Zhang, M., Zhao, Z., & Huang, G. Q. (2026). A double-sided bundle auction mechanism for collaborative additive manufacturing. Journal of Manufacturing Systems, 85, 473-486.
- Wu, H., Li, M., Yu, C., Ouyang, Z., Lai, K. H., Zhao, Z., ... & Huang, G. Q. (2025). Towards cyber-physical internet: A systematic review, fundamental model and future perspectives. Transportation Research Part E: Logistics and Transportation Review, 197, 104051.
- Zhao, Z., Xie, X., Chen, Q., Lin, P., Wu, W., Zhang, M., & Huang, G. Q. (2024). Graduation-inspired manufacturing system and synchronization mechanism for hybrid assembly cell lines. Computers & Industrial Engineering, 198, 110648.

Blockchain-Enabled IoT for Smart Food Traceability: Cost and Performance Insights

Hari Sannamuri
Email: harisannamuri0@gmail.com

***Abstract*—Ensuring transparency in food supply chains requires efficient traceability solutions. This study explores a blockchain-integrated IoT framework for tracking agricultural products from farm to fork. By evaluating multiple IoT devices and blockchain architectures, we assess energy consumption, processing efficiency, and cost-effectiveness in public and private blockchain networks. Our findings highlight the feasibility of low-cost IoT sensors in decentralized traceability systems while providing a detailed cost model for real-world implementation. The results offer valuable insights into optimizing blockchain-based food tracking solutions for enhanced reliability and scalability.**

***Index Terms*—Explainability, Entity Matching, Graph Neural Networks**

***Index Terms*—Blockchain, Internet of Things, Smart Contracts, Traceability, Supply Chain, Cost Modeling**

I. Introduction

The convergence of IoT and blockchain technologies is unlocking new classes of decentralized applications across diverse sectors [**?**]. Compact IoT devices can autonomously sense and transmit detailed information about environmental and process-related phenomena, all at relatively low cost [**?**]. When this data is integrated with blockchain, a decentralized ledger provides immutable and transparent logging, beneficial for inter-organizational business workflows [**?**]. Sectors like manufacturing [**?**], water management [**?**], telecommunications [**?**], remote sensing [**?**], and mining [**?**] have begun to embrace this fusion.

In agriculture, interest is surging in leveraging IoT and blockchain for traceability systems. Under the Industry 4.0 paradigm, supply chain visibility—such as event tracking (harvest, shipping, ownership change) and stakeholder identity (producer, regulator, consumer)—is critical [**?**]. Consumers increasingly seek complete transparency from farm to fork [**?**], prompting research into using digital ledgers to overcome outdated paper-based systems [**?**], [**?**].

Despite the enthusiasm, barriers remain. The economic feasibility of deploying such systems is a subject of ongoing discussion [**?**]. While performance and scalability challenges have been somewhat addressed [**?**], the associated infrastructure expenses—especially for constrained environments—require further scrutiny. Agriculture poses unique challenges due to the need for inexpensive, energy-efficient devices that must function reliably over long durations [**?**].

To address these gaps, this paper extends prior work [**?**] by offering a comprehensive evaluation of six constrained IoT platforms when interfacing with Ethereum and Sawtooth blockchains. It further introduces a cost model to assess public versus private blockchain infrastructure. By applying this model, practitioners can better anticipate system expenses and make informed architectural choices.

II. Related Work

A growing body of literature explores blockchain-enabled traceability in agri-food systems [**?**], [**?**]. However, investigations into the computational and energy demands imposed on IoT devices remain limited.

One of the pioneering efforts in this field was conducted by Tian who utilized RFID and blockchain to improve food safety in China. He later expanded this work to incorporate hazard analysis and critical control point (HACCP) models. Lin et al. examined how blockchain and IoT could integrate with ERP systems in agriculture, while Kim et al. demonstrated traceability using Ethereum smart contracts and GS1 standards.

Salah et al. implemented a traceability mechanism for soybean supply chains and analyzed Ethereum smart contract usage in detail. Caro et al. [**?**] provided a practical implementation for blockchain-based food traceability using IoT. Yet, all these studies neglected the performance bottlenecks of low-power sensors, which often struggle with the heavy computational demands of blockchain [**?**].

Infrastructure costs have also been underexplored. Rimba et al.compared cloud-based and blockchain-based business process execution, while Demir et al proposed a financial assessment framework for blockchain applications. Wu et al.focused on Ethereum smart contract costs, particularly gas consumption.

To summarize, although research supports blockchain's applicability to agri-food traceability, key aspects such as device-level resource impact and infrastructure cost modeling remain inadequately studied.

III. Case Study: Extra Virgin Olive Oil Traceability

To ground our evaluation, we implemented a simplified event-based traceability model for extra virgin olive oil, mirroring existing agri-food models. The process spans four stages: farming, manufacturing, transportation, and market. Figure 1 illustrates this progression.

article tikz

IoT sensors are deployed throughout each phase to collect environmental and process data. However, due to blockchain storage limitations, only select events (e.g., average monthly

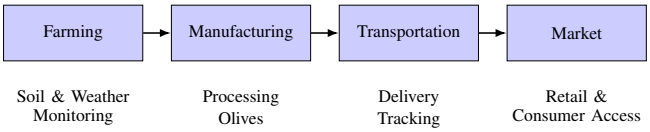


Fig. 1. Simplified farm-to-fork traceability flow for extra virgin olive oil.

soil data, temperature anomalies during transit) are recorded on-chain, while detailed logs are offloaded to cloud platforms or distributed file systems.

For example, at the farming stage, air quality and soil data are gathered per olive tree but only summarized values are stored on-chain. During manufacturing, production metrics per batch are logged. For transportation, location and temperature anomalies are recorded. The retailer stage logs irregularities before final sale. Consumers can trace product history through immutable blockchain entries.

IV. Infrastructure Cost Modeling

To evaluate financial feasibility, we modeled infrastructure expenses for both public and private blockchain implementations. The system comprises multiple actors—human participants (A) and IoT devices (D)—who interact via signed blockchain transactions (T_x).

A. Public Blockchain Costs

In public blockchains (e.g., Ethereum), costs are primarily linked to transactions that alter state (e.g., contract creation, token transfer). We divide costs into:

- **Bootstrap Cost (C_B)** for deploying smart contracts (e.g., device contracts T_x^T, application logic T_x^A).
- **Operational Cost (C_O)** for routine transactions: device interactions T_x^I and value transfers T_x^V.

Thus, the monthly cost is:

$$C_{\text{pub}}(m) = \begin{cases} C_B = C_{T_x^T} + C_{T_x^A}, & \text{if } m = 0 \\ C_O = C_{T_x^I} + C_{T_x^V}, & \text{if } m > 0 \end{cases}$$

Each transaction type $i \in \{T, A, I, V\}$ incurs a cost:

$$C_{T_x^i} = CC_i \cdot \mu \cdot P(m) \cdot Q_{T_x^i}(m)$$

Where:

- CC_i: Computational cost in gas
- μ: Processing factor (gas price)
- $P(m)$: Crypto price (USD/month)
- $Q_{T_x^i}(m)$: Number of transactions per month

B. Private Blockchain Costs

Unlike public networks, private blockchain costs are node-based. The total monthly cost $C_{\text{pri}}(m)$ consists of:

- Initial setup cost: $N \cdot C_{\text{node}}$
- Monthly operational cost: $N \cdot C_{\text{op}}$, where $C_{\text{op}} = F_o \cdot C_{\text{node}}$

The node count N is a function of total actors A and a trust factor F_t:

$$N = A \cdot (1 - F_t)$$

This model demonstrates that private networks have fixed infrastructure costs, while public blockchains scale with transaction volume and token price.

TABLE I
Public Blockchain Cost Parameters

Parameter	Description
A	Number of actors
D	Number of IoT devices
$P(m)$	Crypto price in USD at month m
μ	Processing cost (gas price)
CC_T, CC_A, CC_I, CC_V	Gas costs of transactions
$Q_{T_x^i}$	Monthly transaction count per type

TABLE II
Private Blockchain Cost Parameters

Parameter	Description
C_{node}	Hardware cost per node
F_o	Operation cost percentage
F_t	Trust factor among actors
A	Number of actors

V. Cost Modeling for Blockchain Infrastructure

A. Transaction Quantification in Public Blockchains

The total number of transactions for each type, denoted as $Q_{Tx_i}(m)$ for month m, is categorized as follows:

$$Q_{Tx_i}(m) = \begin{cases} D, & \text{if } i = T \\ N_A, & \text{if } i = A \\ N_I, & \text{if } i = I \\ N_V, & \text{if } i = V \end{cases} \tag{1}$$

Where:

- D is the number of IoT devices (each associated with a TwinContract).
- N_A represents the total number of distributed applications (AppContracts).
- N_I denotes the total number of sensor events captured and written to the blockchain.
- N_V accounts for value transfer events among stakeholders (e.g., from producers to transporters).

Thus, estimating the infrastructure expenditure of a public blockchain-based traceability system requires input for the parameters summarized in Table III.

B. Private Blockchain Infrastructure Costs

In contrast to public blockchains, where cost scales with transaction volume, private blockchain systems incur costs based primarily on infrastructure—particularly the number of nodes in the network. These costs depend on system performance goals (latency, throughput) and trust distribution among stakeholders.

The monthly cost $C_{Pri}(m)$ is defined as:

$$C_{Pri}(m) = \begin{cases} N \cdot C_{node}, & \text{if } m = 0 \\ N \cdot C_{op}, & \text{if } m > 0 \end{cases} \tag{2}$$

TABLE III
PARAMETERS FOR PUBLIC BLOCKCHAIN COST MODEL

Parameter	Description
A	Number of stakeholders
D	Number of IoT devices
$P(m)$	Cryptocurrency price in month m
μ	Processing time factor (gas cost)
CC_T	Computation cost of SmartTwin Tx
CC_A	Computation cost of TwinApp Tx
CC_I	Computation cost of interaction Tx
CC_V	Computation cost of value transfer Tx
N_A	Total AppContracts deployed
N_I	Number of IoT readings on-chain
N_V	Number of inter-actor value transfers

Where:

- C_{node} is the procurement cost per node.
- C_{op} is the operational cost per node per month.
- N is the number of nodes deployed.

Operational cost is assumed to be a fixed proportion of node cost:

$$C_{op} = F_o \cdot C_{node} \tag{3}$$

The node count N is derived from the actor count A and the trust factor F_t:

$$N = A \cdot (1 - F_t) \tag{4}$$

TABLE IV
PARAMETERS FOR PRIVATE BLOCKCHAIN COST MODEL

Parameter	Description
A	Number of actors
C_{node}	Cost per node
F_t	Trust coefficient (0–1)
F_o	Monthly operational factor

VI. PERFORMANCE ASSESSMENT

A. Experimental Setup

To assess resource usage and computational overhead, we conducted tests across six microcontroller platforms, ranging from 8-bit AVR to 32-bit ARM and ESP32. Hardware specs are detailed in Table V.

TABLE V
IOT DEVICES FOR PERFORMANCE EVALUATION

Device	MCU	Arch.	MHz	Flash (kB)	RAM (kB)
UNO	ATMega328P	8-bit AVR	16	32	2
EVERY	ATMega4809	8-bit AVR	20	48	6
L031	STM32L031	Cortex-M0+	32	32	8
F303	STM32F303	Cortex-M4F	72	64	12
L452	STM32L452	Cortex-M4	80	512	96
ESP32	WRover-E	ESP32	80	1024	320

Ethereum nodes used Geth v1.10.1; Sawtooth used v1.2.6. All nodes were hosted on isolated OpenStack VMs (4 GB RAM, 4 vCPUs, 20 GB SSD). Python 3.6 scripts handled contract deployment.

B. Device Memory and Storage Footprint

Program (flash) and memory consumption per device module component is shown in Table VI and VII respectively.

TABLE VI
DISK USAGE (BYTES)

Device	Avail.	Sens.	Eth	Saw	Comms
UNO	32256	2014	26840	35214	2202
EVERY	49152	1956	27428	35735	3105
F303	65536	3040	29940	33404	14240
L031	32768	6168	30796	34432	14724
L452	524288	3032	33120	36568	17336
ESP32	1310720	1496	289318	293674	267270

TABLE VII
RAM USAGE (BYTES)

Device	Avail.	Sens.	Eth	Saw	Comms
UNO	2048	207	1158	3214	188
EVERY	6144	196	611	2435	177
F303	12288	928	1540	3124	908
L031	8192	896	1508	3102	876
L452	163840	936	1560	3144	916
ESP32	327680	13660	14996	16292	13612

C. Transaction Time and Size

Each device performed 100 trials to create, sign, and transmit transactions. Table VIII reports average time and transaction size.

TABLE VIII
PROCESSING TIME (MS) AND TRANSACTION SIZE (BYTES)

Device	Sens.	Eth	Saw	Size (Bytes)
EVERY	32	4170	8388	134 / 925
F303	32	208	421	134 / 925
L452	32	160	278	134 / 925
ESP32	32	83	146	134 / 925

D. Power Consumption

Using the Otii Arc tool, we measured current draw at 5V during idle and working states. Table IX and X present power metrics.

TABLE IX
ETHEREUM POWER CONSUMPTION

Device	Idle (J)	Work (J)	Idle Time (s)	Work Time (s)
UNO	0.86	0.73	5.0	4.30
EVERY	0.73	0.63	5.0	4.22
F303	1.63	0.08	5.0	0.26
L452	1.27	0.04	5.0	0.19
ESP32	1.06	0.03	5.0	0.13

TABLE X
SAWTOOTH POWER CONSUMPTION

Device	Idle (J)	Work (J)	Idle Time (s)	Work Time (s)
EVERY	0.73	1.28	5.0	8.58
F303	1.62	0.2	5.0	0.61
L452	1.35	0.12	5.0	0.47
ESP32	1.06	0.09	5.0	0.34

E. Energy Budget Breakdown

We computed hourly energy budgets assuming a 5-minute sampling rate, with only one in 12 events stored on-chain. The total energy budget is:

$$E_{budget} = E_{idle} + E_{sensing} + E_{comms} + E_{blockchain} \quad (5)$$

Figures 6 show the energy distribution for EVERY (highest usage) and ESP32 (lowest usage).

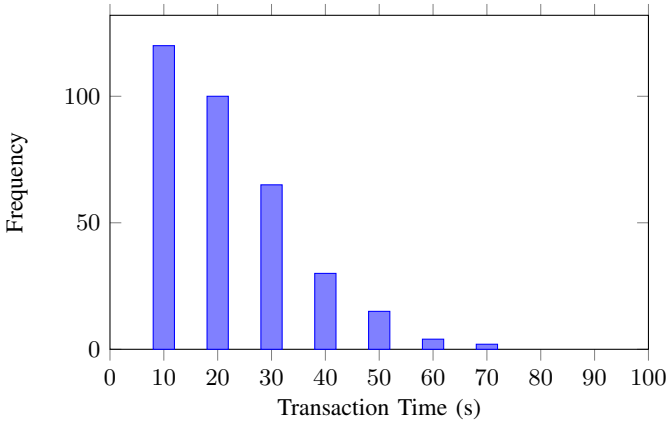


Fig. 2. Distribution of Ethereum Ropsten transaction confirmation times.

Results indicate that idle time dominates energy usage, with blockchain-related operations adding only a marginal increase to the overall budget.

VII. COST EVALUATION OF BLOCKCHAIN-BASED TRACEABILITY

To determine the financial feasibility of deploying the traceability system on both Ethereum and Sawtooth platforms, we evaluate the associated costs using the parameters previously defined in Sections IV-A and IV-B. As outlined in Section III, the architecture involves six stakeholders: provider, producer, processor, distributor, retailer, and end consumer. However, only four IoT-enabled stages (producer, processor, distributor, and retailer) are considered for device integration.

We limit the scope to a single product (a bottle of olive oil) within a transportation-focused traceability use case. However, the model we present is flexible and adaptable to a variety of product categories and deployment sizes by modifying the input parameters.

A. System Parameters and Configuration

The application is configured with $A = 6$ actors and $D = 4$ IoT devices, resulting in the need to deploy four *TwinContracts* (one per device) and a single *AppContract* to manage transportation. The number of IoT-based data logs is set to $N_I = 10$, assuming 10 significant sensor events across production and transport phases. The value transfer transactions between actors total $N_V = 4$, capturing exchanges such as provider-to-producer or transporter-to-retailer.

In the Ethereum implementation, each transaction incurs gas fees, quantified through computational gas cost and the gas price. The gas price μ is set to 60 Gwei ($\mu = 60 \times 10^{-9}$), reflecting a midrange fee that balances speed and cost. Table XI lists the gas requirements for each transaction type in our system.

TABLE XI
GAS CONSUMPTION FOR ETHEREUM TRANSACTIONS

Transaction Type	Gas Required
Create Sensor (`TwinSC`)	557,364
Create App Contract (`AppSC`)	3,343,572
Register Sensor Data (`setValue`)	44,468
Transfer Token (`transferToken`)	63,681

We fix the ETH/USD exchange rate to $P(m) = 1087$ based on historical average values between 2019 and 2021 [**?**]. Node cost for Sawtooth is estimated at $C_{node} = 300$ USD, based on hardware specifications that include a multi-core CPU, 4+ GB RAM, and 320 GB of SSD storage. We assume an operational factor $F_o = 0.4$ and zero trust among actors ($F_t = 0$), implying every actor requires their own node.

TABLE XII
APPLICATION CONFIGURATION PARAMETERS

Symbol	Description	Value
A	Number of actors	6
D	Number of IoT devices	4
$P(m)$	Cryptocurrency price	1087 USD
μ	Gas price	60×10^{-9}
CC_T	Gas for TwinContract	557,364
CC_A	Gas for AppContract	3,343,572
CC_I	Gas for interaction transactions	44,468
CC_V	Gas for value transfers	63,681
N_A	Number of AppContracts	1
N_I	Number of IoT measurements	10
N_V	Number of value transfers	4
C_{node}	Node cost (USD)	300
F_o	Operation cost factor	0.4
F_t	Trust factor among actors	0

B. Comparison of Infrastructure Costs

Based on the configuration above, we evaluate the total deployment cost (bootstrap and operational) for Ethereum and Sawtooth. For Ethereum, we simulate costs for 1, 25, and 50 product traces, keeping bootstrap costs constant. The comparison is shown in Figure 3. pgfplots pgfplotstable

As shown, tracing a single item through Ethereum costs approximately 409 USD. This figure represents around 16%

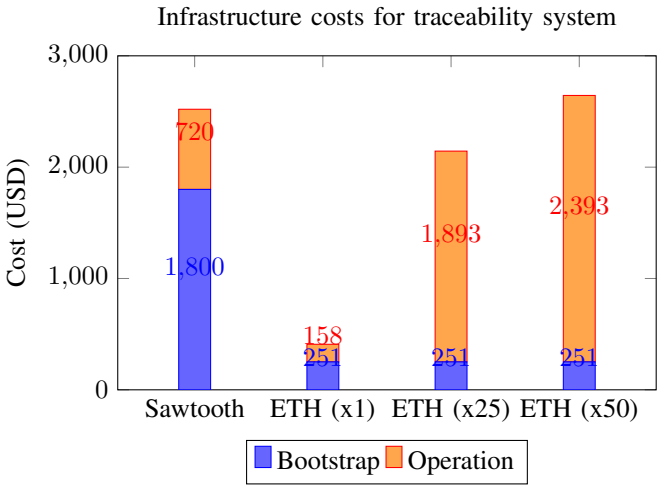


Fig. 3. Infrastructure costs (bootstrap and operation phases) for the traceability application on Sawtooth and Ethereum implementations.

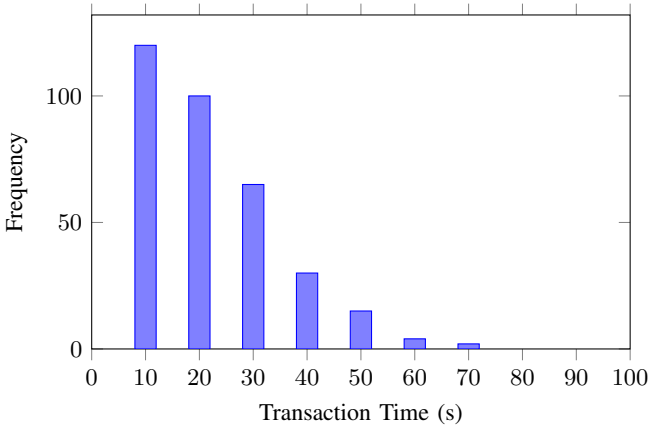


Fig. 4. Distribution of Ethereum Ropsten transaction confirmation times.

of the total cost of a private Sawtooth-based network (2520 USD). Notably, the same amount spent on Sawtooth could enable tracking of nearly 50 products via Ethereum.

However, Ethereum's cost is sensitive to network conditions. For example, at a low ETH price of 307 USD and 30 Gwei gas, the cost drops to just 58 USD. Conversely, during price surges in 2021 (ETH = 2774 USD, gas ¿ 100 Gwei), the cost can spike to 1740 USD per item.

Private blockchains offer advantages in performance and predictability. Studies show Sawtooth can reach throughput levels of up to 2300 tx/s, with optimal performance below 1000 tx/s [**?**]. Since transaction count doesn't affect cost, such infrastructure can handle thousands of items concurrently.

Ethereum's open, decentralized architecture—with roughly 8000 active nodes [**?**]—offers strong resilience, transparency, and integration potential for systems with untrusted, dynamic participants.

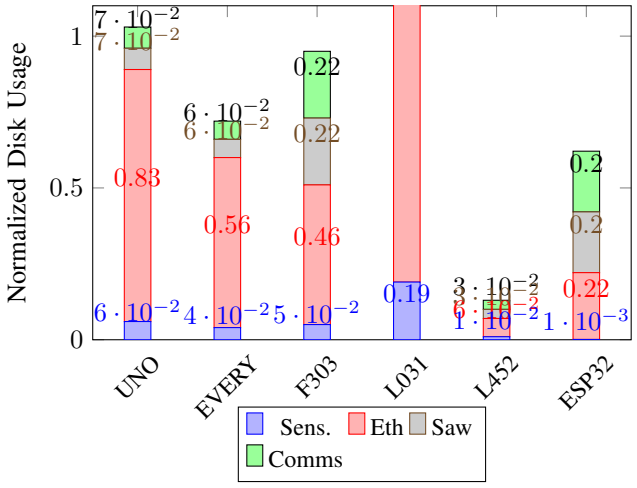


Fig. 5. Normalized disk usage of device module components across IoT boards.

C. Transaction Latency Analysis

To assess network responsiveness, we benchmarked transaction latency using the `setValue()` operation on Ethereum's Ropsten testnet. Transactions were sent every 30 minutes for one week (336 in total) using a gas price of 60 Gwei.

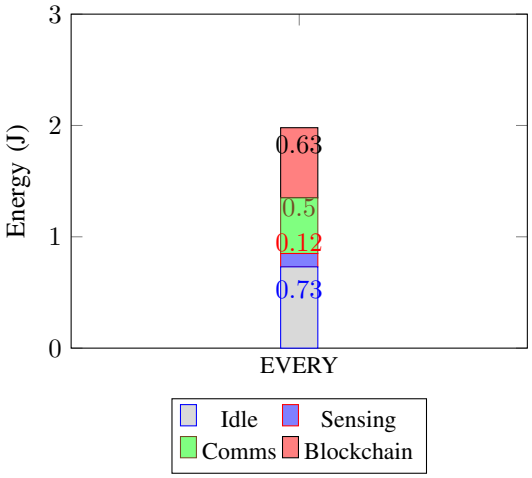


Fig. 6. Energy budget distribution for Arduino Nano EVERY.

Figure 4 shows that the average transaction was confirmed within 17 seconds. Only a small fraction took more than one minute. In contrast, Sawtooth processed similar transactions in under 1 second, showcasing the speed advantage of private deployments.

VIII. Discussion and Future Work

This study systematically examined the technical and economic feasibility of deploying blockchain-enabled traceability using constrained IoT devices. Through a detailed implementation of an event-driven architecture and benchmarking of Ethereum and Sawtooth platforms, the study drew several conclusions:

- Ethereum's lightweight footprint allows it to run on resource-constrained 8-bit boards (e.g., 2 KB RAM, 32 KB flash), making it accessible for low-cost deployments.
- Sawtooth, though more feature-rich, incurs higher memory and storage costs due to batching and encoding schemes. Its transactions are also approximately six times larger.
- Energy-wise, blockchain operations accounted for less than 0.25% of the hourly power budget, with idle states dominating consumption.
- Cost modeling reveals Ethereum to be more scalable for small-to-medium deployments, while Sawtooth's costs remain constant regardless of transaction volume.

Our results suggest that constrained devices can function as first-class actors in blockchain systems, enabling direct, secure data logging and verifiability in traceability systems.

Future Directions:

- Porting the IoT blockchain library to support other platforms (e.g., Hyperledger Fabric, Polkadot).
- Integrating with low-power wide-area networks (LPWAN) such as NB-IoT, LoRaWAN, and 5G.
- Exploring hybrid architectures combining private and public blockchain layers.

These efforts will enhance the robustness, energy efficiency, and scalability of blockchain-based IoT traceability systems in the agri-food supply chain and beyond.

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE
The Hong Kong Polytechnic University 香港理工大學

A Cyber-Physical-Internet-Based Self-Organizing Path Planning Method for Production Logistics

Haoran Liu[1, 2] , Bin Huang[1, 2] , Zhiheng Zhao[1, 2, *] and George Q. Huang[1, 2]

1. Department of Industrial and Systems Engineering, The Hong Kong Polytechnic University, Hong Kong, China
2. Research Institute for Advance Manufacturing, The Hong Kong Polytechnic University, Hong Kong, China

Corresponding author: Zhiheng Zhao (zhiheng.zhao@polyu.edu.hk)

Abstract: *Production logistics are increasingly operating on structured path networks that connect storage areas, workstations, charging points, dispatch interfaces, and intermediate hubs. Under dynamic transportation demands, centralized dispatching and shortest-path control may become unreliable due to path congestion, vehicle availability, and constantly changing local service conditions. This paper proposes a decentralized production logistics architecture based on the Cyber-Physical Internet (CPI), inspired by the self-organizing mechanisms observed in ant recruitment and transport networks. In the proposed framework, each logistics node is modeled as a Cyber-Physical Computer (CPC), which maintains local address and path information, monitors local resource status, and provides positioning services to nearby vehicles. On top of this node-centric CPI layer, we construct a two-level reduced-order transport model that dynamically combines vehicle states with virtual pheromone path reinforcement under a fixed path infrastructure. Path activation is driven by virtual pheromone reinforcement, path utility, travel cost, congestion, location confidence, and path adjustment penalties. Results from simulations indicate that the proposed CPI strategy generates stable path convergence under single-task conditions, produces nearly symmetric dual-path coexistence under balanced conditions, and yields persistent path preferences under asymmetric conditions.*

Keywords: *Cyber-Physical Internet; production logistics; localization; self-organization; routing network.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☐ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☐ *In-Person presentation*

1 Introduction

Over the past two decades, swarm robotic and multi-agent transport systems have attracted increasing attention in intralogistics for collaborative exploration, and decentralized material handling. Although this decentralized structure offers numerous advantages and research has

made some progress, overall, its adoption in practical applications remains slow. One key limitation is that while collective behavior occurs on a macroscopic scale, it arises from the self-organization of robots developed at the microscopic level and a challenge widely recognized as a major hurdle for swarm robotics application in intralogistics (Zhu et al., 2024). As mobile transporters, especially autonomous mobile robots (AMRs), automated guided vehicles (AGVs), and related autonomy carriers, must be assigned to dynamic transport tasks while respecting congestion, aisle availability, and service priorities.

This motivates a routing-oriented view of production logistics. Similar to packet-routing infrastructures, an intralogistics node can maintain local addressing, forwarding preferences, queue status, service capacity, and neighborhood information of nearby transporters. Inspired by the concept of connecting digital and physical spaces of Cyber-Physical Internet (CPI), this paper considers a node-centric intralogistics architecture in which cyber services and physical transportation are tightly integrated (Ouyang et al., 2025). Furthermore, the node is no longer merely a geometric waypoint. Instead, it becomes a cyber-physical decision point that couples information exchange, physical movement, and local service logic. In this work, we formalize that node as a Cyber-Physical Computer (CPC) in the proposed framework, which is a CPI router that integrates local routing intelligence, resource-state awareness, and positioning support for nearby transporters.

To design decentralized coordination laws for such a network, it is natural to draw inspiration from social insects, particularly ants. Ant colonies generate organized transport behavior through local interaction, exploration, and positive feedback without centralized supervision. This mechanism is directly relevant to intralogistics routing, where transport decisions must continuously balance travel cost, service value, and congestion. In addition, decentralized routing in intralogistics requires real-time spatial awareness. For this reason, the proposed CPC nodes are further equipped with BLE-based localization capability (Vaccari et al., 2024; Li et al., 2024), allowing routing decisions to exploit transporter position and local occupancy estimates online. Based on these observations, our contribution can be summarized as follows:

1. We introduce CPCs as nodes in CPI, where routing intelligence, BLE localization, and physical-digital spatial alignment are integrated on the CPCs to improve the real-time information exchange and spatial awareness capabilities of internal logistics.

2. We propose a reduced-order self-organizing route-allocation model to decentralize the organization of AMR transportation within intralogistics.

3. We conduct a simulation to show the concentration, coexistence, and preference of the model.

The remainder of the paper is organized as follows. Section 2 reviews the most relevant literature streams and sharpens the research gap. Section 3 formulates the CPI intralogistics problem and presents the proposed CPI-based model. Section 4 summarizes the experimental setup and illustrative results. Section 5 concludes the paper.

2 Related Work

2.1 Self-organization and transportation in intralogistics

Ant foraging and transport models are canonical examples of distributed coordination. Classical trail-formation studies emphasize exploration, pheromone laying, and positive reinforcement as mechanisms that transform local decisions into global route selection. To emerge multi-

source multi-sink transport-network structures, Lecheval et al. (2021) explicitly reformulated recruitment and trail departure as a differential-equation model dependent on both quality and distance and then connected those local allocation rules. In addition, some recent literature has further explored fleet formation, scheduling, and path planning for automated guided vehicles (AGVs) in internal logistics (Jun et al., 2022). Fragapane et al. (2021) also examined the transition from centrally controlled AGV systems to more autonomous and decentralized autonomous mobile robot (AMR) operations. These reviews justify the need for adaptive routing methods that can exploit dynamic local information.

2.2 CPI network and spatial observation

The CPI concept extends packet routing in communication networks to logistics systems (Montreuil, 2011), integrating network information, physical flows, and routing concepts. For example, recent CPI research in the logistics field (Ouyang et al., 2025) explicitly adopted a CPI-driven logistics infrastructure integration framework, which includes a hierarchical network and transport mechanisms, as well as routing concepts. In this network, logistics infrastructure is no longer viewed merely as passive geometric entities but rather as a network of decision points that integrates physical space with digital systems and possesses routing management capabilities. However, existing CPI research still focuses more on scheduling frameworks and protocol analogies rather than low-dimensional self-organizing routing dynamics. Indoor positioning systems, on the other hand, provide the necessary support for local routing information. A review by Vaccari et al. (2024) consistently highlights that RFID, UWB, Wi-Fi, vision, and BLE are mainstream technologies in logistics-oriented positioning systems, while Li et al. (2024) note that indoor positioning is regarded as a foundational capability for resource, production, and safety management in Industry 4.0. From an Industry 5.0 perspective, indoor positioning is not merely a navigation tool but an enabling layer for resource tracking, production visualization, and safety management.

2.3 Research gap

To the best of our knowledge, these threads have not yet been integrated into a single node-centric model that couples compartmental vehicle-state dynamics, route-memory reinforcement, BLE-based localization confidence, and adjustment-aware routing over a fixed production-logistics graph.

3 Methodology

3.1 Problem formulation

Consider an production logistics network represented by a graph $G = (V, E)$, where each node $v_i \in V$ denotes a logistics service point such as a storage interface, picking station, workstation, charger, consolidation point, or intersection router. Each edge $e_{ij} \in E$ represents a traversable corridor or movement link. Mobile fleet $\mathcal{M}$ consists mainly of AMRs and other automated carriers.

In this network, each node is modeled as a CPC i, which stores an address a_i, a neighbor set $\mathcal{N}_i$, a routing table $R_i(t)$, a local service-state vector $S_i(t)$, and a localization service interface $L_i(t)$. CPI network information is distributed across routers and updated using local observations. The dynamic state can be expressed as:

$$X_i(t) = [a_i \quad \mathcal{N}_i \quad R_i(t) \quad S_i(t) \quad L_i(t)]^{\top}. \tag{1}$$

Transport requests are generated dynamically. A request k is defined by an origin node o_k, a destination node d_k, a priority π_k, and a transport demand w_k:

$$r_k = (o_k, d_k, \pi_k, w_k). \quad (2)$$

For each request, a vehicle must be assigned to one of several feasible paths connecting o_k to d_k through the routing graph. In contrast to conventional shortest-path assignment, the path choice here depends on travel cost, local congestion, estimated route utility, localization reliability, and the need to avoid excessive deviation from the baseline network.

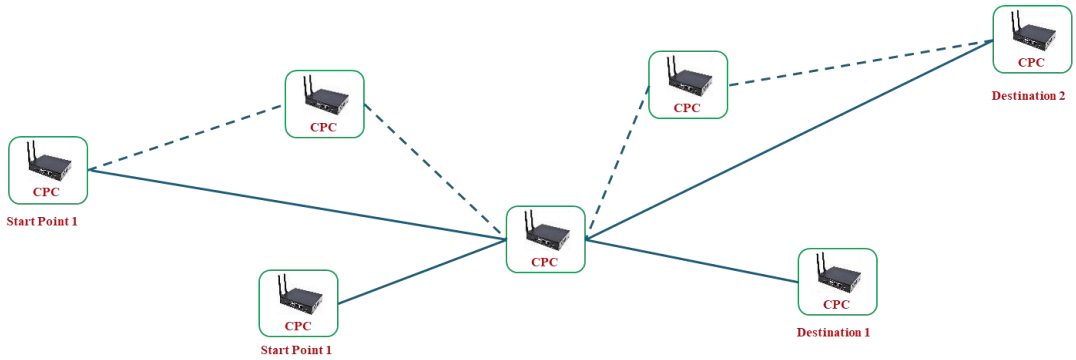


Figure 1: Illustrative warehouse graph used for the CPI routing study. Solid edges denote the shortest path from start point to the corresponding destination. Dashed edges denote the alternative path from start point to the corresponding destination.

Let T denote average request completion time, C a congestion-related cost, A the route-adjustment degree with respect to the baseline network, and U a service-utility term derived from transport priorities. The design goal is to minimize average request completion time T, congestion-related cost C, and route-adjustment degree A, while maintaining high transport utility $U$and stable decentralized route allocation. The problem is therefore not only a routing problem but also a distributed state-estimation and adjustment problem. The routing decision at each node should react to changing local conditions while preserving the operational coherence of the network.

3.2 CPI-driven spatial representation and transportation Model

During system operation, CPCs receive RSSI measurements from tags carried by nearby transport vehicles and mobile resources, and use these measurements to estimate vehicle proximity or location. Through collaboration among multiple CPCs, vehicle location estimation is achieved to support navigation, along with occupancy estimation, estimated arrival time inference, conflict prediction, and local congestion quantification for route selection. The framework is show as Figure 2.

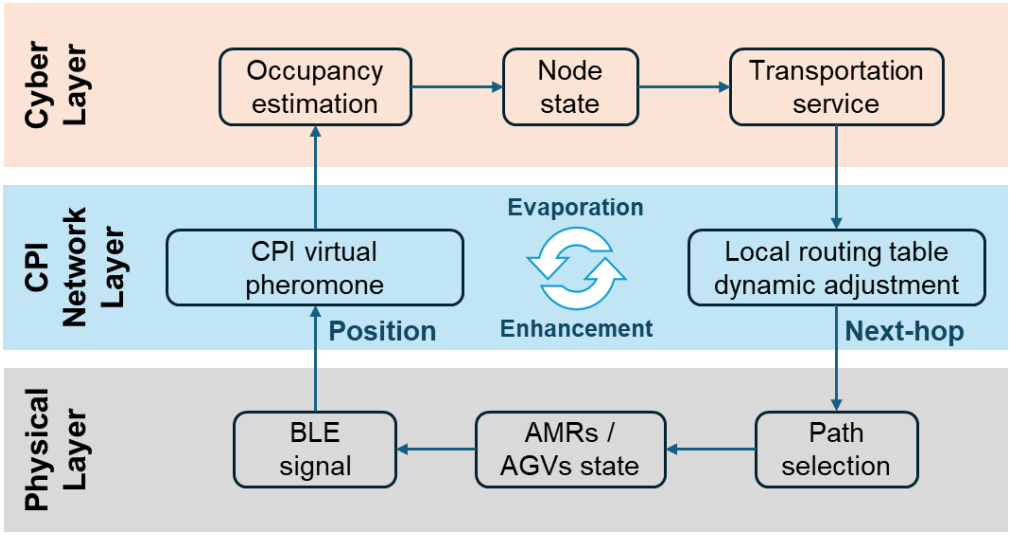


Figure 2: Node-centric CPI architecture. Each router integrates cyber state, BLE sensing, and physical transport coordination. Virtual pheromone reinforcement couples the three layers.

The state of the vehicles and their temporal-spatial represtation are formulated as a two-level reduced-order model. The first level is a compartmental vehicle-state model. The second level is a route-memory and next-hop selection model defined over the CPI graph through virtual-pheromone memory and softmax next-hop decisions. Let $U(t)$ denote the number of idle vehicles available for assignment. In route r, $C_r(t)$ denote the number of vehicles actively committed to route r, and $Q_r(t)$ denote the number of support vehicles, which represent queued reserve vehicles, repositioning traffic, or vehicles shadowing a frequently used route. Notable, we propose an industrial transport-cycle delay term τ_r of route r, which can include loading, loaded travel, unloading, and return or repositioning time. The state equations are:

$$dU(t)/dt = -\Sigma_{r\in R}\Lambda_r(t)U(t) + \Sigma_{r\in R}\mu_r C_r(t-\tau_r) + \Sigma_{r\in R}\nu_r Q_r(t-\tau_r), \quad (3)$$

$$dC_r(t)/dt = \Lambda_r(t)U(t) - \mu_r C_r(t-\tau_r) + \sigma_r Q_r(t-\tau_r), \quad (4)$$

$$dQ_r(t)/dt = \eta_r C_r(t-\tau_r) - \sigma_r Q_r(t-\tau_r) - \nu_r Q_r(t). \quad (5)$$

Equation (3) represents vehicles transforming from the idle pool when routes are activated and returns them when active or support tasks finish. Equation (4) governs active transport commitment on each candidate route. Equation (5) captures the secondary support population that stabilizes frequently used paths through positive feedback. The fleet size remains conserved in the reduced-order model:

$$U(t) + \Sigma_r[C_r(t) + Q_r(t)] = M. \quad (6)$$

The route activation rate $\Lambda_r(t)$ encode route history, task utility, distance, congestion, localization confidence, and adjustment penalty as:

$$\Lambda_r(t) = sig\left(\beta_\varphi \varphi_r(t) + \beta_q q_r - \beta_d d_r - \beta_\kappa \kappa_r(t) - \beta_a a_r + \beta_\ell \ell_r(t)\right), \quad (7)$$

where $\varphi_r(t)$ is virtual pheromone, q_r is task utility or priority, d_r is travel cost, $\kappa_r(t)$ is the estimated congestion level, a_r is the adjustment penalty relative to the reference route, and $\ell_r(t)$ is the BLE-based localization confidence.

To enhance the temporal capability of the system, we propose a memory mechanism that each path or edge therefore carries a virtual pheromone variable that increases with use and decreases through evaporation:

$$d\varphi_r(t)/dt = \rho T_r(t) - \delta\varphi_r(t), \tag{8}$$

where $T_r(t)$ is number or rate of completed/committed transports on route r.

4 Experiments

This section provides a reduced-order proof-of-concept evaluation of the proposed CPI routing mechanism. The goal is to verify whether the local routing rules generate stable route concentration, route coexistence, and route preference under different demand geometries, rather than to reproduce detailed industrial vehicle kinematics.

4.1 Simulation setup

To evaluate the proposed CPI-based self-organizing transportation model, we use a minimal-change lattice surrogate derived from the original ant-foraging simulation framework. In this surrogate, the nest is reinterpreted as a depot CPC, food sources are reinterpreted as task nodes, and the pheromone field is reinterpreted as a virtual route-desirability field. Vehicles are represented by two spatial states, namely exploratory vehicles and committed return-phase vehicles, and their motion is governed by local route reinforcement, task utility, and weak BLE-confidence weighting. The purpose of this spatial surrogate is not to reproduce detailed warehouse kinematics, but to visualize route emergence and task-node competition under fixed routing infrastructure.

Four scenarios are considered: a single-task case, a balanced two-task case, an asymmetric two-task case, and a blocked-routing case. For each scenario, 20 random seeds are evaluated. Performance is summarized using the trail-concentration index Θ, the asymmetry index A, and a coexistence measure that quantifies whether both branches remain active during the late phase of the simulation.

4.2 Evaluation results

Table I summarizes the multi-seed results for the four baseline scenarios. The single-task case provides the clearest baseline behavior. It yields the strongest route concentration, with a final Θ of 0.580 ± 0.007 and a peak Θ of 0.741 ± 0.022, indicating that the route-desirability field consistently condenses onto a single dominant corridor.

Table 1: Multi-seed summary of baseline CPI routing scenarios.

Scenario	Final Θ	Peak Θ	Formation time	Late $\lvert A\rvert$	Late coexistence
Single	0.580±0.007	0.741±0.022	0.285±0.048	--	--
Balanced	0.549±0.021	0.660±0.023	0.278±0.055	0.033±0.018	0.483±0.009
Asymmetric	0.421±0.031	0.597±0.037	0.196±0.039	0.309±0.033	0.345±0.016

Blocked	0.544±0.017	0.655±0.023	0.310±0.038	0.045±0.028	0.478±0.014

Under symmetric dual-task geometry, the proposed routing law maintains late-stage coexistence and near-zero asymmetry, indicating that the reinforcement mechanism does not collapse to a single dominant branch under balanced conditions. Its late signed asymmetry remains close to zero, while the late absolute asymmetry remains small (0.033 ± 0.018), indicating that neither branch systematically dominates across seeds. At the same time, the coexistence level remains high (0.483 ± 0.009), confirming that both branches stay active rather than collapsing to a single route. This result is important because it shows that the corrected CPI routing law can preserve balanced task-node competition under symmetric conditions.

The asymmetric case exhibits the expected preference behavior. Compared with the balanced case, its final and peak Θ values are lower, reflecting a less evenly distributed two-branch structure. More importantly, the late absolute asymmetry increases to 0.309 ± 0.033, while coexistence remains nonzero (0.345 ± 0.016). This indicates that one branch becomes consistently preferred, but the weaker branch is not fully suppressed. Hence, the asymmetric configuration produces persistent route preference rather than winner-take-all collapse.

The blocked case remains close to the balanced case in terms of concentration and coexistence, with a final Θ of 0.544 ± 0.017 and late coexistence of 0.478 ± 0.014. For representative spatial patterns, Figure 3 presents representative late-time route-desirability fields for the single, balanced, and asymmetric scenarios.

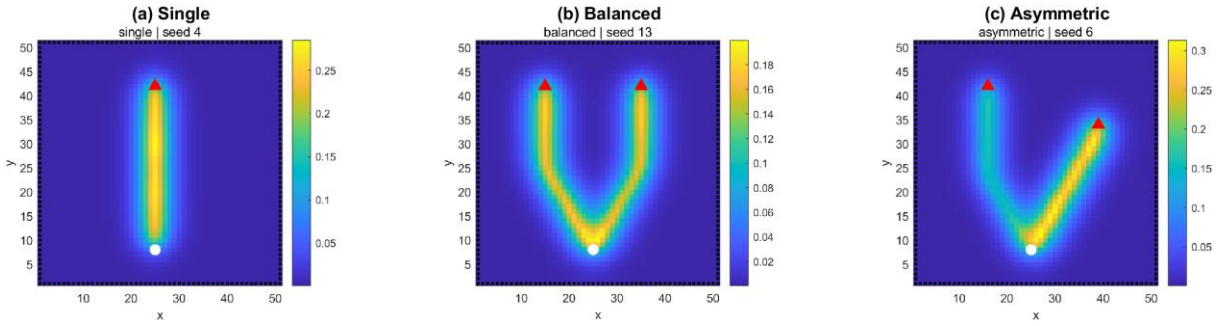


Figure 3: Representative late-time virtual route-desirability fields for the single, balanced, and asymmetric scenarios. Panel (a) shows the single-task case, in which the virtual route-desirability field condenses into one dominant corridor from the depot CPC to the task node. Panel (b) shows the balanced two-task case, where a stable Y-shaped bifurcating structure is formed and both branches remain active. Panel (c) shows the asymmetric two-task case, in which one branch becomes dominant while a weaker secondary branch remains active.

Conclusion

This paper proposes a two-level reduced-order self-organizing routing framework for internal logistics, inspired by ant transport and recruitment mechanisms. The main approach involves defining vehicles as idle, active, and support vehicle states, and combining these states with graph-based route desirability, BLE-based local congestion estimation, and an adjustment penalty term to obtain a self-organizing transportation allocation model. Example simulations demonstrate that the proposed CPI-based transportation allocation strategy retains the adaptive advantages of self-organization while adjusting path changes through the adjustment term and local CPC states. To further enhance applicability, we plan to extend the warehouse to more

complex layouts and consider dynamic obstacles in the future, comparing it with more robust dynamic routing and multi-robot baseline models.

References

- Fragapane G., de Koster R., Sgarbossa F., Strandhagen J.O. (2021): Planning and control of autonomous mobile robots for intralogistics: Literature review and research agenda. European Journal of Operational Research, v294, no2, 405-426.
- Jun, S., Choi, C. H., & Lee, S. (2022). Scheduling of autonomous mobile robots with conflict-free routes utilising contextual-bandit-based local search. *International Journal of Production Research*, *60*(13), 4090-4116.
- Lecheval V., Larson H., Burns D.D.R., Ellis S., Powell S., Donaldson-Matasci M.C., Robinson E.J.H. (2021): From foraging trails to transport networks: how the quality-distance trade-off shapes network structure. Proceedings of the Royal Society B, v288, 20210430.
- Li P., Wu W., Zhao Z., Huang G.Q. (2024): Indoor positioning systems in Industry 4.0 applications: Current status, opportunities, and future trends. Digital Engineering, v3, 100020.
- Montreuil B. (2011): Towards a Physical Internet: Meeting the global logistics sustainability grand challenge. Logistics Research, v3, no2-3, 71-87.
- Ouyang, Z., Yuan, Z., Chen, Q., Li, M., Zhao, Z., & Huang, G. Q. (2025). Shipment scheduling and routing protocols in cyber-physical internet for prefabricated construction modules logistics. IEEE Transactions on Intelligent Transportation Systems.
- Vaccari L., Coruzzolo A.M., Lolli F., Sellitto M.A. (2024): Indoor Positioning Systems in Logistics: A Review. Logistics, v8, no4, 126.
- Zhu, W., Oğuz, S., Heinrich, M. K., Allwright, M., Wahby, M., Christensen, A. L., ... & Dorigo, M. (2024). Self-organizing nervous systems for robot swarms. Science Robotics, 9(96), eadl5161.

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE
CIÊNCIAS EMPRESARIAIS
POLITÉCNICO SETÚBAL
SCHOOL · POLYTECHNIC UNIVERSITY

Digital Twins: Accelerating the Transition of Port Community Systems toward the Physical Internet – The Atlantic Corridor Case Study

João Lemos Nabais[1], Raquel Pereira[2] and Carlos Batista[2]
[1] IDMEC, Instituto Politécnico de Setúbal, School of Business Administration, Setúbal, Portugal
[2] Instituto Politécnico de Setúbal, School of Business Administration, Setúbal, Portugal
Corresponding author: joao.nabais@esce.ips.pt

Abstract: *The rapid evolution of global logistics and synchromodal transport has made the adoption of advanced digital technologies necessary in order to manage complexity and improve operational efficiency. The implementation of Port Community Systems (PCS) has boosted the efficiency of seaport communities. The Physical Internet (PI) is an innovative concept providing new logistics solutions that require integration and interoperability between transport network partners, including those involved in maritime and land transport. PI is based on highly interconnected systems that enable the seamless sharing of information and assets. This leads to the efficient use of available resources through the consolidation of flows within the system. Digital twins offer a highly accurate virtual representation of physical assets, enabling real-time monitoring, predictive analytics, and closed-loop optimisation. In the context of container transport networks, Digital Twin (DT) frameworks are designed to enhance visibility and enable proactive decision-making. This paper presents some of the benefits of using a digital twin in a transport network to monitor activities and make autonomous decisions. The Atlantic Corridor case study is used for illustrative purposes.*

Keywords: Digital Twin, Port Community Systems (PCS), Synchromodal Transport, TEN-T, Physical Internet.

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☒ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☐ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Mode-s*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☒ *In-Person presentation*

1 Introduction

Freight transportation plays a fundamental role in the global economy, enabling the efficient movement of goods and supporting international trade and economic development. However, transport networks are facing an increase in the complexity of physical and information flows. Goods traceability and transparency are now standard service quality requirements worldwide. Seaports are unique locations within transport networks where various stakeholders with different perspectives and conflicting objectives interact. The development of container transport networks based on information technologies has promoted the integration of logistics operations. Investment in information systems is commonplace in seaports, facilitating paperless interactions between port community members. This is a necessary step towards making seaports more competitive. The Port Community System (PCS) supports the

development of a network connecting seaport stakeholders. As an electronic platform, the PCS links multiple systems of distinct companies and organisations within the port community, guiding the seaport towards business digitisation (Rodon & Ramis-Pujol, 2008; Carlan et al., 2016). In major European and global seaports, the PCS is becoming the primary technological platform connecting shipping agents, shippers, freight forwarders, transport providers, terminals, logistics platforms and public entities involved in cargo services crossing the seaport (Caldeirinha et al., 2022).

The ongoing globalisation of transport networks further highlights the need for integrated, efficient freight transport systems that can handle increasing volumes of goods. This expansion puts considerable pressure on existing infrastructure, emphasising the importance of coordinated, multimodal logistics solutions. In this context, using multiple transport modes is a promising way to overcome logistical constraints and promote more sustainable freight flows. A key consideration in logistics decision-making is shifting freight flows from faster, more polluting modes such as road transport to slower, less carbon-intensive alternatives such as rail transport (Flodén, 2007). This modal shift is widely recognised as a vital part of the decarbonisation strategy promoted by the European Green Deal, which aims to reduce transport-related emissions by changing the way transport networks work (Bogojević and Drenovak-Ivanović, 2020). Scholars increasingly emphasise that such transitions have the potential to reconcile economic efficiency with environmental sustainability in freight logistics systems. However, achieving this transformation requires improved coordination among stakeholders and more effective data-sharing mechanisms across transport networks. When available, logistics data is often fragmented across organisations and stored in heterogeneous formats, which complicates integration and coordination efforts. Similar challenges related to data availability, quality, reliability, temporal constraints and cost limitations are frequently encountered in logistics data collection across different sectors (Brochado et al., 2024).

The concept of the Physical Internet (PI) was first introduced by Montreuil (2011), who defined it as an open logistics system designed for the sustainable and efficient transportation of physical objects around the world. It is a metaphor for the digital internet applied to logistics service networks. The PI is based on highly interconnected systems that enable the seamless sharing of information and assets, leading to the efficient use of available resources through the consolidation of flows within the system. In transport networks, the standardised container resembles a PI-container, and the available transport can be seen as PI-movers. The Physical Internet involves integration and interoperability between supply chain actors. Seaports and terminals are part of global supply chains. The Physical Internet is not just a reality; it also contains a vision for the near future and guidelines for behavioural change. The European Technology Platform ALICE (Alliance for Logistics Innovation through Collaboration in Europe) focuses on defining a research and innovation strategy for supply chain management. Research and statements from this platform could lead to the adoption of the Physical Internet. The goal is to gradually transition from a gateway perspective to Physical Internet implementation. This implementation uses the one-stop-shop concept to: take transport orders and guarantee deliveries using available network resources; and be able to autonomously react to local unforeseen events such as congestion, among others (ALICE-ETP, 2020). In order to realise the vision of a PI and build a digital logistics system, it is fundamental to be able to monitor the PI in real time (through IoT advances, this is now a reality) and improve the visualisation and traceability functions. In a cyber-physical internet (CPI), information flows in the cyber layer and items are transported in the physical layer (Wu et al., 2024; You et al., 2025).

Digital Twins (DTs) aim to integrate technologies, processes and human expertise to create highly realistic digital systems that support the management of transport networks, improving cost-efficiency, quality, reliability and responsiveness. The DT concept has recently gained considerable attention, as organisations increasingly recognise its value in supply chain management (Le & Fan, 2024). Similarly to the CPI, the DT is also based on two layers: a virtual twin (the digital component) and a physical twin (the real world), which are connected through the IoT. DTs provide deep operational insights, incorporate multiple system dynamics, and offer predictive capabilities:

Real-Time Monitoring and Prediction: DTs provide high-fidelity visualisations of physical assets using numerous IoT sensors, with data analysed through machine learning algorithms to support strategic decision-making. This approach enables stakeholders to transition from merely monitoring deviations to proactively predicting and managing long-term logistics operations. For instance, in port environments, DT can generate performance forecasts based on the current system state to assist dispatch operators with planning and resource allocation.

Simulation and Optimization: DT enables scenario simulation, allowing stakeholders to evaluate operational outcomes and identify optimal strategies without exposing the real system to risk. This capability is critical for logistics decision support and relies on real-time access to current system data.

Data Integration and Management: DT facilitates the seamless integration of heterogeneous data from multiple sources and operational dimensions. They address data fragmentation, providing comprehensive and accurate analysis support. This promotes robust communication, collaboration and planning among stakeholders, particularly when managing complex datasets from diverse origins.

This paper focuses particularly on the current development of digital twins that can operate alongside PCS, providing access to data monitoring and offering valuable insights to improve the performance of transport networks in the context of the future implementation of a cyber-physical internet.

2 Methodology

This work takes the perspective of a global operator that can be seen as a regulator of a container transport network. The focus is on solving the problem of routing containers to ensure they are all delivered, rather than managing the problem of booking containers. In terms of transport demand, all containers delivered to the system are accepted for transport, thus fulfilling the market's transport needs. Transport demand comes from the local market, other transport networks, or international maritime trade. In contrast to a corridor, the network level addresses the possibility of accomplishing transshipment operations. It is also important to note the capability to account for operations at an international level, as well as for maritime operations, as opposed to work that only accounts inland networks (Zhang et al., 2025). In terms of the simultaneous consideration of maritime routes and inland networks, this work is similar to that discussed by Guo et al. (2025).

A container transport network can be represented as a network of nodes (such as seaports, ports, terminals and warehouses) and connections (such as road, rail and waterway transport connections). Nodes can be related to potential, as goods are stored there, and transport can be related to flows, as movement between locations is accomplished. From a cargo flow perspective, the various economic parties involved in container transport networks (Rodrigue et al., 2009) can be categorised into two main groups:

Intermodal hub managers: Manage the storage of cargo and facilitate the transfer of transport modes for cargo en route to its final destination;

Transport operators: offer transport capacity via various transport modes between existing intermodal hubs.

Although all parties contribute to the main objective of the container transport network, which is to deliver commodities to the agreed location at the agreed time in the agreed quantity, each stakeholder has its own objectives, which can lead to conflicts (Ishfaq and Sox, 2010). Due to these conflicting objectives and the resulting competition, a lack of confidence exists among the different parties. This has a significant impact on the type of information available to each stakeholder regarding the cargo to be transported, resulting in different transport paradigms: i) Merchant haulage, in which the shipper or forwarder is responsible; ii) Carrier haulage, in which the transport provider organises the transport; and iii) Terminal haulage, in which the terminal co-determines the transport.

This work assumes that the transport organization follows the terminal haulage approach. Terminals are special locations on the transport network where multiple stakeholders interact. Terminal operations are related to cargo handling, specifically the loading and unloading of cargo from and to transport connections, as well as the reshuffling of cargo within the terminal yard to prepare future operations. A terminal model should capture the necessary information to support analysis of the transport network, with the aim of achieving a more collaborative and sustainable transport organization (Figure 1). It is critical to distinguish the containers taking into account the destination. The container can be an empty container with no destination or a loaded container with cargo with a known destination. This information is relevant to proceed with the transport planning such that the container is moved towards its destination. Reducing transport costs depends on choosing the most appropriate route for each container within the network, depending on the context. A terminal can be described in terms of the available parking area (which can be divided into smaller, dedicated areas), the number of transport modalities available, the operating hours of the modalities, and the number of gates for road and rail modalities, and quay length for waterway modalities.

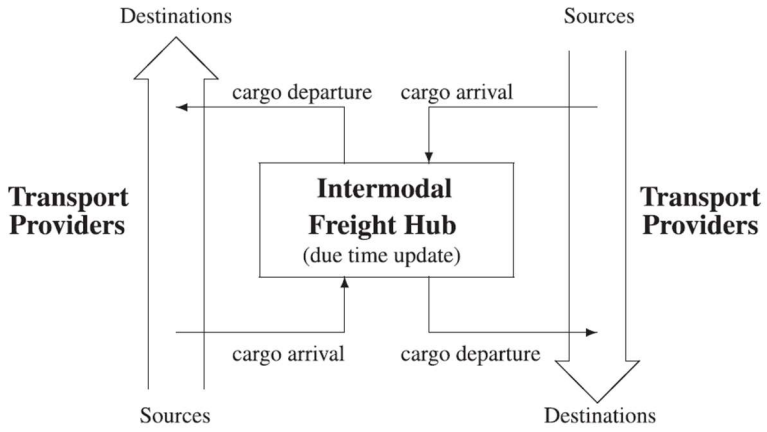


Figure 1: Node and connection interaction in a container transport network (Nabais et al., 2015).

There are two types of operations at a terminal: i) Load cargo operations: decisions to move a container into a means of transport are made by the terminal in close collaboration with transport operators. ii) Unload cargo operations: these operations are the result of a previous loading operation at the origin of the connection. Assuming that resources are available (handling equipment or storage area), the cargo will be unloaded as expected. If the terminal is congested, the cargo can be placed in a waiting queue until resources become available.

Connections allow us to overcome the distance between locations. The main parameter that describes a connection is therefore the average estimated time it takes to travel between two destinations. Other relevant information includes transport capacity, available schedules, and the transport route, which helps to identify reachable destinations. The transport capacity of the vehicles allocated to the connection, together with the transport frequency, determines the constraints regarding container transport capacity over a specific time interval. The fundamental components of transport networks — nodes and connections — are implemented in Matlab and form the basis of the Transport Network Library (TNL) (Figure 2). The TNL is composed of the following sections: i) Node, ii) Connection, iii) Market and iv) Planning and Control.

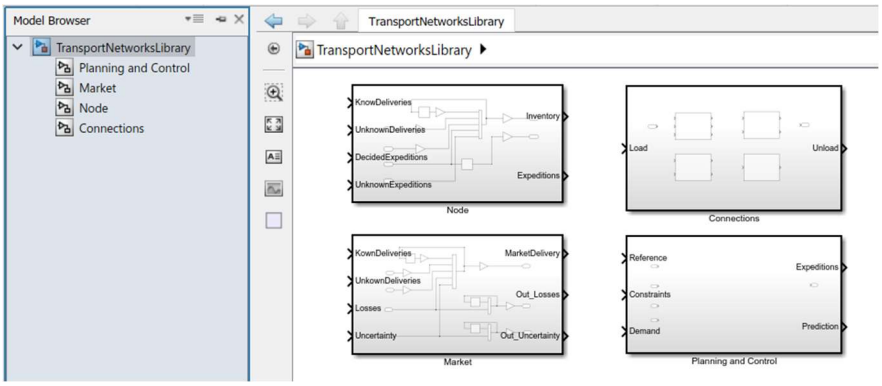


Figure 2: Transport Network Library (TNL) implemented in Matlab.

The Container Allocation Problem (CAP) is addressed at each terminal by the corresponding Control Tower, which follows the terminal haulage approach. Each control tower implements Sustainable Awareness Cargo Allocation (SACA) in terminal operations, which follows Transport Modal Split Heuristics (see Figure 3, Nabais et al., 2025). This heuristic is applied locally at each control tower, which has access only to the existing cargo at the terminal and the visiting connections (and, respectively, the destinations reachable by each connection). Once an unloading operation starts at the terminal, the containers are routed to a reserved area in the central yard corresponding to the transport modality to be used for the next transport step. Using this approach, containers arriving at the terminal instantly receive the transport modality indication and are moved to join the queue of containers for that modality.

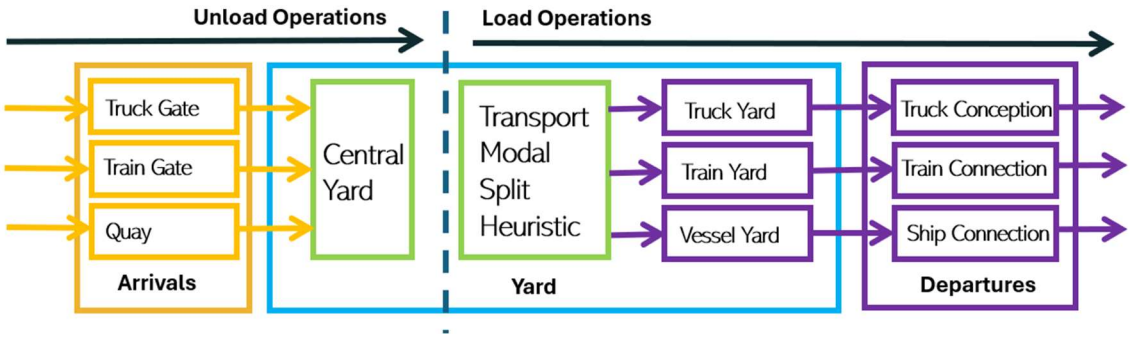


Figure 3: General representation of the Sustainable Awareness Cargo Allocation (SACA) at the terminal following a push approach.

The procedure for assigning a transport modality to arriving containers is based on the Transport Modal Split (TMS). Taking a transport modal split as a reference, the containers are divided between the transport modalities according to the defined shares. This push-based approach involves loading containers onto connections towards the destination as soon as a

connection becomes available, thereby bringing the cargo closer to its destination. A transport modal split guideline is included for each destination at each terminal to promote sustainability awareness. Access to the desired destination at each terminal is critical for the efficient allocation of cargo, enabling containers to be moved closer to their final destination (provided there are no network congestion or bottleneck issues at that time).

3 Case Study: Atlantic Corridor

The Atlantic Corridor is a significant railway project running along the European Atlantic coast. It is part of the Trans-European Transport Networks (TEN-T). Its primary objective is to facilitate the transportation of goods and promote logistical efficiency and competitiveness. Consisting of sections of existing and planned railway infrastructure, the Atlantic Corridor connects major cities and ports, including Sines, Lisbon, Madrid, Bilbao, Nantes, Paris, Strasbourg and Mannheim (see Figure 4).

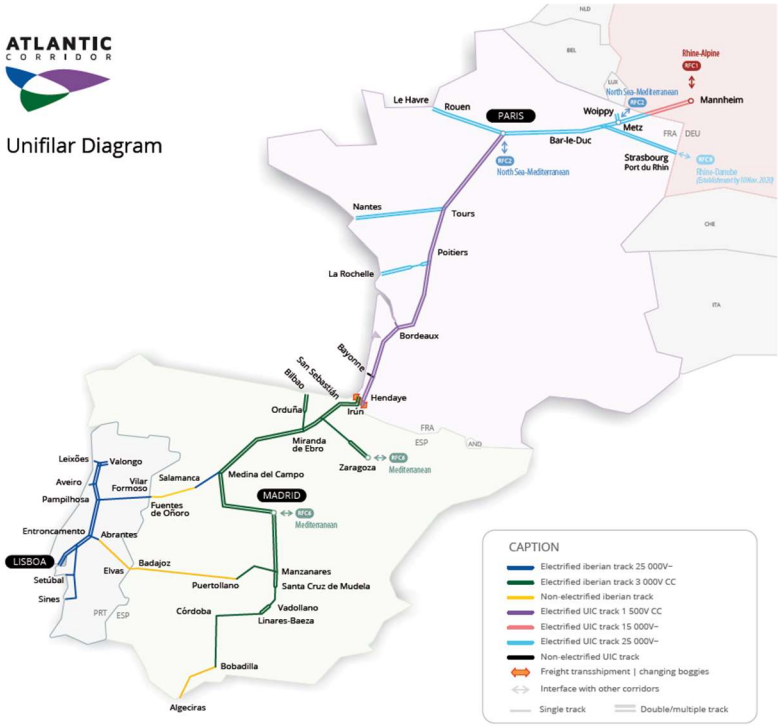


Figure 4: Railway network of the Atlantic Corridor (Atlantic Corridor, 2026).

The aim is to overcome technical and operational barriers and promote interoperability among the railway systems of the participating countries. The objective is to apply the "One Stop Shop" concept for cargo owners. From the cargo owner's perspective, the system will handle the delivery with no need for the owner to provide any further technical details, as is the case with the internet. In Europe, the Atlantic Corridor plays a crucial role in the trans-European transport network by aiming to enhance connectivity and transport capacity between coastal regions. For illustrative purposes, the Atlantic Corridor has been simplified to include seven nodes (seaports and terminals) in three countries: Portugal (Sines and Entroncamento), Spain (Madrid, Bilbao and Irún) and France (Le Havre and Paris). This simplified configuration captures the Atlantic Corridor's key characteristics: i) international maritime transport is assured via the seaports of Sines, Bilbao and Le Havre; ii) connection to other trans-European transport networks is assured via Madrid and Paris (to the Mediterranean and the North Sea–Rhine respectively) and iii) Irún is a critical location for railway infrastructure, ensuring a smooth transition between the Iberian and European gauges (see Figure 5).

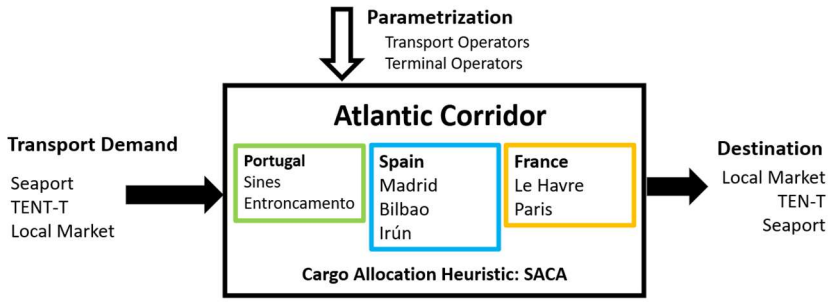


Figure 5: Transport within the Atlantic Corridor in what resembles a Physical Internet.

3.1 Digital Twin Setup

Transport demand is posted to the network from the local market, other TEN-T networks, and international seaports. The system allows all origin-destination pairs and therefore makes no distinction between import and export flows. The Digital Twin was constructed using containers categorised into eight categories, corresponding to the seven nodes plus empty containers. According to the dynamics of the Atlantic Corridor, the simulation was based on a sample time of one hour. Connection times between locations are given as follows:

Road: Truck trips of more than 16 hours are considered to require two drivers, allowing for 16 hours of travel and eight hours of rest per day. It is assumed that trucks are always available in the vicinity of the seaport, but loading operations are restricted to six days per week, from 08:00 to 20:00. Transport time was determined based on connection distance and an average velocity of 70 km/h (see Table 1).

Railways: Train connections are restricted to pre-defined schedules, using the available infrastructure. Loading operations are available six days per week. The connection between Iberia and France is made at Irún, which is connected to five terminals (see Table 2).

Waterways: There are maritime connections between the seaports of Sines, Bilbao and Le Havre. Ship connections are less frequent and follow a pattern similar to a weekly maritime service, with loading operations seven days a week.

Table 1: Time table for road connections (time expressed in hours).

	Portugal		Spain			France	
	Sines	Entron.	Madrid	Bilbao	Irún	Le-Havre	Paris
Sines	-	3	8	14	17	28	26
Entroncamento	3	-	7	11	12	24	23
Madrid	8	7	-	5	6	18	17
Bilbao	14	11	5	-	2	14	13
Irún	17	12	6	2	-	13	12
Le-Havre	28	24	18	14	13	-	3
Paris	26	23	17	13	12	3	-

Table 2: Time table for railway connections (time expressed in hours).

	Portugal		Spain			France	
	Sines	Entron.	Madrid	Bilbao	Irún	Le-Havre	Paris
Sines	-	5					
Entroncamento	5	-	6	15	19		
Madrid		6	-	6	7		
Bilbao		15	6	-	2		
Irún		19	7	2	-	19	17
Le-Havre					19	-	4
Paris					17	4	-

To illustrate this, different scenarios were created using synthetic data inspired by the available information on the Atlantic Corridor (Atlantic Corridor, 2026). Two dimensions were considered for testing:

Transport Demand: Transport demand can be interpreted as transport requests from agents outside the Atlantic Corridor. Depending on whether it arrives at seaports (Sines – PT, Bilbao – SP and Le Havre – FR) or intermodal locations (Madrid and Paris), different demand profiles were built. Transport demand arrives at seaports during the first week of the simulation. For hinterland gateways Madrid and Paris, transport demand occurs during the first 30-day period, from Monday to Saturday. After the first 30 days (720 hours), no more transport demand is added to the transport network to allow time for all containers to be delivered to their destinations, enabling fair comparisons to be made between the tests. Scenarios 2 and 3 are obtained by applying factors of 1.25 and 1.5 to the transport demand arriving at Sines and Le Havre, respectively (Table 3). To test the impact of demand stochasticity variability on system performance, 70 demand profiles were considered for Madrid, Bilbao and Paris.

Cargo Allocation (SACA): Using the TMS indicated in Table 4 as a reference, both the Sines and Le Havre terminals apply different transport modal splits for destinations such as Madrid, Bilbao and Paris, creating new sets named TMS 2 and TMS 3 (see Table 5). The evolution from TMS 1 to TMS 3 ensures that the most sustainable transport modalities are favoured. For TMS 3, cargo destined for Madrid or Bilbao can depart from either Sines or Le Havre using only one transport modality. For TMS 3 at Sines, all cargo destined for Paris is loaded onto a vessel that will either go to Bilbao or Le Havre, and from there the final destination will be reached.

Table 3: Transport demand for Scenario 1 (In brackets, the variability interval).

	Source	Destination						
		Portugal		Spain			France	
		Sines	Entron.	Madrid	Bilbao	Irún	Le-Havre	Paris
Foreland	Sines	[40 ; 120]	[20 ; 60]	[180 ; 350]	[80 ; 240]	[20 ; 60]	[900 ; 1500]	[300 ; 700]
	Bilbao	[101 ; 139]	[31 ; 88]	[31 ; 89]	[302 ; 497]	[61 ; 100]	[102 ; 300]	[52 ; 150]
	Le-Havre	[600 ; 1000]	0	[375 ; 500]	[150 ; 300]	0	[300 ; 500]	[750 ; 1000]
Hinterland	Entroncamentc	0	0	0	0	0	0	0
	Madrid	[531 ; 635]	0	0	[535 ; 660]	0	0	[650 ; 816]
	Irún	0	0	0	0	0	0	0
	Paris	[369 ; 414]	0	[517 ; 660]	[422 ; 496]	0	0	0

Table 4: TMS for configuration of the SACA approach, with impact in cargo flows.

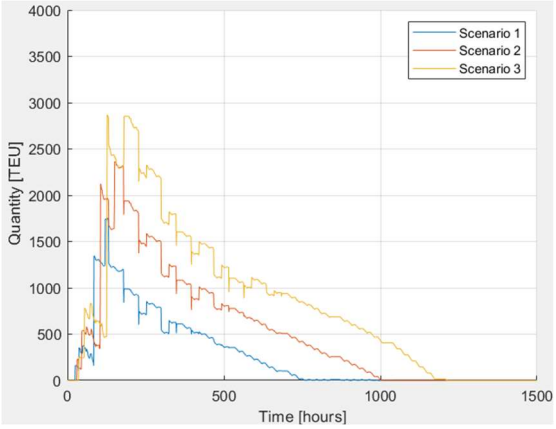
	Portugal		Spain			France	
	Sines	Entron.	Madrid	Bilbao	Irún	Le-Havre	Paris
Sines	(1.0 ; 0.0 ; 0.0)	(0.0 ; 1.0 ; 0.0)	(0.5 ; 0.5 ; 0.0)	(0.0 ; 0.3 ; 0.7)	(0.0 ; 1.0 ; 0.0)	(0.0 ; 0.0 ; 1.0)	(0.0 ; 0.2 ; 0.8)
Entron.	(0.2 ; 0.8 ; 0.0)	(1.0 ; 0.0 ; 0.0)	(0.2 ; 0.8 ; 0.0)	(0.2 ; 0.8 ; 0.0)	(0.2 ; 0.8 ; 0.0)	(0.0 ; 1.0 ; 0.0)	(0.5; 0.5 ; 0.0)
Madrid	(0.2 ; 0.8 ; 0.0)	(0.1 ; 0.9 ; 0.0)	(1.0 ; 0.0 ; 0.0)	(0.2 ; 0.8 ; 0.0)	(0.2 ; 0.8 ; 0.0)	(0.2 ; 0.8 ; 0.0)	(0.2 ; 0.8 ; 0.0)
Bilbao	(0.2 ; 0.3 ; 0.5)	(0.2 ; 0.8 ; 0.0)	(0.2 ; 0.8 ; 0.0)	(1.0 ; 0.0 ; 0.0)	(0.2 ; 0.8 ; 0.0)	(0.2 ; 0.3 ; 0.5)	(0.2; 0.6 ; 0.2)
Irún	(0.2 ; 0.8 ; 0.0)	(0.2 ; 0.8 ; 0.0)	(0.2 ; 0.8 ; 0.0)	(0.2 ; 0.8 ; 0.0)	(1.0 ; 0.0 ; 0.0)	(0.2 ; 0.8 ; 0.0)	(0.0 ; 1.0; 0.0)
La-Havre	(0.0 ; 0.0 ; 1.0)	(0.3 ; 0.7 ; 0.0)	(0.3 ; 0.7 ; 0.0)	(0.1 ; 0.2 ; 0.7)	(0.2 ; 0.8 ; 0.0)	(1.0 ; 0.0 ; 0.0)	(0.2 ; 0.8 ; 0.0)
Paris	(0.2 ; 0.8 ; 0.0)	(0.2 ; 0.8 ; 0.0)	(0.2 ; 0.8 ; 0.0)	(0.2 ; 0.8 ; 0.0)	(0.2 ; 0.8 ; 0.0)	(0.2 ; 0.8 ; 0.0)	(1.0 ; 0.0 ; 0.0)

Table 5: Variable Transport Modal Splits at Sines and Le-Havre, with impact in cargo flows.

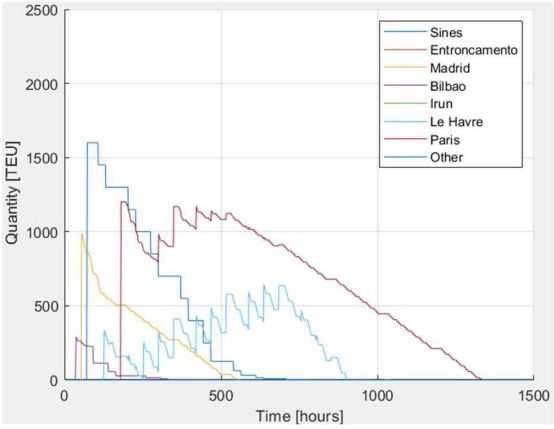
	Sines			Le-Havre		
	Madrid	Bilbao	Paris	Madrid	Bilbao	Paris
TMS 1	(0.5 ; 0.5 ; 0.0)	(0.0 ; 0.3 ; 0.7)	(0.0 ; 0.2 ; 0.8)	(0.3 ; 0.7 ; 0.0)	(0.1 ; 0.2 ; 0.7)	(0.2 ; 0.8 ; 0.0)
TMS 2	(0.2 ; 0.8 ; 0.0)	(0.0 ; 0.2 ; 0.8)	(0.0 ; 0.1 ; 0.9)	(0.1 ; 0.9 ; 0.0)	(0.0 ; 0.1 ; 0.9)	(0.1 ; 0.9 ; 0.0)
TMS 3	(0.0 ; 1.0 ; 0.0)	(0.0 ; 0.0 ; 1.0)	(0.0 ; 0.0 ; 1.0)	(0.0 ; 1.0 ; 0.0)	(0.0 ; 0.0 ; 1.0)	(0.0 ; 1.0 ; 0.0)

3.2 Numerical Results

The experiments were conducted using Matlab 2021a and Simulink on an 11th generation Intel Core i7-1165G7 processor with 16 GB of RAM. Once the scenario demand and TMS have been set, the cargo allocation runs automatically according to the SACA approach. Transport demand delivered at a terminal is automatically handled towards the destination, with all cargo allocation decisions being taken automatically. Information at each terminal and transport connection is accessible and visible to all agents within the Atlantic Corridor Digital Twin.



a) Inventory level at Sines for TMS 1.



b) Inventory level (per type of container) at Le-Havre for TMS 1 with Scenario 3.

Figure 6: Monitoring inventory levels.

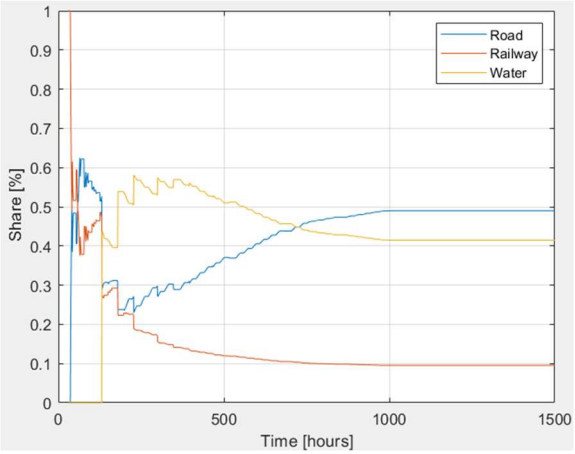
It is possible to monitor inventory levels at terminals, which have a clear impact on loading operations, particularly for vessels (Figure 6a). Inventory levels at each location are accessible by destination, as this information is critical for transporting cargo to its destination (Figure 6b), and for planning routes and schedules. Yard occupancy at terminals can be inspected using the daily sum of inventory levels (Table 6). Transport operations can also be monitored using the Digital Twin, for example to monitor the cumulative transport modal split for loaded cargo (Figure 7a) and to monitor cargo loading operations (Figure 7b). Table 7 shows the time taken to complete cargo allocation. It is clear that operations take more time to complete at seaports, followed by TEN-T gateways, with passing-by locations being the least sensitive.

Table 6: Sum of daily inventory at Irún, in brackets the interval values for the 70 simulations.

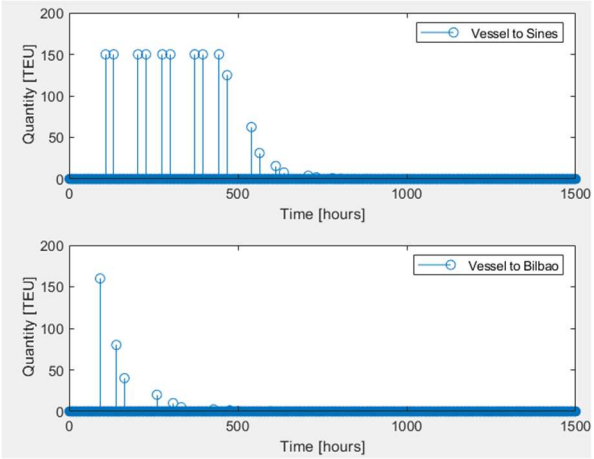
		Portugal		Spain			France	
		Sines	Entron.	Madrid	Bilbao	Irún	Le-Havre	Paris
TMS 1	Scenario 1	[2655 ; 3034]	0	[9835 ; 14332]	[3646 ; 4316]	[258 ; 392]	[365 ; 1102]	[23957 ; 79885]
	Scenario 2	[2655 ; 3034]	0	[16743 ; 25228]	[4038 ; 4974]	[484 ; 604]	[365 ; 1102]	[37347 ; 100780]
	Scenario 3	[2655 ; 3034]	0	[25678 ; 39405]	[4566 ; 5827]	[926 ; 1042]	[365 ; 1102]	[50642 ; 120610]
TMS 2	Scenario 1	[2655 ; 3034]	0	[13727 ; 20168]	[3281 ; 3824]	[258 ; 392]	[365 ; 1102]	[14336 ; 54895]
	Scenario 2	[2655 ; 3034]	0	[24232 ; 37055]	[3406 ; 3985]	[484 ; 604]	[365 ; 1102]	[16321 ; 64834]
	Scenario 3	[2655 ; 3034]	0	[38856 ; 61022]	[3555 ; 4183]	[926 ; 1042]	[365 ; 1102]	[21618 ; 74166]
TMS 3	Scenario 1	[2655 ; 3034]	0	[15862 ; 23438]	[3045 ; 3575]	[258 ; 392]	[365 ; 1102]	[9522 ; 31331]
	Scenario 2	[2655 ; 3034]	0	[28753 ; 44235]	[3045 ; 3575]	[484 ; 604]	[365 ; 1102]	[9522 ; 31331]
	Scenario 3	[2655 ; 3034]	0	[46842 ; 74237]	[3045 ; 3575]	[926 ; 1042]	[365 ; 1102]	[9522 ; 31331]

Different transport demands put pressure on transport network performance, but this is not equally apparent in all locations, as cargo flows also depend on the transport modalities chosen by the SACA approach. The seaports of Sines and Le Havre are more sensitive to demand variations as they receive close to 80% of the foreland transport demand. This can be seen in the daily sum of inventory levels (Table 6) and the number of trains loaded, as well as the time

required to ship all cargo at those terminals to their destinations, i.e. to set inventory levels to zero (Table 7 and 8). Using Scenario 3 as a reference point, changing from TMS 1 to TMS 3 improves sustainable awareness when choosing a transport modality. This reduces truck kilometres from 2,137,910 km (Table 9) to 1,482,247 km (Table 10), which is a 30% reduction in kilometres. This has an impact on road congestion and CO_2 emissions.



a) Transport Modal Split at Sines for TMS 1 with Scenario 1.



b) Loaded cargo into vessels at Le-Havre for TMS 3 with Scenario 3.

Figure 7: Transport Modal Split and loaded trains monitoring.

Table 7: Time required to deliver all containers to final destination, measured as the time to inventory reach zero.

		Portugal		Spain			France	
		Sines	Entron.	Madrid	Bilbao	Irún	Le-Havre	Paris
TMS 1	Scenario 1	993	995	917	777	1085	757	1112
	Scenario 2	1002	995	994	777	1121	1061	1138
	Scenario 3	1212	995	1137	801	1163	1379	1377
TMS 2	Scenario 1	993	995	945	777	1043	827	1065
	Scenario 2	1002	995	1114	777	1115	1217	1211
	Scenario 3	1212	995	1282	810	1283	1577	1572
TMS 3	Scenario 1	993	995	974	777	989	947	994
	Scenario 2	1002	995	1162	777	1163	1385	1379
	Scenario 3	1212	995	1377	825	1373	1793	1788

Table 8: Number of train operations considered, measured as the number of trains that were loaded with at least one container.

		Portugal		Spain			France	
		Sines	Entron.	Madrid	Bilbao	Irún	Le-Havre	Paris
TMS 1	Scenario 1	228	990	900	258	1074	732	756
	Scenario 2	282	990	900	258	1116	1050	756
	Scenario 3	366	990	900	258	1116	1050	756
TMS 2	Scenario 1	228	990	900	258	1038	822	756
	Scenario 2	306	990	900	258	1110	1212	756
	Scenario 3	414	990	900	258	1278	1572	756
TMS 3	Scenario 1	258	990	900	258	978	942	756
	Scenario 2	390	990	900	258	1158	1380	756
	Scenario 3	486	990	900	258	1338	1788	756

Table 9: Configuration for TMS 1 with Scenario 3 (2.137.910 km by truck).

	Portugal		Spain			France	
	Sines	Entron.	Madrid	Bilbao	Irún	Le-Havre	Paris
Sines	0	0 0	190400	0	0	0	0
Entroncamento	48720	0 0	29232	14573	13216	0	0
Madrid	123760	0 0	0	72436	0	0	297180
Bilbao	24000	7670 0	31840	0	2096	40000	18880
Irún	63550	0 0	130702	15327	0	10440	0
Le-Havre	0	0 0	397500	32000	0	0	77164
Paris	143364	0 0	231140	122720	0	0	0

Table 10: Configuration for TMS 3 with Scenario 3 (1.482.247km by truck).

	Portugal		Spain			France	
	Sines	Entron.	Madrid	Bilbao	Irún	Le-Havre	Paris
Sines	0	0	0	0	0	0	0
Entroncamento	48720	0	58464	0	13216	0	0
Madrid	123760	0	0	72436	0	0	297180
Bilbao	24000	7670	31840	0	2096	40000	18880
Irún	64575	0	158122	13624	0	10440	0
Le-Havre	0	0	0	0	0	0	0
Paris	143364	0	231140	122720	0	0	0

4 Conclusions

The Port Community Systems have promoted digitalisation at seaports and terminals, leading to seamless data exchange and accessibility and the creation of digital logistics systems. Digital Twins capture the main dynamics at locations and have a structure similar to the Cyber-Physical Internet with a cyber and a physical layer. This paper demonstrates that access to data in the Digital Twin ensures the ability to monitor transport operations and make local decisions wisely. The transport network structure addressed is set at a network level, allowing transshipment operations without restrictions on import or export flows. The perspective taken is from a global operator focused on assuring the delivery of all cargo to the destination using a Terminal Haulage approach. Local decisions at each control tower are made using a Sustainable Awareness Cargo Allocation (SACA), where the configuration set for transport modal splits at each location can have a significant impact on how cargo flows evolve over the transport network, either leading to or avoiding congestion and promoting sustainable transportation. Both import and export operations are addressed in the proposed framework, resembling Physical Internet behavior with no limitations regarding origin and destination pairs. Future work will focus on: i) developing additional components for information exchange between stakeholders in the transport network; ii) including predictive analysis of the impact of decisions on network performance; and iii) including containers due time.

5 Acknowledgements

Content produced within the scope of the Agenda "NEXUS - Pacto de Inovação – Transição Verde e Digital para Transportes, Logística e Mobilidade", financed by the Portuguese Recovery and Resilience Plan (PRR), with no. C645112083-00000059 (investment project no.º 53). The authors acknowledge Fundação para a Ciência e a Tecnologia (FCT) for its financial support via LAETA (project https://doi.org/10.54499/UID/50022/2025).

References

ALICE-ETP (2020). Roadmap to The Physical Internet, Executive Version.

Atlantic Corridor (2026). https://www.atlantic-corridor.eu/, last accessed 2026/02/12.

Bogojević, S., & Drenovak-Ivanović, M. (2020). EU climate change litigation and the European Green Deal: The role of law in climate governance. Review of European, Comparative & International Environmental Law, 29(3), 339–347.

Brochado, Â. F., Rocha, E. M. & Costa, D. (2024). A Modular IoT-Based Architecture for Logistics Service Performance Assessment and Real-Time Scheduling towards a Synchromodal Transport System, Sustainability (Switzerland), vol. 16(2).

Caldeirinha, V., Nabais, J.L., Pinto, C. (2022). Port Community Systems: Accelerating the Transition of Seaports toward the Physical Internet—The Portuguese Case. Journal of Marine Science Engineering, 10, 152. https://doi.org/10.3390/jmse10020152

Carlan, V., Christa, S. & Vanelslander, T. (2016). Port Community Systems Costs and Benefits: From Competition to Collaboration within the Supply Chain; Helenic Institute of Transport: Thessaloniki, Greece, 2016.

Flodén, J. (2007). Environmental assessment of transport modes in logistics systems. International Journal of Logistics Research and Applications, 10(2), 123–138.

Guo, W., Atasoy, B., van Blokland, W.B. & Negenborn, R.R. (2021). Global sybchromodal transport with dynamic and stochastic shipment matching. Transportation Research Part E: Logistics and Transportation Review, 152, 102404.

Ishfaq, R. & Sox, C.R. (2010). Intermodal logistics: The interplay of financial, operational and service issues. Transportation Research Part E: Logistics and Transportation Review, Elsevier, vol. 46(6), pages 926-949.

Le, T. V., & Fan, R. (2024). Digital twins for logistics and supply chain systems: Literature review, conceptual framework, research potential, and practical challenges. Computers & Industrial Engineering, 187, 109768.

Montreuil B. (2011): Towards a Physical Internet: Meeting the Global Logistics Sustainability Grand Challenge, Logistics Research, v3, no2-3, 71-87, https://doi.org/10.1007/s12159-011-0045-x

Nabais, J. L., Batista, C. M. O., Pereira, R. F. & Gonçalves, M. (2025). Container intermodal transport along the Atlantic corridor with sustainable awareness using a digital twin. In Proceedings of the 25th International Scientific Conference: Business Logistics in Modern Management. Osijek: Business Logistics in Modern Management Conference. ISSN 1849-5931; ISSN 1849-6148.

Nabais, J.L., Negenborn, R.R., Carmona Benítez, R.B. & Botto, M.A. (2015). Achieving transport modal split targets at intermodal freight hubs using a model predictive approach. Transportation Research Part C, 60, pp. 278-297.

Rodon, J. & Ramis-Pujol, J. (2008). Exploring the intricacies of integrating with a port community system. In Proceedings of the 19th Bled eConference eValues, Bled, Slovenia, 5–7 June 2008.

Rodrigue, J.P., Comtois, C., Slack, B., 2009. The Geography of Transport Systems, second ed. Routledge, London.

You, Y., Yu, C., Ouyang, Z. & Li, M. (2025). Four-dimensional spatial-temporal packet encapsulation in the Cyber-Physical Internet: A reinforcement learning Approach. Transportation Research Part E: Logistics and Transportation Review, 211, 104843.

Wu, H., Huang, L., Li, M. & Huang, G.Q. (2024). Cyber-Physical internet (CPI)-enable logistics infrastructure integration framework in the greater bay area. Advanced Engineering Informatics, 60, 102551.

Zhang, Y., Tan, X., Gan, M., Liu, X. & Atasoy, B. (2025). Operational Synchromodal Transport Planning Methodologies: Review and Road Map, Transportation Research Part E, 194.

IPIC 2026
12th International Physical Internet Conference
June 8-10, 2026 | Bordeaux, FRANCE
Physical Internet Center

Physical Internet Across Sovereign Borders: A Formalized Design Framework and the Price of π-Sovereignty

Oscar Hasburn-Babich[1,2], Craig Tovey[2], and Benoit Montreuil[1,2]
1. Georgia Tech Physical Internet Center, Supply Chain & Logistics Institute
2. H. Milton Stewart School of Industrial and Systems Engineering
Georgia Institute of Technology, Atlanta, GA, USA
Corresponding author: babich@gatech.edu

Abstract: *Why would directly competing firms join open Physical Internet (PI) networks—relinquishing proprietary infrastructure, exposing operational patterns, and introducing coordination dependencies vulnerable to manipulation? And if they join, do sovereign borders necessarily weaken, or can they strengthen network viability under coordination? Fifteen years of PI research have convincingly argued from logistics efficiency, resilience, and sustainability (Montreuil, 2011) yet have lacked formal tools to determine when such participation is individually rational, when it is welfare-enhancing, or when sovereignty enables coordination that unrestricted interaction cannot sustain. We provide, to our knowledge, the first formalized framework that reframes these open questions from advocacy to design. It integrates three components: 1) a mapping of jurisdictional pairs to admissible cross-border flows; 2) a disclosure hierarchy governing π-actor exposure; and 3) the Price of π-Sovereignty (PoS), defined as the ratio of best jurisdiction-free to best sovereign equilibrium welfare. PoS = 1 means sovereignty has no effect; PoS > 1 means it erodes welfare; PoS < 1 means it enhances welfare. On an illustrative network connecting Shanghai, Rotterdam, and New York, a single regulatory asymmetry drives PoS = 0.8 at the welfare-maximizing disclosure level: sovereignty thus increases network welfare by 25%.*

Keywords: *Physical Internet, Price of π-Sovereignty, disclosure level, π-actor, governance, open network design, cross-border logistics, welfare creation, regulatory heterogeneity.*

Physical Internet (PI) Roadmap Fitness*: ☒ PI Networks, ☒ System of Logistics Networks, ☒ Access and Adoption, ☒ Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☒ In-Person presentation.*

1 Introduction

The Physical Internet (PI, π) vision (Montreuil, 2011) proposes an open, globally hyperconnected logistics system built on modularity, open routing, decentralized coordination, and universal interconnection. Fifteen years of PI research have convincingly argued from logistics efficiency, resilience, and sustainability, without formalizing the conditions under which competing firms rationally join open networks across sovereign jurisdictions or whether sovereignty enables coordination that no jurisdiction-free design can sustain. PI exhibits superior performance in many ways, but it is unclear when π-actors —firms operating within PI networks— who are fierce competitors can nonetheless materially benefit if their participation is in a strategic equilibrium and feasible under sovereignty constraints.

The horizontal collaboration literature documents that logistic firms view proprietary infrastructure as a competitive advantage that *"simply must not be shared"* and consider

cooperation with direct competitors *"inconceivable"* (Plasch et al., 2017). The CO3 project concluded that horizontal collaboration requires *"a neutral trustee"* to establish sustainable collaboration between competing parties (Cruijssen, 2012). PI experts have identified governance as the most likely bottleneck for adoption (Fahim et al., 2021).

These concerns are the long-standing barriers our contribution addresses.

Throughout, we distinguish *jurisdictional sovereignty* —the exogenous regulatory regime: what each jurisdiction permits to cross its borders— from *firm sovereignty* —the strategic participation decision: whether and how it joins an open network or operates proprietary routes. The framework takes the first as given and analyzes the second as an equilibrium outcome.

A firm running its own supply chain across jurisdictions —or its Physical *Intra*net— already faces jurisdictional sovereignty costs: operational frictions it absorbs and classically optimizes around. There is no coordination equilibrium because there is no other player. The moment that a firm joins PI, two things change at the lowest structural level: it is now coordinating with others —consolidating scheduling, pooling container flows across borders— and PI's open interfaces create cross-firm data flows that sovereignty constrains. Participation gives the firm something it did not have before: *the option to* ***defect***. A π-actor can defect from equilibrium coordination by dropping to autarky, generating no flows. If its compliance cost on a disclosure category —e.g. AI model weights under the EU AI Act— exceeds its equilibrium step-up, it exercises that option, not because PI is inefficient, but because the sovereignty cost of PI's openness exceeds the surplus that holds π-actors at the table. In the worked example (Section 3), this is a π-actor deciding the juice is not worth the squeeze. The mechanism is not sovereignty alone (firms already handle it) and not PI alone (single-jurisdiction PI faces no sovereignty). It is sovereignty acting on PI's open, multi-firm interfaces, creating a wedge between the coordination and autarky payoffs that the jurisdiction-free game did not have: what we call the *Price of π-Sovereignty (PoS)*.

Our contribution provides the first formal game-theoretic framework that addresses this structural interaction and yields three results. The equilibrium concept is generalized within a *Sovereignty-Partitioned Game (SPG)* (Hasburn-Babich et al., 2026a): a non-cooperative game under exogenous institutional authority and regulatory regimes that alter player action sets and payoffs, with the decomposition of base-game incentives, jurisdictional constraints, and institutional modifications maintained as a structural primitive. In the worked example, Nash and Stackelberg equilibria are applied, while the sovereignty partition operates through the payoff channel only.

First, ***(Participation breakdown)*** the framework explains when the answer to *"should π-actors join a cross-border network?"* is *no*, even when the network is operationally superior by a large margin. Sovereignty acting on open interfaces can destroy coordination equilibria whose costs the overlaying logistics surplus would otherwise absorb, because the sovereignty costs of participation —compliance on cross-jurisdictional flows— exceed the coordination benefit.

Second, ***(Sovereignty-enabled coordination)*** the framework explains when the answer is not only *yes* but *yes, precisely because of sovereignty*. We show that PoS < 1 is achievable. Sovereignty can ban a deviation strategy outright: if antitrust regulation removes the predatory pricing that would let one π-actor undercut another after learning its routing patterns, the ban stabilizes cooperation that the jurisdiction-free game could not sustain (Hasburn-Babich et al., 2026a). Sovereignty can degrade proprietary alternatives more than network cooperation: restrictions on a critical bilateral link can block a defection strategy, creating a coordination equilibrium that no jurisdiction-free design sustains. π-actors join not despite sovereignty but because π-sovereignty makes defection *unprofitable*.

Third, *(Optimal disclosure level)* the framework tells network designers where to set the dial. The disclosure level hierarchy admits a welfare-maximizing level that is neither insufficient for coordination nor excessive sovereignty exposure. The framework identifies it —or determines that no viable level exists, in which case the honest answer is that the network should not be built under existing conditions— before infrastructure is committed.

These results reframe the PI adoption question from advocacy to design: *"at what disclosure level does participation become an equilibrium?"*. The framework enables the answer.

SPG foundations are developed in Hasburn-Babich et al. (2026a); regulatory composition and welfare bounds in Hasburn-Babich et al. (2026b); and the PI instantiation for dynamic network settings, in Hasburn-Babich et al. (2026c), where networks are expressed as sequences of static instances, each representable with the machinery of Sections 2–3. SPG relates to established game-theoretic and economic coordination models which partition the strategic environment along different dimensions. In SPG, it is partitioned by exogenous institutional authority: jurisdiction-indexed strategy sets and payoff structures arise as *inherited, composable institutional constraints*, not as objects chosen by players, designed by a mechanism designer, or resolved by information exchange.

Our paper introduces the static PI instantiation framework and a trilateral network example.

1.1 Why Opt In: Four Barriers Addressed

The PI adoption literature identifies three structural barriers to participation by competing firms, and a fourth institutional concern that crosscuts all three. The framework addresses each.

Proprietary advantage is measured, not assumed. The conventional argument treats proprietary infrastructure as an unquantified competitive moat. The framework makes the comparison explicit. A firm operating proprietary cross-jurisdictional routes faces sovereignty costs —export controls, customs duties. The same firm as π-actor in a network faces additional costs from cross-firm data flows but also gains a coordination benefit from shared routing and bilateral operational alignment. The *participation margin* μ_i computes the comparison firm by firm: if $\mu_i \geq 0$, the proprietary alternative is not a profitable deviation from participation, and its advantage is a priced quantity rather than an unquantified veto. Whether μ_i stays positive depends on the natural coordination benefit and on whether sovereignty taxes deviation more than cooperation. The PoS diagnoses whether sovereignty strengthens or weakens the network.

Exposure is bounded and optimizable. Concern about operational exposure assumes ungoverned information flow. The disclosure level hierarchy provides that governance. At each level, the admission set specifies which categories cross which jurisdictional borders. The π-actor sees the admission set before committing. Exposure is a design variable to be optimized. The operational protocols for managing flows within the governed topology —hub consolidation, dispatch, and dynamic directional routing— already exist (Shaikh et al., 2025); the orchestration of sovereign PI networks through performance commitments (Shaikh et al., 2026) complements the disclosure hierarchy by governing what operational data networks expose during coordination.

Defection is blocked by the game structure. In the base game without sovereignty, a π-actor can earn more by declining the network and operating proprietary cross-jurisdictional routes than by joining —the defection incentive that destabilizes coordination. If sovereignty degrades the proprietary alternative more than it degrades network cooperation the deviation becomes unprofitable. PI adoption does not require trust between competitors, but a governance regime removing the strategies that make trust necessary.

Neutrality of the operator is replaced by incentive compatibility. The framework shows that what matters is not who runs the infrastructure, but *what the game structure permits*. If the disclosure level bounds what is visible, and sovereignty constrains what crosses borders, and the equilibrium is individually rational for each π-actor, then the operator's identity is irrelevant to the incentive structure.

2 Design Framework for Sovereign PI Networks

The framework comprises formal objects that answer specific network design questions. Figure 1 maps the pipeline from design inputs to network viability diagnosis.

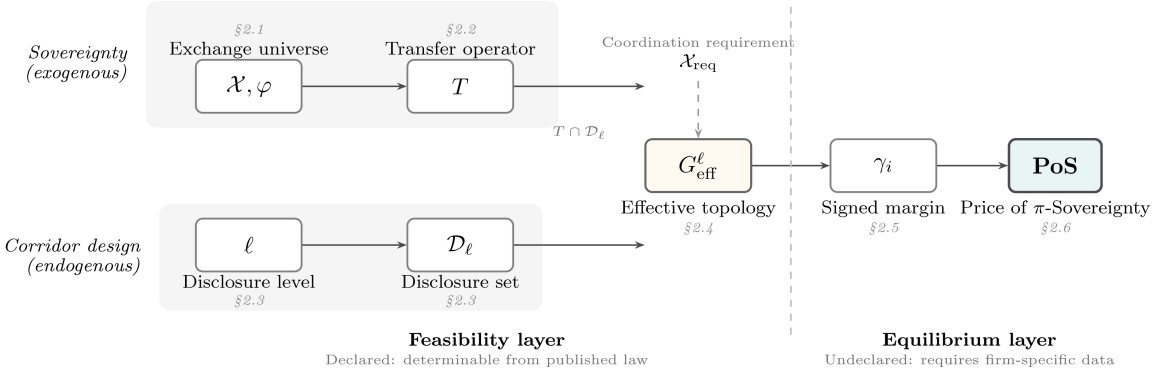


Figure 1. *Framework pipeline.*
*The **feasibility layer** determines which bilateral links exist from published law; the **equilibrium layer** determines whether participation pays from firm-specific data. Sovereignty and network design jointly set the effective topology: a link exists if the transfer operator permits it and the disclosure demands it.*

2.1 What must cross jurisdictional boundaries?

The *exchange universe* X is the catalog of everything that may need to cross a jurisdictional boundary for coordination to function. We decompose X into axes reflecting the independence of the authorities that govern them. A single cross-border shipment may clear customs (physical axis), pass data transfer assessment (data axis), and settle payment (financial axis) under distinct authorities. Within each, X is partitioned into categories: minimally governed units of exchange whose admissibility, degradation, or prohibition may vary across jurisdictions. The partition granularity depends on the institutional environment and the resolution required for the network under study. A *sensitivity index* φ ranks each category by intensity, from level 1 (least restricted) to level 4 (most restricted). Table 1 shows an illustrative baseline instantiation.

Table 1: *Sensitivity ordering for PI exchange universe (φ = 1 least, 4 most restricted)*

φ	Cyber (Data)	Physical	Financial	Human	Intellectual Property (IP)
1	d_1 tracking & telemetry	p_1 general cargo	f_1 settlement	—	—
2	d_2 trade compliance	p_2 certified goods	f_2 insurance	h_1 credentials	q_1 standards
3	d_3 operational governance	p_3 licensed goods	f_3 payment	h_2 expertise	q_2 patents
4	d_4 proprietary routing	p_4 regulated goods	f_4 investment	h_3 personnel	q_3 trade secrets

The number of sensitivity tiers per axis reflects the regulatory gradation. Not all axes may exhibit the same number of tiers; the framework accommodates heterogeneous tier structures. The instantiation shown is illustrative —

human and IP categories purposely begin at 2 (they are not as freely transferable as basic operational data or general cargo); a specific network may reveal finer or coarser gradations depending on the jurisdictions involved.

2.2 What does each jurisdiction permit to cross?

A *transfer operator* T maps each ordered pair of jurisdictions to the set of categories permitted to flow. It is generically asymmetric: what the EU permits toward the US need not equal what the US permits toward the EU. The operator distinguishes binary permission (whether a category can cross at all) from degradation (the compliance cost conditional on permission). A tariff on cross-border containers degrades payoffs ($\zeta > 0$) without severing the link; an embargo severs it (the category exits T). T preserves the axis decomposition. Examples illustrate:

- *Cabotage (United States–Mexico–Canada Agreement, USMCA).* A Mexican or Canadian carrier cannot move freight within the US.
- *SWIFT disconnection (EU–Russia).* A network may be physically and informationally connected yet operationally disconnected because the financial axis is severed.

2.3 How much cross-boundary exchange does the network demand?

The *disclosure level* ℓ is the take-it-or-leave-it governance parameter. At level ℓ, the network's governance rules require that every category with sensitivity $\varphi \leq \ell$ be available for cross-jurisdictional flow. The *disclosure set* at level ℓ, D_ℓ, is the informational precondition for network participation: what a π-actor must be willing to expose. In the worked example (Section 3), the four levels progress to full transparency. At $\ell = 0$ the network is disconnected.

A category *c* crosses a jurisdictional boundary if two conditions are met: sovereignty permits it (*c* is in T(j,j′)) *and* the network's governance requires it (*c* is in D_ℓ). The *effective transfer* at level ℓ is the intersection: $T_\ell(j,j') = T(j,j') \cap D_\ell$. Raising the disclosure level cannot overcome a sovereignty prohibition. If sovereignty blocks a category, the network designer cannot force it through. *Sovereignty sets the ceiling; the disclosure level sets the floor.*

2.4 Which bilateral links actually work for a given service?

Every network service has a *coordination requirement* X_{req}: the minimum set of exchange categories that must flow across every bilateral link for the service to function. A container tracking service requires operational data; a customs-integrated service requires both operational and trade compliance data. The *effective topology* at level ℓ for a service requiring X_{req} is a directed graph over jurisdictions: an edge from j to j′ exists if and only if $X_{req} \subseteq T_\ell(j,j')$. A single missing category severs the edge. The same jurisdictions produce different graphs for different services—and different graphs at different disclosure levels.

2.5 Does the coordination benefit exceed the compliance cost for each firm?

The *participation margin* μ_i answers whether π-actor *i* stays in the network or exits: it is the difference between *i*'s sovereign payoff at the candidate equilibrium σ* and its best deviation. If $\mu_i \geq 0$ for every firm, coordination survives as an equilibrium and the proprietary cross-jurisdictional operation alternative is dominated. If any firm's μ_i turns negative, coordination is destroyed. Whether μ_i stays positive depends on whether sovereignty taxes the proprietary deviation more than network coordination. The worked example (§3.3) demonstrates the mechanism: π-actor$_{CN}$'s defection is degraded more than its coordination, raising μ_{CN} above zero and stabilizing the welfare-maximizing equilibrium that the jurisdiction-free game could not sustain. Two firms on the same network, under the same regulations, at the same disclosure level, can face opposite survival outcomes because their compliance costs differ.

2.6 The Price of π-Sovereignty

The *Price of π-Sovereignty (PoS)* captures the best equilibrium welfare without sovereign constraints relative to the best equilibrium welfare that emerges when π-actors operate under jurisdictional regimes. PoS = 1 means sovereignty has no effect; PoS > 1 means it erodes welfare; PoS < 1 means it enhances welfare. Whether PoS exceeds or falls below unity on a given network is a *sovereignty-effect diagnostic*. The PoS admits a multiplicative decomposition: a ***selection effect*** (does sovereignty shift which equilibrium is played?) and a ***degradation effect*** (does compliance extract value from the equilibrium that is played?):

$$\textbf{\textit{PoS}}(\ell) = [W^*(\Gamma) / W(\hat{o})] \times [W(\hat{o}) / W^S(\hat{o})] \qquad (1)$$

where W*(Γ) is the best base-game equilibrium welfare, ô is the welfare-maximizing sovereign equilibrium, W(ô) is its base-game welfare, and W^S(ô) is its sovereign welfare. The decomposition is *root-cause diagnostic*. A key result —the *Disclosure-Sovereignty Tradeoff*— establishes that sovereignty exposure and effective connectivity are both monotone in ℓ, but PoS is not. More openness is not always better; less is not always safer.

2.7 Operational Correspondence with PI Natives

We illustrate correspondence on π-containers, π-hubs, and π-networks. The **π-container** is PI's physical unit of flow, carrying physical goods, data payloads, and financial settlement triggers. It is compliant when T admits every category it bundles on every directed edge it traverses.

π-hubs —physical nodes for routing and consolidation— play a key role under the framework: *governance intermediary*. In the worked example, Rotterdam enables data flow that the direct channel prohibits. *Hub selection is also a regulatory question, not just a geographic one.* SPG determines the effective topology; directional routing protocols such as Shaikh et al. (2025) would navigate remaining edges where Rotterdam is the viable intermediary.

The PI literature describes **π-networks** along physical, operational, transactional, personal, legal, and informational connectivity. The framework adds *governance connectivity*, determined by T and the jurisdiction set J. The effective topology G_{eff} is *the π-network that exists after sovereignty acts on it*. A link that is physically, operationally, and informationally connected can be effectively disconnected —the SWIFT example severs the financial axis while other links remain intact. PI network design is a multi-layer problem, and the governance layer, unlike physical infrastructure, can change overnight through regulatory action.

3 Shanghai–Rotterdam–New York: A Worked Example

We illustrate on a trilateral network connecting China (Shanghai), the EU (Rotterdam), and the US (New York/Newark). The network exhibits regulatory heterogeneity across all bilateral pairs and a critical asymmetry that drives the entire welfare creation result. For expositional clarity, the illustration isolates the cyber axis where the asymmetry bites; the general multi-axis case is treated in Hasburn-Babich et al. (2026c). The structure of Table 1 applies to networks.

3.1 Setup: three π-actors and the transfer operator

Three π-actors, one per region, evaluate whether to join the network. Each chooses between Open (O: join, share data per the disclosure level) and Closed (C: operate proprietary routes, no network data sharing). The network's coordination requirement is $X_{req} = \{d_1, d_2\}$: d_1 (tracking & telemetry) and d_2 (trade compliance). Table 2 shows the pairwise transfer structure. The number of permitted categories on each directed edge reflects the regulatory posture of the sending jurisdiction.

Table 2: *Transfer operator T (cyber axis, data categories)*

From \ To	→ EU	→ US	→ CN
EU →	—	$\{d_1, d_2, d_3\}$	$\{d_1, d_2\}$
US →	$\{d_1, d_2, d_3, d_4\}$	—	$\{d_1, d_2\}$
CN →	$\{d_1, d_2\}$	$\{d_1\}$	—

Critical asymmetry: the CN→US edge permits only $\{d_1\}$ —operational tracking. Personal Information Protection Law (PIPL) and the Data Security Law impose compliance procedures on outbound transfer of trade compliance data (d_2) to the US, which the illustration treats as de facto severance for expositional clarity; in practice, the burden functions as degradation. No comparable US restriction applies in the reverse direction. This single asymmetry drives the entire welfare creation result.

3.2 Effective topology across disclosure levels

At **level 1** ($D_1 = \{d_1\}$): the coordination requirement $X_{req} = \{d_1, d_2\}$ demands d_2, but $d_2 \notin D_1$. No edges exist. The network is disconnected. At **level 2** ($D_2 = \{d_1, d_2\}$): five of six directed edges open. The sole missing edge is CN→US: PIPL blocks d_2 outbound to the US. Shanghai can receive from both Rotterdam and New York, and can send to Rotterdam, but cannot send trade compliance data directly to New York. At **levels 3 and 4**: connectivity is unchanged (X_{req} only needs d_1, d_2, already available at level 2).

3.3 Payoffs and the welfare creation mechanism

Payoffs are illustrative, calibrated to demonstrate the mechanism; Section 4 presents a methodology for eliciting real-world payoffs. The key pattern in Table 3 is CN's defection incentive: CN gains 2 by defecting from trilateral coordination, making (O, O, O) unstable. Sovereignty reverses this by degrading the defection payoff more than the coordination payoff. Table 3 shows the base-game payoff structure.

Table 3: *Base-game payoffs (EU, US, CN), 3 players choosing O or C, 8 strategy profiles. Nash and Stackelberg Equilibria (NE, SE) marked. (O, O, C) is SE under any single-leader ordering.*

Profile	EU	US	CN	Total	NE	SE
(O, C, C)	2	3	4	9		
(C, O, C)	3	2	4	9		
(C, C, O)	3	3	3	9		
(C, C, C)	**3**	**3**	**4**	**10**	✓	

Profile	EU	US	CN	Total	NE	SE
(O, O, O)	12	10	8	30		
(O, O, C)	**6**	**4**	**10**	**20**	✓	✓
(O, C, O)	8	3	5	16		
(C, O, O)	3	7	6	16		

The base game without sovereignty has two pure-strategy Nash equilibria: (C, C, C) with welfare 10 —a coordination trap in which no firm joins because joining alone is unprofitable— and (O, O, C) with welfare 20, in which EU and US coordinate bilaterally while CN defects. The welfare-maximizing outcome (O, O, O) yields welfare 30 but is not an equilibrium: CN's proprietary route through the restricted CN→US channel yields 10, exceeding its trilateral payoff of 8. The defection incentive is visible in Table 3: CN gains 2 units (from 8 to 10) by defecting from trilateral coordination. No other player has a profitable deviation at (O, O, C), making it the Pareto-superior NE.

Under sovereignty at level 2, the missing CN→US edge transforms the game. CN's proprietary Shanghai–New York route requires trade compliance data (d_2) on the CN→US edge, but T(CN,

US) = {d_1} does not contain d_2. The proprietary route is severely degraded: CN's payoff from defection drops from 10 to 3. Meanwhile, the network routes CN's data through Rotterdam (CN→EU→US), bypassing the restricted direct channel. A proprietary route through Rotterdam would require coordination with EU-based infrastructure, precisely the participation CN declines. CN's network payoff is degraded from 8 to 5 —but 5 exceeds the degraded proprietary payoff of 3. The defection incentive is reversed. Of the 5 total degradation units at (O, O, O) under level 2, CN absorbs 3 (from 8 to 5) and the EU–US pair absorbs 2 in compliance costs on the bilateral edges, yielding sovereign welfare 25. Base-game equilibrium (O, O, C) is destroyed: CN's degraded proprietary payoff of 3 falls below its network payoff of 5, so CN deviates to Open. *Sovereignty simultaneously destroys the defection equilibrium and provides the welfare-maximizing one.* The coordination trap (C, C, C) persists under sovereignty — autarky generates no cross-border flows and therefore no degradation— but the created equilibrium (O, O, O) Pareto dominates it: every π-actor earns strictly more under sovereign trilateral coordination (11, 9, 5) than under autarky (3, 3, 4). Sovereignty thus resolves both failures of the base game: it eliminates the defection equilibrium by degrading CN's proprietary route, and it renders the coordination trap irrelevant by creating a strictly superior alternative.

Rotterdam is a *governance intermediary*: it enables data flows between Shanghai and New York to bypass the restrictive direct channel. The intermediary role arises from asymmetry in transfer permissions —PIPL permits d_2 outbound to the EU but restricts it to the US; the EU→US edge permits d_2 under the Data Privacy Framework— not from port infrastructure (§2.7).

3.4 The PoS trajectory

Table 4 summarizes the PoS trajectory. Figure 2 displays it with the institutional topology.

Table 4: *PoS trajectory across disclosure levels. Sovereign Nash Equilibrium (NE) and welfare.*
** Created NE (not a NE in the jurisdiction-free base game).*
***At ℓ=4, π-actor$_{CN}$'s sovereign coordination payoff falls below its autarky payoff, destroying the equilibrium*

ℓ	Sovereign NE	$W^{\pi S}$	PoS	Selection	Degradation	Regime
1	(C, C, C)	10	2.00	2	1	Disconnected
2	**(O, O, O)***	**25**	**0.80**	**2/3**	**6/5**	**Welfare enhancement**
3	(O, O, O)*	22	0.91	2/3	15/11	Welfare enhancement
4	(C, C, C)	10	2.00	2	1	Collapse**

The welfare-maximizing level is ℓ* = 2, where PoS = 0.80: sovereignty increases network welfare by 25% over the best jurisdiction-free equilibrium. The non-monotonicity is the *Disclosure-Sovereignty Tradeoff* made visible: sovereignty exposure rises monotonically with ℓ as more categories are exposed; effective connectivity also rises monotonically as more edges open; but welfare is non-monotone —the interaction produces a maximum at an interior level, with welfare declining beyond it as exposure costs outpace connectivity gains. Applying decomposition (1) at level 2:

$$PoS(2) = [20/30] \times [30/25] = [2/3] \times [6/5] = 4/5 = \mathbf{0.80}$$

The selection effect (2/3 < 1) reflects sovereignty enabling a superior equilibrium —the welfare-maximizing trilateral coordination that defection prevented. The sovereign game sustains (O, O, O) with base welfare 30, while the base game could only sustain (O, O, C) with base welfare 20. The degradation effect (6/5 > 1) reflects compliance costs extracting 5 units from the created equilibrium. The selection benefit is larger than the degradation disbenefit from compliance costs.

The degradation effect escalates from 6/5 to 15/11 as the EU–US network absorbs governance-level compliance costs (d_3), but not enough to overtake the selection benefit. At level 4, sovereignty costs exceed the coordination benefit for the weakest π-actor: the juice is not worth the squeeze, and the network collapses.

4 From Framework to Field: Eliciting Network Parameters

The framework's objects partition into two classes with sharply different observability. The transfer operator T and jurisdiction set J are ***declared***: determinable from published law —trade agreements, data adequacy decisions, export control schedules. The degradation ζ, the participation margin μ, and the survival regime are ***undeclared***: determinable from firm-specific operational data. No authority publishes interaction costs between GDPR compliance and customs compliance because no single authority controls both axes. This partition is a structural feature of sovereignty itself. The framework's contribution is to identify precisely what must be learned —the undeclared objects: firm-specific degradation, interaction costs, and participation margins— and from what sources.

Natural experiments. Sovereignty events with before/after observability provide direct identification of degradation on real networks. Brexit is paradigmatic: pre-January 2021, the UK was inside the EU single market ($|J| = 1$, PoS = 1); post-January 2021, $|J| = 2$ with the same firms, technology, and base-game payoffs. Post-Brexit compliance cost surveys report 8–9% non-tariff barrier costs (Barfield, 2022) —direct estimates of ζ for specific firm classes. SWIFT disconnection of Russian banks (financial axis severance) and US–China Section 301 tariffs (physical axis degradation) provide analogous identification.

Network operational data. PI pilot programs —Alliance for Logistics Innovation through Collaboration in Europe (ALICE), EU-funded initiatives, Trans-European Transport Network (TEN-T)— generate operational data on coordination benefit (cost savings from joint routing and scheduling versus independent operation) and compliance costs ζ (customs processing, data transfer assessments, rules-of-origin documentation). The framework tells practitioners which operational metric maps to which formal object: compare net network performance against the proprietary alternative at disclosure levels —the difference is μ, and its sign the survival regime.

Expert elicitation. For prospective networks without historical data, the network planner needs two numbers per firm per disclosure level: coordination benefit *("What is your estimated annual cost saving from shared scheduling with your network partner?"*) and compliance cost *("What is your estimated annual cost for transferring routing metadata across the EU–UK border?"*). The participation margin is the difference.

5 Related Work

The PI literature (Montreuil, 2011; Ballot et al., 2012; Sarraj et al., 2014; Pan et al., 2017; Boysen et al., 2025) develops the vision and operational theory of open logistics. Boysen et al. (2025) find that split transports offer negligible practical benefit, bearing on the routing dimension of network governance —the worked example's effective topology (Section 3.2) formalizes when and why split routing across jurisdictions is structurally constrained rather than merely inefficient. Lafkihi et al. (2019) review freight transportation service procurement but without jurisdictional constraints. Li et al. (2025) propose a six-level National Digital Twin (NDT) architecture; the sovereignty layer extends their framework along a governance dimension. This paper concisely introduces SPG (Hasburn-Babich et al., 2026a) for PI systems across heterogeneous institutional environments, a setting PI's doctrine generates but the existing literature had not formalized. Cohen and Huchzermeier (1998) study global supply chain management under operational and financial uncertainty, but without institutional

topology or equilibrium characterization. Generalized Nash equilibrium problems (Facchinei and Kanzow, 2010) can embed sovereignty constraints but the embedding discards the structure that makes SPG operational.

6 Conclusion

The framework reframes PI adoption from advocacy to design. It asks at what disclosure level participation becomes an equilibrium —and provides the machinery to answer. The practical methodology for network designers is concrete. Map the coordination requirement for each proposed service. Compute the task-dependent effective topology. Evaluate firm-specific participation margins. Identify the disclosure level at which participation becomes an equilibrium or determine that no viable level exists.

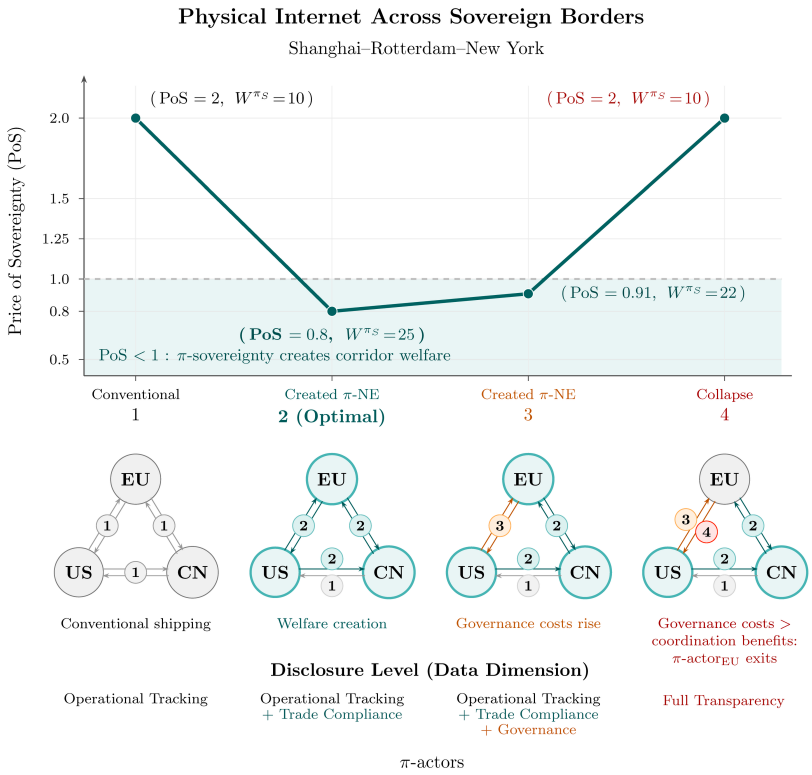


Figure 2. *Price of π-Sovereignty trajectory and institutional topology across four disclosure levels.*

The Shanghai–Rotterdam–New York network example shows that sovereignty can enhance network welfare. A single regulatory asymmetry blocks the defection strategy that destabilizes welfare-maximizing trilateral coordination. Rotterdam's role as a governance intermediary, routing data through the EU to bypass the restricted direct channel, demonstrates that the institutional topology, not geography, determines the viable path.

The framework reveals a structured learning need: the laws governing cross-jurisdictional transfer are publicly declared and readable in a day; the firm-specific margins determining whether coordination survives must be learned from operational data. Natural experiments (Brexit, SWIFT, Section 301 tariffs) and PI network pilots provide the empirical basis.

The question for this community is *"at what disclosure level does participation become an equilibrium?"*. The companion machinery (Hasburn-Babich et al., 2026a,b,c) extends the framework to regulatory composition and multi-network, multi-period PI instantiation.

References

- Ballot E., B. Montreuil, R.D. Meller (2012): The Physical Internet: The Network of Logistics Networks. Bureau de la Recherche, Mines ParisTech.
- Barfield R. (2022): The new costs for business post-Brexit. UK in a Changing Europe, 1 December 2022. Available at: https://ukandeu.ac.uk/the-new-costs-for-business-post-brexit/ (accessed May 10, 2026).
- Boysen N., D. Briskorn, B. Montreuil, L. Zey (2025): The π-transportation problem: On the value of split transports for the Physical Internet concept. European Journal of Operational Research, 324, 629–643.
- Cohen M.A., A. Huchzermeier (1998): Global Supply Chain Management: A Survey of Research and Applications. In: Tayur S., Ganeshan R., Magazine M. (eds.), Quantitative Models for Supply Chain Management. Kluwer Academic Publishers, Boston, MA, 669–702.
- Cruijssen F. (2012): CO3—Collaboration Concepts for Co-modality. Final Report, European Commission FP7 Project, Brussels, Belgium.
- Facchinei F., C. Kanzow (2010): Generalized Nash equilibrium problems. Annals of Operations Research, 175(1), 177–211.
- Fahim P.B.M., M. Martinez de Ubago Alvarez de Sotomayor, J. Rezaei, A. van Binsbergen, M. Nijdam, L. Tavasszy (2021): On the evolution of maritime ports towards the Physical Internet. Futures, 134, 102834.
- Hasburn-Babich O., C. Tovey, B. Montreuil (2026a): Sovereignty-Partitioned Games. Working paper, H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology.
- Hasburn-Babich O., C. Tovey, B. Montreuil (2026b): Regulatory Composition in Sovereignty-Partitioned Games: Feasibility and Welfare Bounds. Working paper, H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology.
- Hasburn-Babich O., C. Tovey, B. Montreuil (2026c): Physical Internet and Jurisdictional Sovereignty: Game-Theoretic Foundations. Working paper, H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology.
- Lafkihi M., S. Pan, E. Ballot (2019): Freight transportation service procurement: A literature review and future research opportunities in omnichannel E-commerce. Transportation Research Part E, 125, 348-365.
- Li M., T. Le, L. Tavasszy (2025): Six-maturity-level national digital twins for logistics and supply chain systems. SSRN Electronic Journal. Available at: https://ssrn.com/abstract=5183145 (accessed May 10, 2026).
- Montreuil B. (2011): Toward a Physical Internet: Meeting the global logistics sustainability grand challenge. Logistics Research, 3(2–3), 71–87.
- Pan S., E. Ballot, G.Q. Huang, B. Montreuil (2017): Physical Internet and interconnected logistics services. International Journal of Production Research, 55(9), 2603–2609.
- Plasch M., S. Pfoser, M. Gerschberger, R. Gattringer, O. Schauer (2017): From Horizontal Collaboration to the Physical Internet—A Case Study from Austria. International Journal of Transport Development and Integration, 1(2), 129–136.

- Sarraj R., E. Ballot, S. Pan, D. Hakimi, B. Montreuil (2014): Interconnected logistic networks and protocols. International Journal of Production Research, 52(11), 3185–3208.
- Shaikh S.J., P. Muthukrishnan, Y. Lai, B. Montreuil (2025): Dynamic Directional Routing of Freight in the Physical Internet. arXiv preprint arXiv:2504.02722.
- Shaikh S.J., B. Montreuil, A. Erera, T. Zhang, M. Spanuth, P. Nourmuhammadi (2026): Orchestrating Sovereign Logistics Networks Through Performance Commitments in the Physical Internet. Proceedings of the 12th International Physical Internet Conference (IPIC 2026), Bordeaux, France.

IoT Enabled Smart Dry Containers as Foundational Infrastructure for the Physical Internet: A Standards-Aligned Architecture for Fire Risk and Door Event Detection in Multimodal Transport

Marianna Levtov, PhD
Hoopo Systems, Zurich, Switzerland
Corresponding author: marianna.l@hoopo.tech

Abstract

The digitalisation of containerised logistics has accelerated through the deployment of IoT-enabled devices on dry containers, yet their systemic role remains insufficiently theorised. Existing literature predominantly treats smart containers as visibility tools rather than as infrastructural components within logistics systems. This article reconceptualises IoT-enabled smart dry containers as cyber-physical infrastructure nodes within the Physical Internet, transforming physical container states into standardised and interoperable event signals. The study adopts a conceptual and integrative research design, synthesising engineering studies, maritime digitalisation literature, and international standards frameworks, including ISO, UN/CEFACT, and DCSA. Three functional domains are analysed: track-and-trace, door event detection, and fire-related anomaly monitoring. The findings demonstrate asymmetrical while fire detection remains emergent. The article develops a standards-aligned event architecture and argues that smart container value depends on semantic interoperability, governance, and system integration rather than sensing capability alone.

Keywords

Physical Internet, smart containers, IoT, cyber-physical systems, multimodal logistics, data governance

***Targeted Delivery Mode-s**: ☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

1. Introduction

The standard freight container is one of the most consequential infrastructural innovations in global trade, enabling the modularisation and scalability of international logistics systems (Levinson, 2016). While containerisation significantly reduced physical friction through standardised handling and intermodal compatibility, it did not resolve the persistent problem of informational opacity. Containers continue to move across maritime, terminal, rail, and road environments with limited real-time visibility of their condition, integrity, and operational status.

The emergence of IoT-enabled smart containers represents an attempt to address this gap. A smart container can be defined as a standardised freight container equipped with IoT hardware and software enabling visibility of operational data. However, current academic and industry discourse largely conceptualises these technologies as visibility-enhancing tools rather than as structural components of logistics systems. This framing fails to explain how container-level sensing contributes to system-wide coordination and does not sufficiently account for the role of standards, interoperability, and governance in enabling scalable deployment.

IoT-enabled smart dry containers should instead be understood as foundational infrastructure for the Physical Internet. Within this paradigm, containers evolve from passive transport units into cyber-physical nodes capable of generating standardised event signals that support

decentralised coordination across heterogeneous logistics actors (Cortes-Murcia et al., 2022). The central theoretical contribution of this article is to reconceptualise smart containers as distributed cyber-physical infrastructure nodes that transform physical container states into interoperable event signals enabling coordination, optimisation, and risk management across multimodal transport networks.

Two constraints shape the analysis. First, quantitative operational data remain largely unavailable due to confidentiality considerations, particularly in large-scale commercial deployments. Second, the analysis must therefore rely on engineering evidence, standards documentation, and bounded empirical illustrations. The paper proceeds by establishing a cyber-physical systems framework, reviewing relevant literature, analysing functional maturity, and developing a standards-aligned event architecture.

2. Theoretical Framing

2.1 Cyber-Physical Systems

A cyber-physical systems perspective extends beyond the conventional connected asset view. Within cyber-physical logistics systems, system-level value might not emerge directly from sensing capability itself, but from the contextualisation of physical container states reading into actionable events. This transformation occurs through multiple sequential layers comprising physical state detection, communication transmission, semantic encoding, standards-based event classification, system interpretation, and operational response.

A door opening, for example, only becomes operationally meaningful when the detected physical state transition is reliably transmitted, encoded according to shared event semantics, interpreted through comparison of the predefined normality and abnormality and linked to operational workflows (See Diagram 1). Consequently, smart container intelligence depends not solely on sensor capability, but on the coherence of the entire cyber-physical event chain.

This distinction shifts the analytical focus from device capability to system coherence. Logistics IoT systems operate across multiple layers, including perception, connectivity, processing, and application (Song et al., 2021). A location update or door alarm becomes operationally meaningful only when interpreted within an operational context and embedded within decision-making processes. Cyber-physical systems thinking, therefore, clarifies the difference between sensing capability and system-level value.

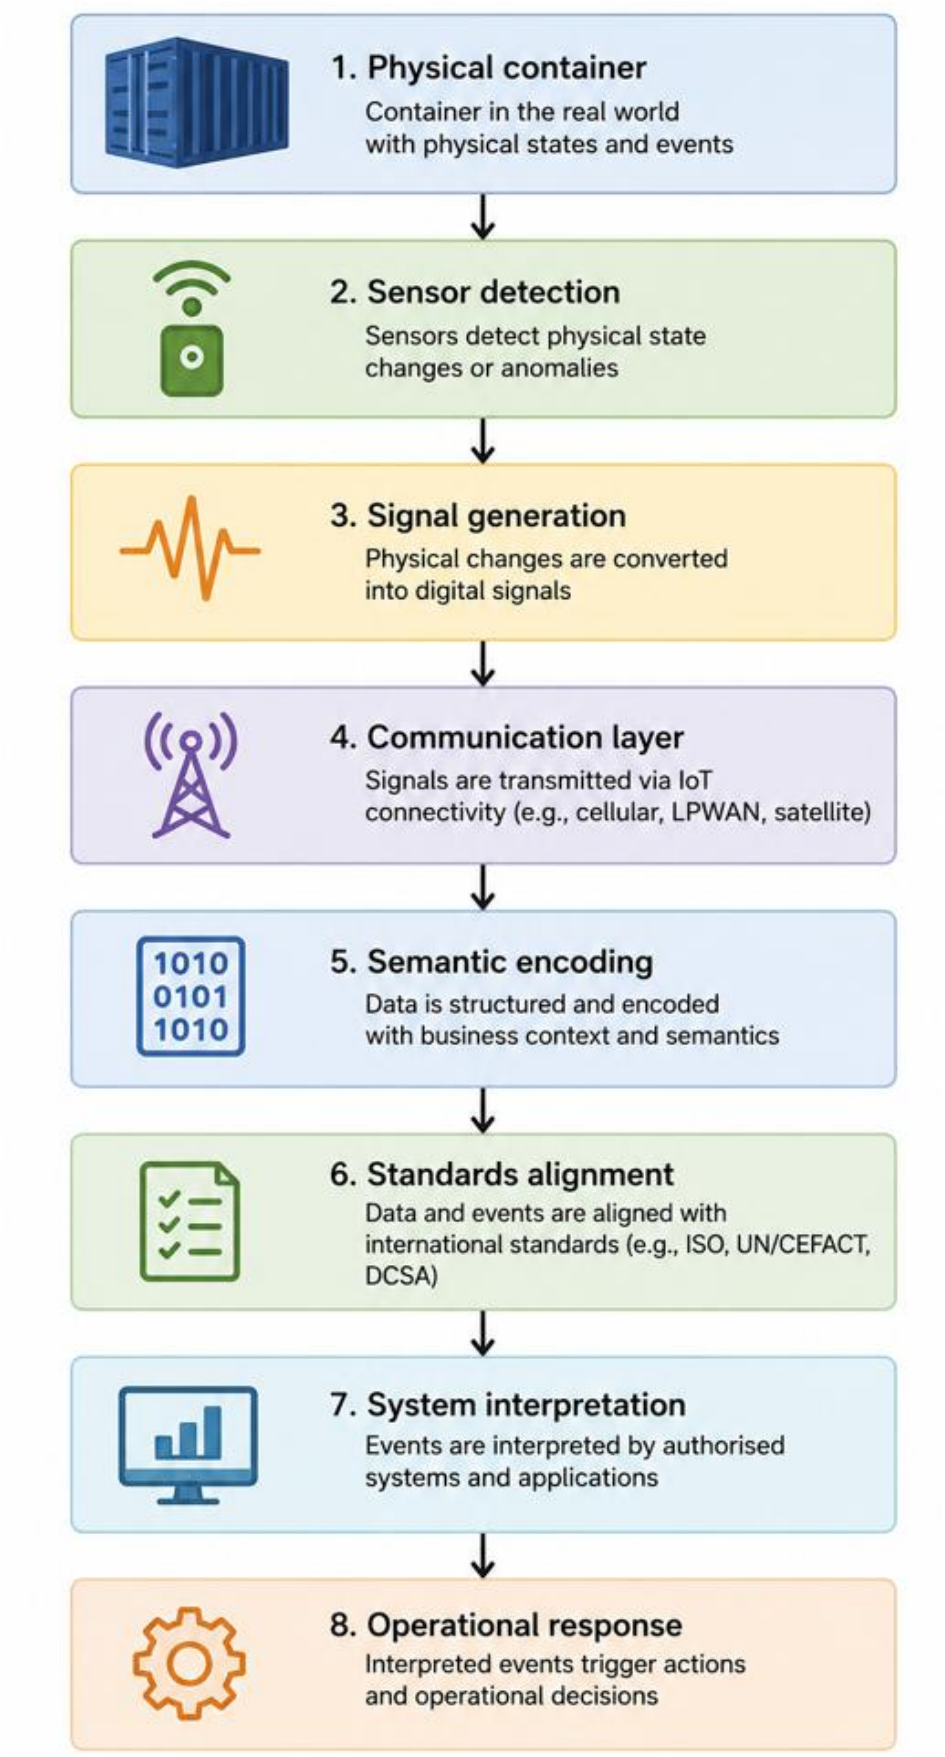


Diagram 1: Smart Container Logic

2.2 Informational Friction

Within cyber-physical logistics systems, system-level value does not emerge directly from sensing capability itself, but from the transformation of physical container states into interoperable and actionable events. This transformation occurs through multiple sequential layers comprising physical state detection, signal generation, communication transmission, semantic encoding, standards-based event classification, system interpretation, and operational response (Voorspuij & Becha, 2021).

A door opening, for example, only becomes operationally meaningful when the detected physical state transition is reliably transmitted, encoded according to shared event semantics, interpreted within authorised systems, and linked to predefined operational workflows. Consequently, smart container intelligence depends not solely on sensor capability, but on the coherence of the entire cyber-physical event chain.

In the multimodal container transportation, intermodal change, terminal dwell, route deviations, and operational anomalies often remain invisible or inconsistently captured. Informational friction arises where physical continuity is not matched by reliable and interoperable event visibility. Therefore, it could lead to delayed exception handling, inefficient asset utilisation, and limited coordination capacity. A smart container could reduce informational friction only when physical states are captured, transmitted reliably, and translated into events interpretable across systems. This reinforces the importance of analysing the full chain from physical state to operational action.

2.3 Physical Internet

The Physical Internet conceptualises logistics as a system of modular, interoperable, and coordinated flows (Cortes-Murcia et al., 2022). While containers already provide physical standardisation, digital legibility remains incomplete.

Smart containers address this gap by enabling the transformation of physical states into standardised event signals. Their significance lies not in enhancing the container itself, but in enabling distributed coordination within logistics networks.

3. Literature Review

3.1 Digitalisation in Logistics

Digitalisation research highlights that technological adoption alone does not produce transformation. Integration, organisational readiness, and system alignment determine value (Wang and Sarkis, 2021; Ghobakhloo et al., 2023). IoT systems in logistics operate as interconnected chains linking sensing, communication, processing, and application (Song et al., 2021). This implies that the value of IoT-enabled smart containers cannot be assessed at the level of sensing or connectivity alone but should be evaluated in terms of their integration into system-level processes that enable coordinated and actionable logistics operations.

3.2 Maritime Digitalisation and Fragmentation

Maritime logistics digitalisation remains uneven due to fragmented implementation across actors and systems (Raza et al., 2023). Despite widespread adoption of digital tools in the recent years, as it is shown, for example, by the Hapag-Lloyd and ZIM Shipping Line (Bruno, 2022; Uliel- Mirkin, 2024) system-level integration is constrained by heterogeneous infrastructures and governance complexity. This fragmentation implies that the deployment of digital technologies, including smart containers, does not automatically result in system-level

coordination, but instead exposes the necessity of interoperable architectures and governance mechanisms across supply chain actors to enable integration.

3.3 Smart Containers and Standardisation

Smart container research emphasises the role of global data exchange standards (Becha et al., 2021). At the level of formal international standardisation, ISO TC 104/SC 4/WG 2 (ISO, 2025) addresses smart container identification and functionality, providing a regulatory and technical foundation for defining core capabilities. Complementing this, UN/CEFACT frameworks define structured event models and use cases, demonstrating that value emerges through interoperable event exchange rather than isolated device deployment.

3.4 Gaps

Three gaps remain. Limited distinction between connected assets and cyber-physical nodes, insufficient differentiation across functional maturity, and under-theorisation of standards as constitutive components of system value. Together, these gaps indicate the need for an integrated analytical framework that links cyber-physical systems logic, differentiated functional maturity, and standards-based semantics to explain how smart containers generate system-level value.

4. Methodology

This study adopts a conceptual and integrative research design combining engineering studies, standards documentation, and industry materials. The analytical framework evaluates smart container capabilities across three dimensions: sensing reliability, communication continuity, and event standardisation. In this study, functional maturity refers to the degree of convergence between technical reliability, communication continuity, semantic interoperability, standards alignment, and operational adoption.

Due to confidentiality constraints, the analysis relies on publicly available evidence. The objective is to assess under which conditions smart containers function as cyber-physical infrastructure rather than isolated devices.

5. Functional Maturity

5.1 Track-and-Trace

Track-and-trace represents the most mature domain. Engineering studies demonstrate robust multimodal communication architectures, including hybrid terrestrial and satellite systems (Falcitelli et al., 2024; Ledesma and Lamo, 2025). While performance varies across environments, location tracking shows strong convergence between technical capability and operational deployment. Together, these studies demonstrate that location visibility is increasingly supported through architectures explicitly designed for multimodal logistics. This is reflected in commercial deployments, where key providers such as Hoopo Systems, Nexxiot and ORBCOMM rely primarily on cellular communication technologies, including 4G LTE, NB-IoT, and GSM fallback, to ensure global connectivity for dry containers (Hoopo, 2026; Nexxiot, 2026; Orbcomm, 2026). These architectures enable continuous data transmission across heterogeneous geographic environments. In parallel, short-range communication technologies such as BLE and LoRaWAN are increasingly used in controlled infrastructure environments, including vessels and terminals, particularly in reefer container operations where local network integration is required (Birkel & Hartmann, 2020). This hybrid communication

approach reflects the necessity of combining wide-area and local connectivity to maintain communication continuity across multimodal transport contexts.

5.2 Door Event Detection

Door event detection is supported by institutional standards such as ISO 17363:2013 (ISO, 2013) and ISO 18185-1 (ISO, 2007). Empirical studies demonstrate improvements in security and operational efficiency through electronic seals and RFID integration (Kadir et al., 2016).

The value of door events lies in contextualisation. A door opening becomes meaningful only when combined with location, time, and operational expectations. This illustrates the transformation from sensor output to actionable system event. In the presented example (Figure 1; Figure 2), the identification of the door opening in combination with the precise location of the event and a time-stamp provides operational evidence to authorised or unauthorised penetration into a container. The alert in near real-time on this event enables a direct action to prevent an illegal activity around the container and the associated cargo. Therefore, these observations demonstrate that door event detection transcends its role as a security mechanism and operates as a cyber-physical system function, in which standardised, context-rich events enable coordinated, real-time decision-making across distributed logistics networks.

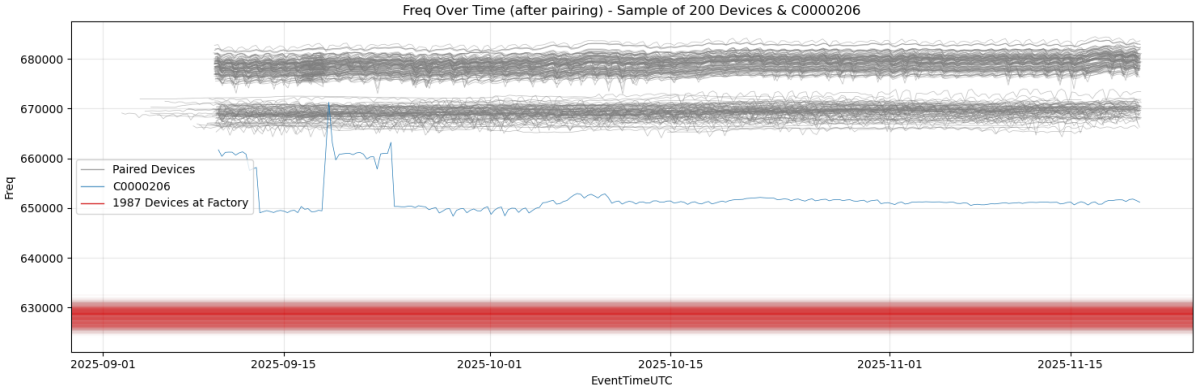


Figure 1: Dry Container Door Sensor Implementation. Hoopo©, provided by Ilan Lozovsky, PhD; Rashel Tublin, PhD

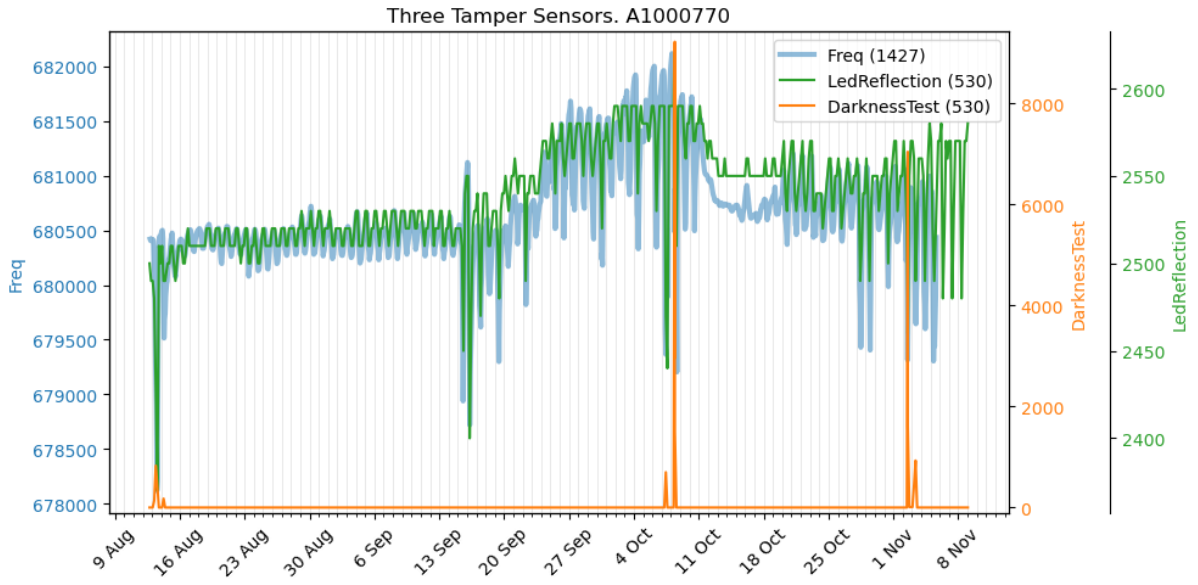


Figure 2: Dry Container Door Open/Close Identification. Hoopo©, provided by Ilan Lozovsky, PhD; Rashel Tublin, PhD

5.3 Fire-Related Monitoring

Fire-related anomaly monitoring remains less mature. Existing research indicates technical feasibility but limited standardisation and operational validation (Radicos, 2021). Detection requires complex interpretation of temperature patterns and environmental conditions (Konon, 2023). This domain highlights the limits of current smart container capabilities and reinforces the argument that system intelligence is asymmetrical across functions.

In the presented example (Figure 3 and Figure 4), temperature anomaly registration from an IoT-enabled container illustrates the variability and contextual dependence of thermal signals, requiring interpretation beyond simple threshold detection. Unlike location or door state, temperature data are inherently ambiguous and demand integration with environmental and operational context to produce reliable alerts.

This demonstrates that a cyber-physical system is not defined by the presence of sensors alone, but by the quality of coupling between physical state, signal reliability, semantic interpretation, and operational response. Consequently, fire-related monitoring exemplifies the limits of current system integration and illustrates that smart container intelligence must be understood as functionally asymmetrical rather than homogeneous.

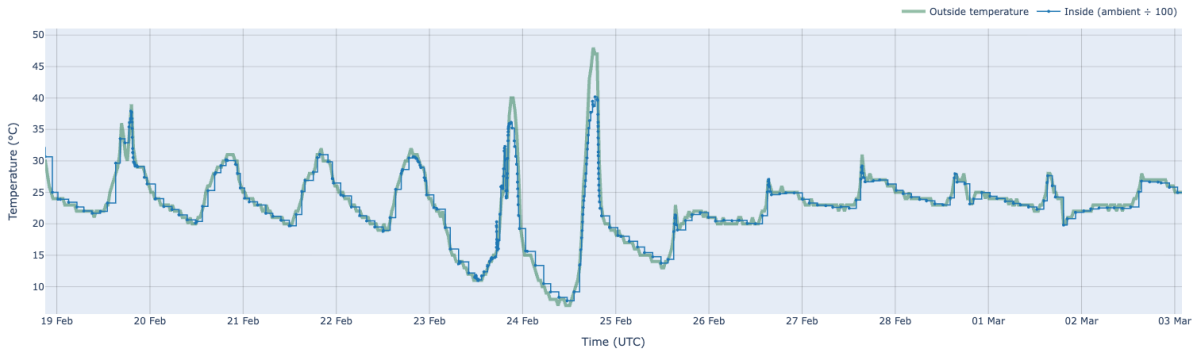


Figure 3: Indoor/Outdoor Temperature (°C) Monitoring in Miami, FL, Dry Container. Hoopo©, provided by Ilan Lozovsky, PhD; Rashel Tublin, PhD

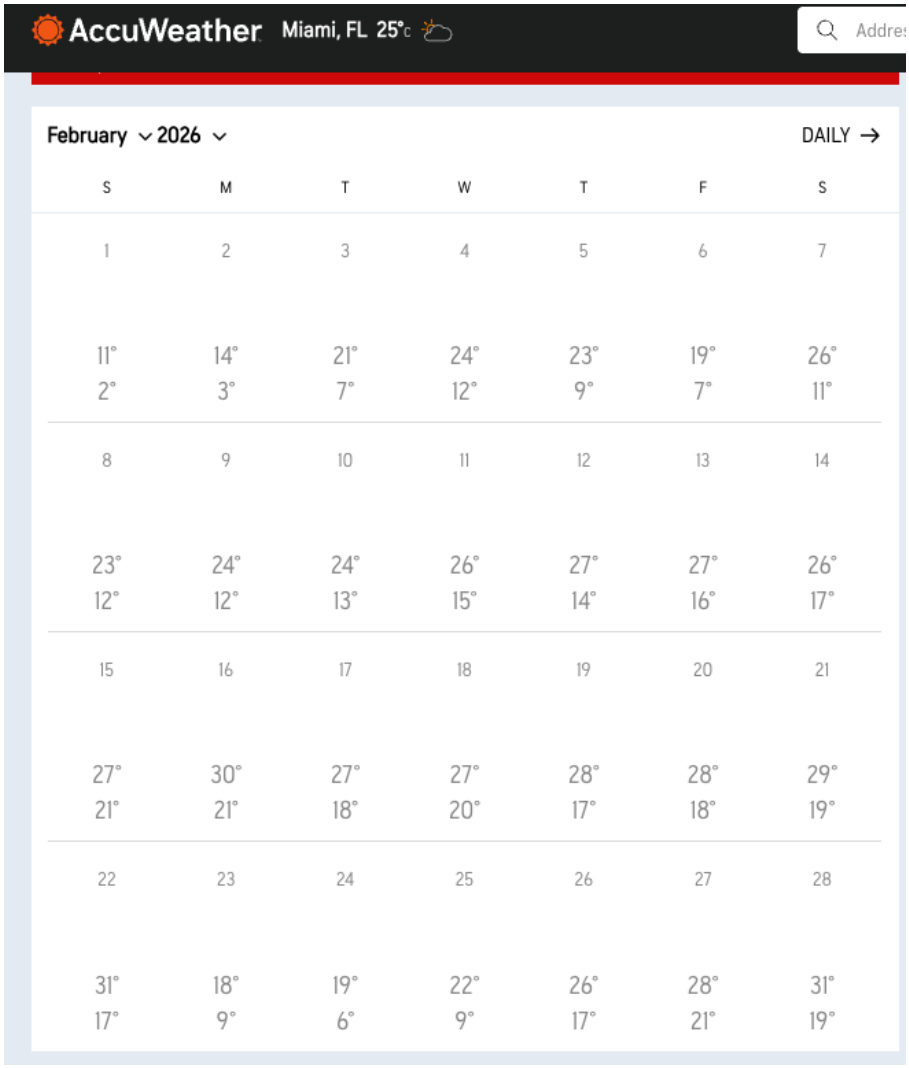


Figure 4: Historical temperature (°C) data in Miami, FL. February 2026.

6. Standards and Semantics

Standards convert sensor outputs into interoperable system events and therefore constitute a foundational component of smart container systems. Semantic interoperability requires that events are not only transmitted but consistently classified across heterogeneous systems. This includes shared definitions of event type, temporal structure, asset identity, geospatial reference, and event status. A door-open event, for example, acquires operational meaning only when multiple systems interpret the event identically in relation to time, location, and authorised operational expectations. At the level of formal international standardisation, ISO TC 104/SC 4/WG 2, including the ongoing development of ISO/AWI TS 25287 on smart container functionality and identification (ISO, 2025), defines core capabilities and establishes a structured link between physical asset characteristics and their digital representation. Building on this foundation, UN/CEFACT frameworks define structured data models and event semantics that enable cross-actor communication (e.g. UN ECE/UN/CEFACT, 2021). Without such shared semantics, sensor data remain confined to proprietary systems and cannot support system-level coordination.

Industry initiatives such as Digital Container Shipping Association (DCSA) further illustrate the operationalisation of these principles, translating standardised event definitions into

deployable track-and-trace frameworks. Event integrity therefore depends not only on data capture, but on consistent encoding, reliable time-stamping, and institutional trust across systems. The standardisation efforts across the industry confirm that standards are not auxiliary enablers, but constitutive elements that determine whether sensor outputs can function as interoperable and actionable system events within distributed logistics networks. Consequently, the value of smart container does not arise from isolated telemetry generation, but from the integration of trusted event streams into distributed logistics decision-making processes.

7. Empirical Illustration

Commercial smart container illustrates convergence towards integrated architectures combining tracking, condition monitoring, and event generation. Smart container reflects a broader transition from data collection to integration into the logistics operational logic.

For instance, the HoopoSense Solar represents a combined solution for track-and-trace and condition monitoring embedded within a single IoT device (Hoopo, 2026). Such integration enables a layered use-case approach, where condition-related events, including door openings or temperature anomalies, are contextualised through location and time data. This coupling transforms discrete sensor outputs into operationally meaningful events that can be interpreted across multiple supply chain actors.

The analytical relevance of this example lies not in vendor-specific performance, but in demonstrating architectural convergence across the industry. The integration of sensing, communication, and event generation reflects the emergence of smart containers as cyber-physical nodes, capable of supporting near real-time visibility, exception handling, and coordinated decision-making. Therefore, it could be suggested that the value of smart container is realised through their ability to embed physical assets into interoperable and event-driven logistics processes, rather than through isolated device functionality.

8. Governance

Smart container systems operate within sensitive data environments. Event streams may reveal operational patterns, requiring governance mechanisms to ensure controlled access.

Data sovereignty frameworks emphasise control, access, and accountability (Hummel et al., 2021). However, increased interoperability creates tension between transparency and competitive sensitivity. While data sharing enables coordination, it may also expose vulnerabilities.

Trusted ecosystems therefore require role-based access, data minimisation, and contractual agreements to balance sharing and protection. Therefore, this suggests that governance is not an external constraint but a constitutive condition of smart container systems, determining whether interoperable event data can be shared, trusted, and operationalised across actors without compromising competitive integrity.

9. Implications

The article reframes smart containers as cyber-physical infrastructure and highlights functional asymmetry across capabilities. It could be assumed, that standards and governance are central to system value.

For practitioners, system evaluation must prioritise interoperability and integration rather than hardware features. For policymakers, coordinated development of standards and governance frameworks is necessary for scalable deployment. Without such alignment, smart container deployments risk remaining fragmented and unable to contribute meaningfully to the coordinated, interoperable logistics systems envisioned by the Physical Internet.

10. Conclusion

Smart dry containers function as cyber-physical infrastructure when they transform physical states into interoperable system events. Their value lies in enabling coordination within logistics networks.

The analysis illustrates uneven maturity across functions, reinforcing the need for system-level integration. Smart containers contribute to the Physical Internet when they operate as trusted and interoperable nodes within multimodal systems. This reframing shifts the focus from device-level innovation to system-level coherence, indicating that the future impact of smart containers will depend not on sensing capabilities alone, but on their integration into standardised, governable, and interoperable logistics architectures.

References

Becha, H., Schroeder, M., Voorspuij, J., Frazier, T., & Lind, M. (2021). Global data exchange standards: The basis for future smart container digital services. In M. Lind, M. Michaelides, R. Ward, & R. T. Watson (Eds.), Maritime informatics (pp. 293-312). Springer. https://doi.org/10.1007/978-3-030-50892-0_18

Birkel H.S., & Hartmann, E. (2020). Internet of Things- the future of managing supply chain risks. *Supply Chain Management: An International Journal, 25*(5), 535-548. https://doi.org/10.1108/SCM-09-2019-0356

Cortes-Murcia, D. L., Guerrero, W. J., & Montoya-Torres, J. R. (2022). Supply chain management, game changing technologies, and Physical Internet: A systematic meta-review of literature. *IEEE Access, 10,* 61721-61732. https://doi.org/10.1109/ACCESS.2022.3181154

Digital Container Shipping Association. (2025, April 17). DCSA's Track & Trace standard ready for use in the Port of Rotterdam. https://dcsa.org/newsroom/track-trace-ready-for-use-in-port-of-rotterdam

Dragulinescu, A. M. C., Constantin, M. G., Ionescu, B., Vochin, M., & Tamas, R. (2023). Maritime Internet of Things meets LoRaWAN: A container testbed, measurement campaign and dataset analysis. *IEEE 9th World Forum on Internet of Things*. https://doi.org/10.1109/WF-IoT58464.2023.10539500

Falcitelli, M., Misal Noto, S., & Pagano, P. (2024). Development of a multi radio device for dry container monitoring and tracking. *IoT, 5*(2), 187-211. https://doi.org/10.3390/iot5020010

Falcitelli, M., Noto, S., Pagano, P., Gharbaoui, M., Isca, A., Fresi, F., Mancina, A., Toffetti, M., Amatruda, A., Bendoni, N., Sarandrea, E., & Scalambro, P. (2024). Full scale assessment of the 5GT system for tracking and monitoring of multimodal dry containers. *IoT, 5*(4), 922-950. https://doi.org/10.3390/iot5040042

Ghobakhloo, M., Iranmanesh, M., Foroughi, B., Tseng, M. L., Nikbin, D., & Khanfar, A. A. A. (2023). Industry 4.0 digital transformation and opportunities for supply chain resilience: A

comprehensive review and a strategic roadmap. *Production Planning & Control*, *36,* 1-31. https://doi.org/10.1080/09537287.2023.2252376

Hoopo Systems Ltd. (2025). Maritime container tracking. https://www.hoopo.tech/maritime

Hoopo Systems Ltd. (2025). Technology. https://www.hoopo.tech/technology

Hoopo Systems Ltd. (2026). HoopoSense Solar. https://www.hoopo.tech/hooposense-solar

Huang, S. W., Lee, M. T., & Gong, D. C. (2012). Implementing a passive RFID e-seal system for transit container security: A case study of Kaohsiung Port. *Advances in Transportation Studies, 26*, 69-88.

Hummel, P., Braun, M., Dabrock, P., & Tretter, H. (2021). Data sovereignty: A review. *Big Data & Society, 8*(1), 1-17. https://doi.org/10.1177/2053951720982012

International Organization for Standardization. (2013). ISO 17363:2013 Supply chain applications of RFID: Freight containers. https://www.iso.org/standard/57596.html

International Organization for Standardization. (2015). ISO 18185-3:2015 Freight containers: Electronic seals. Part 3: Environmental characteristics. https://www.iso.org/standard/62268.html

International Organization for Standardization. (2025). *ISO/AWI TS 25287 Smart container functionality and identification.* ISO. https://www.iso.org/standard/90215.html

Kadir, E. A., Rosa, S. L., & Gunawan, H. (2016). Application of RFID technology and e-seal in container terminal process. *2016 4th International Conference on Information and Communication Technology*. https://doi.org/10.1109/ICoICT.2016.7571926

Kavuri, S., Moltchanov, D., Ometov, A., Andreev, S., & Koucheryavy, Y. (2020). Performance analysis of onshore NB-IoT for container tracking during near-the-shore vessel navigation. *IEEE Internet of Things Journal, 7*(5), 4288-4301. https://doi.org/10.1109/JIOT.2020.2964245

Konon, V. (2023). Image processing method for cargo container identification in a stack within the cargo temperature control and fire safety system on container ships. *Polish Maritime Research, 30*(2), 114-120. https://doi.org/10.2478/pomr-2023-0027

Lambrou, M., Watanabe, D., & Iida, J. (2019). Shipping digitalization management: Conceptualization, typology and antecedents. *Journal of Shipping and Trade, 4*(11), 1–17. https://doi.org/10.1186/s41072-019-0052-7

Ledesma, O., & Lamo, P. (2025). Hybrid IoT network for real time monitoring of maritime containers. *Computer Networks, 259*, Article 111627. https://doi.org/10.1016/j.comnet.2025.111627

Levinson, M. (2016). The box: How the shipping container made the world smaller and the world economy bigger (2nd ed.). Princeton University Press.

Lis, D., & Otto, B. (2020). Data governance in data ecosystems: Insights from organisations. AMCIS 2020 Proceedings. https://aisel.aisnet.org/amcis2020/strategic_uses_it/strategic_uses_it/12

McCormack, E., Jensen, M., & Hovde, A. (2010). Evaluating the use of electronic door seals on shipping containers. *International Journal of Applied Logistics, 1*(4), 13-29. https://doi.org/10.4018/jal.2010100102

Nexxiot. (2026). Edge Data Sheet. https://nexxiot.com/wp-content/uploads/2024/07/2024-Nexxiot-Edge-Data-Sheet.pdf

Orbcomm. (2026). CT 1010. https://www.orbcomm.com/smart-start/direct-customer/ct-1010

Radicos. (2021). What can be done to improve the fire detection in a container cargo under deck and on deck. *TIS GDV*. https://www.tis-gdv.de/tis_e/gof/radicos.pdf

Raza, Z., Woxenius, J., Vural, C. A., & Lind, M. (2023). Digital transformation of maritime logistics: Exploring trends in the liner shipping segment. *Computers in Industry, 145*, Article 103811. https://doi.org/10.1016/j.compind.2022.103811

Safetytech Accelerator. (2022). A feasibility study into using energy harvesting sensors to detect fires in cargo holds. https://safetytechaccelerator.org/case-studies/energy-harvesting-sensors-to-detect-fires-in-cargo-holds/

Scheider, S., Lauf, F., Möller, F., & Otto, B. (2023). A reference system architecture with data sovereignty for human centric data ecosystems. Business & Information *Systems Engineering, 65*, 577-595. https://doi.org/10.1007/s12599-023-00816-9

Song, Y., Yu, F. R., Zhou, L., Yang, X., & He, Z. (2021). Applications of the Internet of Things in smart logistics: A comprehensive survey. *IEEE Internet of Things Journal, 8*(6), 4250-4274. https://doi.org/10.1109/JIOT.2020.3034385

Tijan, E., Jovic, M., Aksentijevic, S., & Pucihar, A. (2021). Digital transformation in the maritime transport sector. *Technological Forecasting and Social Change, 170*, Article 120879. https://doi.org/10.1016/j.techfore.2021.120879

United Nations Centre for Trade Facilitation and Electronic Business. (2019). Smart container business requirements specification. *UNECE.* https://unece.org/fileadmin/DAM/cefact/cf_plenary/2019_plenary/ECE_TRADE_C_CEFACT_2019_010E.pdf

United Nations Economic Commission for Europe. (2019). Trade facilitation white paper on real-time smart container data for supply chain excellence. *UNECE.* https://unece.org/fileadmin/DAM/trade/Publications/ECE_TRADE_446E_SmartContainers.pdf

United Nations Centre for Trade Facilitation and Electronic Business. (2021). White paper on integrated track-and-trace for multimodal transportation. *UNECE*. https://unece.org/sites/default/files/2021-03/ECE_TRADE_C_CEFACT_2021_INF4E-Integrated-T-T.pdf

Voorspuij, J., & Becha, H. (2021). Digitalisation in maritime regional and global supply chain. In M. Lind, M. Michaelides, R. Ward, & R. T. Watson (Eds.), *Maritime informatics* (pp. 65–80). Springer. https://doi.org/10.1007/978-3-030-50892-0_5

Wang, Y., & Sarkis, J. (2021). Emerging digitalisation technologies in freight transport and logistics: Current trends and future directions. *Transportation Research Part E: Logistics and Transportation Review, 148*, Article 102291. https://doi.org/10.1016/j.tre.2021.102291

World Customs Organization. (2025). Study report on smart security devices. https://www.wcoomd.org/-/media/wco/public/global/pdf/topics/facilitation/instruments-and-tools/tools/safe-package/smart-security-devices-report_en.pdf

ZIM. (2025). ZIM's smart container. https://www.zim.com/landing-pages/hoopo

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Cyber-Physical Internet-Enabled Spatial-Temporal Traceability for Green Methanol Production and Logistics

Qi Liu[1] , Leidi Shen[1], Zhichen Wu[1], Zhiheng Zhao[1,2*], and George Q. Huang[1,2]
1. Department of Industrial and Systems Engineering, The Hong Kong Polytechnic University, Hong Kong, China
2. Research Institute for Advanced Manufacturing, The Hong Kong Polytechnic University, Hong Kong, China
Corresponding author: zhiheng.zhao@polyu.edu.hk

Abstract: Sustainable supply chains require auditable lifecycle traceability across both cross-organizational logistics and intra-node material conversion. However, existing traceability approaches typically focus on securing isolated event records, mirroring physical assets, or capturing sensor data, without jointly representing shipment execution and material-state lineage in continuous-process networks. To address this gap, this paper proposes a Cyber-Physical Internet (CPI)–token-integrated framework for end-to-end spatial-temporal traceability in sustainable supply chains. In the proposed framework, CPI-based cyber-physical synchronization captures planned-versus-observed logistics trajectories across PI-aligned nodes, while tokenization records material transformation, splitting, merging, and ownership transfer through explicit state-transition and lineage rules. To operationalize the framework, the study further defines four lightweight validation indicators: traceability completeness (TC), spatial-temporal consistency (STC), mass-balance consistency (MBC), and a traceability feasibility index (TFI). A green methanol supply chain scenario is used to demonstrate the workflow, and a process-consistent simulated dataset is employed for proof-of-concept evaluation. The results show that the integrated CPI-token setting outperforms token-only and CPI-only baselines by providing more complete lineage continuity, stronger logistics consistency, and better compatibility between digital records and physical conversion yields. The study contributes a PI-aligned cyber-physical traceability architecture for sustainable continuous-process supply chains and offers an implementable basis for provenance verification and compliance auditing.

Keywords: *Cyber-Physical Internet, Green Methanol Supply Chain, Spatial-Temporal Traceability*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan:* ☐ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☐ *Transportation Equipment,* ☒ *PI Networks,* ☒ *System of Logistics Networks,* ☐ *Vertical Supply Consolidation,* ☐ *Horizontal Supply Chain Alignment,* ☒ *Logistics/Commercial Data Platform,* ☐ *Access and Adoption,* ☐ *Governance.*

Targeted Delivery Modes*:* ☒ *Paper,* ☐ *Poster,* ☐ *Flash Video,* ☐ *In-Person presentation*

1 Introduction

Advanced digitalization, sustainability, and resilience have emerged as the defining characteristics of modern manufacturing and logistics under the paradigms of Industry 4.0 and Industry 5.0 (Xu et al., 2021). Specifically, regarding sustainable supply chains such as those involving environmentally friendly fuels, their "green" premium and market value do not merely lie in the physical molecules themselves but rely fundamentally on rigorous lifecycle

traceability and sustainability certification. Because these products undergo complex material transformation and multi-tier logistics, any gap in the chain of custody can compromise their verifiable environmental attributes. Consequently, how to achieve end-to-end traceability and transparency has become a critical challenge. However, current sustainability certification across multi-stakeholder networks still relies heavily on manual documentation, isolated declarations, and periodic audits (Ares-Sainz et al., 2025). These conventional practices inevitably create organization-centric data silos and fragmented visibility, rendering the supply chain highly vulnerable to misreporting, data tampering, and "greenwashing".

To overcome the aforementioned traceability challenges in sustainable supply chains, the solution must fundamentally address two interconnected dimensions: the management of complex logistics networks and the tracking of continuous material conversions. Regarding the logistics dimension, the Cyber-Physical Internet (CPI) introduces a novel paradigm that aims to transport physical goods just like transmitting data packets over the internet (Ouyang et al., 2025). By standardizing routing and scheduling protocols, CPI provides a highly efficient architecture for physical distribution. However, CPI primarily focuses on the movement of discrete logistics entities (e.g., containers) and currently lacks the capability to trace the complex morphological changes and material conversions inherent in continuous process industries. Concurrently, regarding the material tracking dimension, tokenization has emerged as a powerful mechanism for creating programmable digital representations of physical assets, capable of enforcing strict mass balance rules during state transitions. Nevertheless, tokenization alone functions as a digital abstraction; without a standardized physical logistics framework to anchor these cyber-tokens to real-world shipments, it cannot achieve end-to-end traceability. Consequently, there is a critical need for a comprehensive framework that synergizes the physical routing capabilities of CPI with the state-transition tracking of tokenization to ensure authentic certification in continuous process industries.

To bridge these gaps, this paper proposes an integrated traceability framework that combines CPI-based logistics management and a tokenization method to achieve end-to-end traceability in sustainable supply chains. Within this framework, each component fulfills a distinct yet complementary role. On the one hand, CPI is utilized to achieve cyber-physical synchronization; by strictly aligning the physical movement of goods with their digital counterparts, it enables the seamless tracking of logistics processes across the supply chain. On the other hand, the tokenization method is designed to digitize material transformation processes. By mapping physical materials into cyber tokens and employing a "burn-and-mint" state transition logic, this method accurately records the step-by-step transformation of materials. This digitized conversion process lays a solid foundation, providing essential traceability evidence for accounting mechanisms, such as mass-balance rules. Ultimately, the proposed framework seamlessly links the spatial-temporal tracking of logistics provided by CPI with the discrete material-transformation records captured by tokens, ensuring that final sustainable products can be transparently traced back to their exact inputs of verified green feedstocks.

Methodologically, this study adopts a design-oriented proof-of-concept approach. It develops a CPI–token integrated traceability architecture, instantiates it in a green methanol supply chain, and evaluates it through simulation-based comparison across token-only, CPI-only, and integrated settings. Accordingly, this study addresses two research questions: 1) how CPI-based cyber-physical synchronization and tokenization can be integrated to enable end-to-end traceability in continuous-process sustainable supply chains; and 2) how the effectiveness of such an integrated framework can be operationalized and evaluated.

In summary, this study proposes an Industry 5.0-aligned integrated traceability framework by combining CPI-based cyber-physical synchronization with tokenization. The contributions of

this paper are threefold: 1) it develops a PI-aligned CPI–token framework that links logistics trajectories with discrete material-state transitions for end-to-end traceability; 2) it establishes a lightweight validation layer with four computable indicators (TC, STC, MBC, and TFI) to assess traceability completeness, spatial-temporal consistency, and feasibility; and 3) through a green methanol case study with simulation-based comparison, it shows that the integrated approach provides more robust and auditable traceability than single-mechanism baselines.

The remainder of this paper is organized as follows: Section 2 reviews related literature. Section 3 details the proposed traceability framework. Section 4 demonstrates its application through the green methanol case study, and Section 5 concludes the paper with future research directions.

2 Literature Review

With the transition towards Industry 5.0, the emphasis on sustainability has made the traceability of green fuels, such as green methanol, increasingly critical. However, the premium associated with sustainable products has led to significant "greenwashing" risks and persistent information asymmetry across supply chains (Appolloni and Chakraborty, 2025; Phan et al., 2026). Current certification practices heavily rely on fragmented documentation and centralized record-keeping, which are vulnerable to data manipulation and lack end-to-end dynamic verification (Alqaryuti et al., 2026; Gale et al., 2024; Kshetri, 2018). To address these challenges, recent studies have explored the integration of the Internet of Things (IoT) and blockchain technology to establish an immutable audit trail, ensuring data integrity from physical sensors to digital ledgers (Ibrahim et al., 2026; Saberi et al., 2018; Yessenbayev et al., 2024). Despite these advancements, most existing frameworks focus on discrete manufacturing or logistics. They still lack effective integration mechanisms for the complex continuous-flow logistics and material-state transformations inherent in process industries, which calls for a more granular and anti-tampering traceability paradigm.

To address dynamic uncertainties in complex supply chains, the Cyber-Physical Internet (CPI) constructs a cyber layer to optimize physical logistics entities (Huang, 2022; Ouyang et al., 2025). Inspired by the TCP/IP model, CPI encapsulates physical goods into Protocol Shipment Units (PSUs), enabling them to be dynamically routed across logistics nodes like data packets (Ouyang et al., 2025; Qu et al., 2024). This protocol-based framework significantly enhances network resilience and just-in-time (JIT) delivery during disruptions (Ouyang et al., 2025). However, current CPI research primarily focuses on the spatial-temporal routing of discrete entities, lacking fine-grained mechanisms to track complex asset state transformations within industrial supply chains.

On the other hand, tokenization provides a robust platform for digitizing physical assets, enabling the precise tracking of discrete transformation processes and ownership transfers across supply chains (Enayati, 2024; Westerkamp et al., 2020). By representing goods as programmable digital entities, this approach excels at recording compositional changes and discrete transaction events at specific nodes (Westerkamp et al., 2020). However, while tokenization effectively captures static snapshots of asset transitions, it inherently lacks the capability to natively trace continuous logistics events and dynamic spatial-temporal variations over time.

3 The Traceability Framework

3.1 CPI-Token Integrated Traceability Framework

This section presents the proposed CPI–token integrated framework for end-to-end traceability in sustainable supply chains. Unlike existing blockchain-based traceability systems that mainly secure isolated event records, digital-twin approaches that primarily mirror physical assets, and

IoT-enabled certification systems that focus primarily on data acquisition, the proposed framework is designed for continuous-process supply chains in which both inter-node logistics execution and intra-node material conversion must be jointly captured. As illustrated in Figure 1, the architecture consists of two tightly coupled realms: the physical space at the bottom and the cyber space at the top. The physical space covers the actual lifecycle and flow of materials, while the cyber space maintains synchronized digital representations of logistics states and material-state transitions. The key novelty of the framework lies in coupling CPI-based synchronization with token lineage management so that traceability is preserved not only across movements between nodes but also across transformations within nodes.

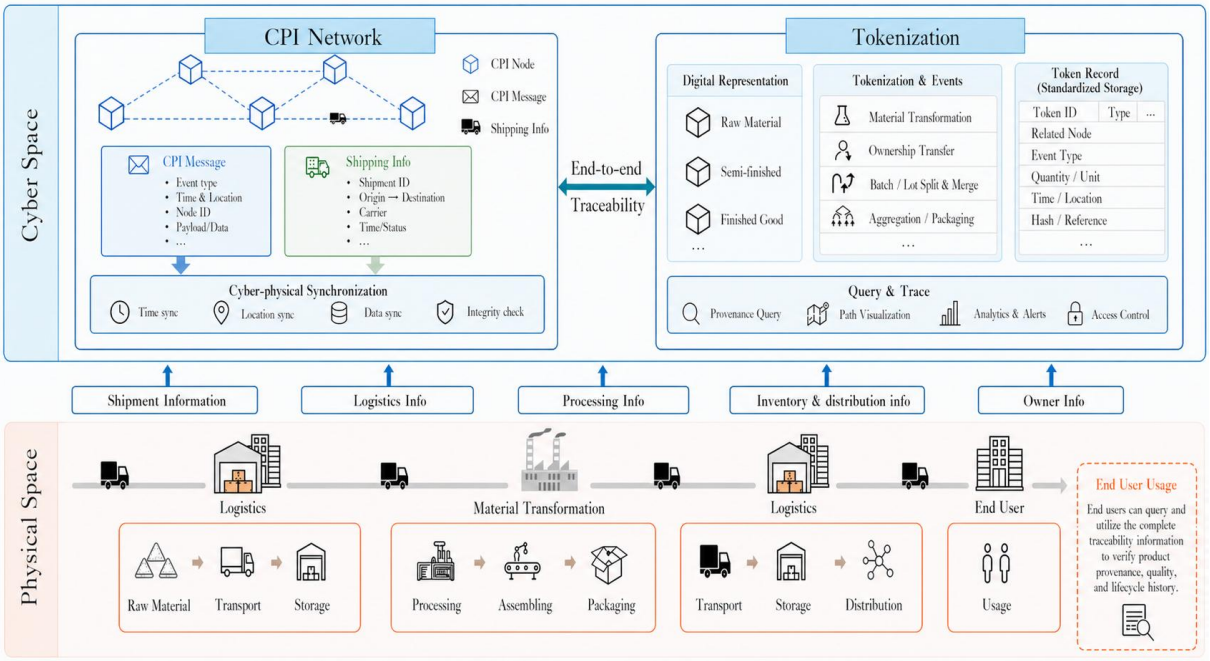


Figure 1: End-to-End Supply Chain Traceability Framework Based on Cyber-Physical Integration and Tokenization

Within the physical space, raw materials are transported to pre-processing and production facilities, transformed into intermediate and final products, and then distributed to end users. In PI terms, these facilities, storage points, and transfer locations can be interpreted as π-nodes, while shipments, lots, and transport units can be viewed as π-container-like traceable logistics units linked through standardized CPI event messages and shared identifiers. During these processes, heterogeneous data are generated, including shipment orders, routing plans, logistics events, processing records, inventory changes, and ownership transfers. Instead of remaining fragmented in organization-centric databases, these data are systematically extracted and uploaded into the cyber space to support interoperable end-to-end traceability across multiple stakeholders.

The first core mechanism in the cyber space is CPI-based cyber-physical synchronization, which is responsible for tracking the spatial-temporal execution of logistics activities. Shipping information, such as shipment ID, origin, destination, carrier, and planned schedule, is first registered to establish an expected logistics baseline. As physical goods move through the network, distributed CPI nodes generate standardized CPI messages containing event type, node ID, timestamp, location, and payload-related data. The synchronization mechanism then reconciles planned shipment information with actual execution records, thereby updating the digital trace according to real transport, handover, delay, return, or delivery events. Compared with conventional traceability systems that record only coarse logistics milestones, this

mechanism provides a PI-compatible and network-level view of how goods actually move across the supply chain.

While the CPI network monitors physical movements across logistics checkpoints, the tokenization method is employed to record and manage crucial discrete events within the sustainable supply chain. Each physical batch or material state is mapped into a digital token categorized as raw material, semi-finished product, or finished product. When operational events occur within a node, including transformation, splitting, merging, aggregation, packaging, and ownership transfer, the corresponding tokens are updated according to explicit state-transition rules. In particular, material conversion is represented through a burn-and-mint logic, whereby predecessor tokens are consumed and successor tokens are created while their relational references are preserved in the token records. Each token record is linked to the corresponding logistics execution records through shared batch/shipment identifiers, node identifiers, and event timestamps, enabling bidirectional traceability between transport events and material-state transitions. This design differs from generic digital certification registries because it does not merely attest that an event has occurred; it explicitly records lineage, quantity changes, and predecessor–successor relations, thereby enabling auditable mass-balance reasoning across conversion steps.

Ultimately, the originality of the proposed framework lies in the explicit coupling of CPI-based synchronization and token lineage management rather than in the isolated use of either mechanism. CPI-based synchronization anchors the digital trace to the actual spatial-temporal execution of logistics activities across PI-aligned nodes, whereas tokenization preserves material genealogy, quantity evolution, and ownership transitions during intra-node processing. By linking logistics execution records and token lineage through shared identifiers and cross-references, the framework unifies where and when a shipment moved with what material state it became in a single auditable traceability chain. This integrated design enables provenance queries, path reconstruction, compliance verification, and mass-balance auditing from raw-material origin to final product delivery.

3.2 Validation Metrics for End-to-End Traceability

To operationalize the two core mechanisms introduced in Section 3.1, this study defines a lightweight validation metric set for end-to-end traceability assessment. Specifically, the metrics are designed to evaluate (i) lineage completeness from origin to destination, (ii) spatial-temporal consistency between planned logistics baselines and observed CPI event trajectories, and (iii) mass-balance validity of tokenized material conversions at process nodes. Based on these dimensions, four parsimonious indicators are defined: TC, STC, MBC, and TFI.

- **Traceability Completeness (TC)**
  The TC metric measures whether each final batch can be fully traced back to verified upstream origins through continuous predecessor-successor links. Here, we define the TC as:

$$TC = \frac{1}{|B|}\sum_{b\in B} I(b)$$

$$I(b) = \begin{cases} 1, if\ batch\ b\ has\ an\ unbroken\ traceability\ chain \\ 0, otherwise \end{cases}$$

  where B is the set of final batches and $I(\cdot)$ is the indicator function.
- **Spatial-Temporal Consistency (STC)**
  The STC evaluates the alignment between planned routes/schedules and observed CPI event trajectories.

$$STC = 1 - \frac{1}{|B|}\sum_{b\in B} \alpha \frac{d(R_b^{obs}, R_b^{plan})}{d_{max}} + (1-\alpha)\frac{|T_b^{obs} - T_b^{plan}|}{T_{max}}$$

where R_b^{obs}, R_b^{plan} are observed/planned routes, T_b^{obs}, T_b^{plan} are observed/planned lead times, $d(\cdot,\cdot)$ denotes a normalized route-deviation function between the observed and planned node sequences, and $\alpha \in [0,1]$ balances spatial and temporal deviations.

- **Mass-Balance Consistency (MBC)**
  The MBC checks whether token burn-and-mint transitions remain consistent with expected physical conversion yields at process nodes.

$$\varepsilon_n = \frac{|\sum q_n^{in}\eta_n - \sum q_n^{out}|}{\sum q_n^{in}\eta_n}, MBC = 1 - \frac{1}{N}\sum_{n\epsilon N}\varepsilon_n$$

  where q_n^{in}, q_n^{out}, and η_n denote the input quantity, output quantity, and expected yield factor at node n, respectively.
- **Traceability Feasibility Index (TFI)**
  Finally, the TFI provides a compact overall score for cross-setting comparison and managerial interpretation.

$$TFI = w_1 TC + w_2 STC + w_3 MBC$$

  where w_1, w_2, w_3 are predefined weights reflecting governance priorities. In this study, we set $w_1 = 0.4, w_2 = 0.3, w_3 = 0.3$, giving a slightly higher priority to end-to-end chain completeness because it is the primary requirement of traceability-oriented certification.

4 Case Study: Green Methanol Supply Chain Demonstration

4.1 Green Methanol Scenario and CPI-token Workflow

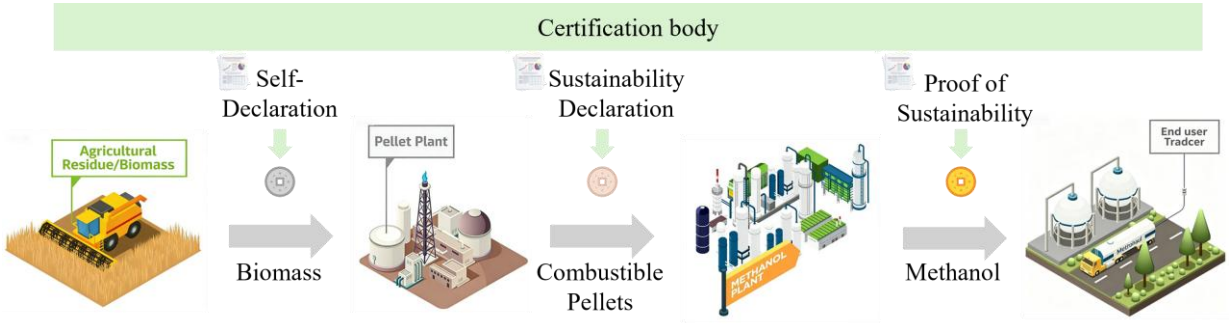


Figure 2: The physical architecture of the green methanol supply chain.

To demonstrate the applicability of the proposed framework, this study uses a green methanol supply chain as a representative case. This scenario is suitable because it combines two traceability challenges: frequent inter-organizational logistics handovers and intensive intra-node material conversions. This section first outlines the overall green methanol supply chain and its compliance requirements and then zooms in on the pellet production stage to illustrate how the dual-engine framework operates at the local level.

The physical chain includes four key stages under certification oversight: biomass collection, pellet processing, methanol synthesis, and final delivery to end users, as shown in Figure 2. The material flow begins with biomass collectors, who gather fragmented biomass waste and register the initial batch information. Upon transfer to a downstream pellet plant, the raw material undergoes inspection, weighing, and physical densification into standardized combustible pellets. These pellets are subsequently transported to a green methanol plant, where they are chemically synthesized into liquid methanol. Finally, the resulting methanol

product is distributed to end users. To ensure compliance across these successive material transformations, every transfer of physical custody must be accompanied by corresponding certifications issued under the purview of the certification body. The compliance chain progresses from initial Self-Declarations at the collection point to Sustainability Declarations (SD) after pelletization and ultimately to a Proof of Sustainability (PoS) for the final green methanol product. At the case level, the two engines serve complementary roles: the CPI network synchronizes verifiable logistics events across organizational interfaces, whereas tokenization represents material-linked asset states and preserves the lineage of material transformation across processing stages.

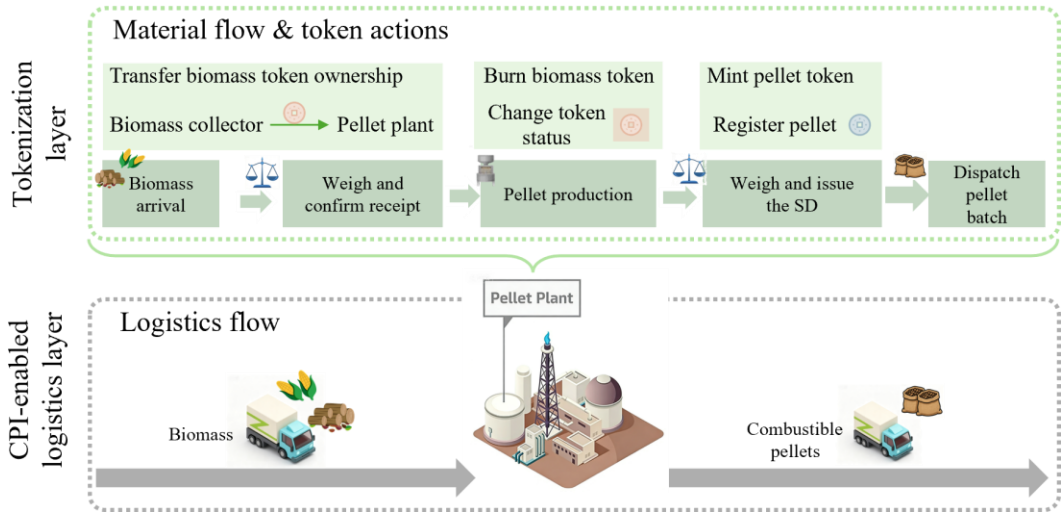


Figure 3: Cyber-physical integration and tokenization workflow for the pellet production stage.

Among these stages, the pellet production stage is selected for a detailed illustration because it contains both cross-organizational logistics handover and in-plant material transformation, making it a representative local unit for demonstrating the dual-engine framework. Figure 3 provides a zoomed-in view of this stage. The zoom-in covers a complete local chain: biomass is received and converted into pellets in the pellet plant, and the resulting pellet batches are then transported to the methanol plant. More specifically, this local chain is governed by two tightly coupled but functionally distinct layers. The CPI network synchronizes logistics events across organizational boundaries, while the tokenization layer represents ownership changes, asset-state transitions, and the transformation lineage associated with each batch.

The process follows an inbound logistics–material transformation–outbound logistics sequence. First, on the inbound side, the CPI network records and synchronizes the delivery and arrival of biomass at the pellet plant as verifiable inter-node logistics events, while the tokenization layer records the transfer of biomass token ownership from the biomass collector to the pellet plant together with the weighing and receipt confirmation of the incoming batch. Second, within the plant, the material flow takes place as the received biomass is processed into pellets; correspondingly, the tokenization layer captures this conversion by burning the biomass token and updating its status to indicate material consumption. Third, on the outbound side, tokenization records the production outcome by minting a pellet token for the newly generated pellet batch and registering the corresponding asset information. The pellet batch is then dispatched for downstream transport to the methanol plant, and the corresponding outbound logistics events are synchronized and recorded through the CPI network. In this way, the CPI network provides verifiable synchronization of logistics events across plant interfaces, while tokenization maintains the state continuity, ownership traceability, and transformation lineage

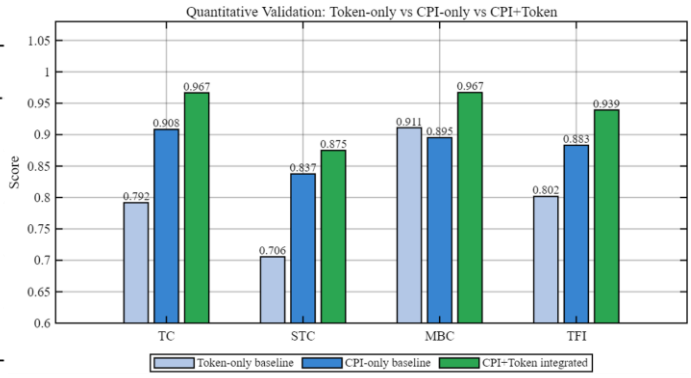
of the material-linked assets, together establishing a complete and auditable local traceability chain.

4.2 Quantitative Validation and Discussion

To provide a proof-of-concept evaluation of the proposed framework, this study constructs a process-consistent simulated dataset based on the green methanol workflow described in Section 4.1. Each simulated batch is represented as a sequence of operational events with fields such as batch identifier, plant stage, event type, timestamp, location, and output quantity, together with token predecessor–successor relations where applicable. To better approximate realistic supply-chain conditions, the simulation further injects stochastic anomalies, including random broken lineage chains, route deviations, lead-time deviations, and local mass-balance inconsistencies caused by record loss, synchronization failures, or abnormal token burn-and-mint operations. Figure 4(a) provides a compact illustration of five representative event records.

Batch	Stage	Event	Time_h	Location	Qty_out
B001	PelletPlant	Receive Biomass	2	Collector	120
B002	PelletPlant	Convert To Pellet	6	PelletLineA	115
B003	Transport	Dispatch	10	PelletGate	115
B004	Transport	Arrival	14	MethanolGate	110
B005	Methanol Plant	Burn Pellet Token	18	Reactor	0

(a)



(b)

Figure 4: Quantitative validation using simulated green methanol workflow data. (a) Compact illustrative sample of the simulated event schema. (b) Comparison of traceability metrics (TC, STC, MBC, and TFI) across three settings: the token-only baseline, the CPI-only baseline, and the CPI-token integrated setting.

Following Section 3.2, we generate 150 simulated batches and evaluate four indicators, namely TC, STC, MBC, and TFI. For each batch, planned route, transport lead time, and conversion-related parameters are first initialized, and then three comparison settings are constructed: (i) a token-only baseline, where discrete token state transitions are recorded without CPI-based logistics synchronization; (ii) a CPI-only baseline, where transport and handover evidence are captured through CPI events but material lineage and conversion consistency are represented less explicitly; and (iii) a CPI-token integrated setting, where token transitions are systematically linked with CPI checkpoint evidence. In addition, a small number of anomaly cases are injected into all settings to test robustness under operational disturbances.

Figure 4(b) compares the metric values across the three settings. The Token-only baseline shows relatively acceptable mass-balance consistency because transformation events are digitally recorded, but its TC and STC are lower due to the lack of logistics synchronization and end-to-end route evidence. The CPI-only baseline improves spatial-temporal consistency by better aligning planned and observed transport processes, but it remains weaker in lineage continuity and process-level conversion control. By contrast, the CPI-token integrated setting achieves the best performance across all four indicators, leading to the highest overall TFI. This result indicates that integrating CPI-based transport synchronization with token-based transformation records can provide more complete origin-to-destination traceability, better

planned-versus-observed execution consistency, and stronger compatibility between digital records and physical conversion yields.

Overall, the results suggest that CPI and tokenization are complementary rather than interchangeable. CPI mainly strengthens the observability of inter-organizational logistics and handover processes, whereas tokenization provides explicit digital lineage for material transformation inside process nodes. Using only one of the two mechanisms leaves blind spots in the traceability chain, while their integration offers a more robust end-to-end solution under random disruptions. From a managerial perspective, the proposed metric set also serves as a compact diagnostic tool: TC reveals whether custody chains remain complete, STC reflects execution deviations in logistics, MBC captures inconsistencies in physical-digital conversion, and TFI provides an interpretable overall score for cross-setting comparison.

5 Conclusion and Future Work

In conclusion, this paper proposes a CPI–token integrated framework for end-to-end spatial-temporal traceability in sustainable supply chains, using green methanol as the demonstration scenario. By combining CPI-based logistics synchronization with token-based lineage management, the framework jointly captures inter-node logistics execution and intra-node material conversion within a unified auditable traceability chain. The quantitative evaluation based on TC, STC, MBC, and TFI shows that the integrated setting outperforms token-only and CPI-only baselines, suggesting that CPI and tokenization are complementary rather than interchangeable.

This study should be regarded as a proof of concept, since the evaluation is based on a process-consistent simulated dataset rather than field deployment data. Although demonstrated using green methanol, the framework is potentially transferable to other continuous-process sustainable supply chains characterized by repeated custody handovers and material transformations, such as sustainable aviation fuel, hydrogen derivatives, and bio-based chemicals. Future work will focus on pilot implementation in real multi-stakeholder supply chains, refinement of the validation indicators under different certification settings, and integration of blockchain or smart-contract mechanisms to enhance tamper resistance and compliance assurance.

References

- Alqaryuti A., H. Aljaghoub, K. Salah, A. Mayyas (2026): Blockchain-based traceability, certification, and maintenance of FCEV components. Green Technologies and Sustainability, v4, 100396.
- Appolloni A., D. Chakraborty (2025): Addressing the drivers and concerns towards Agri-Food Strategy: Prioritizing environmental management business models. Business Strategy and the Environment, v34, no6, 7161-7179.
- Ares-Sainz J.L., A. Arias, N. Matovic, L. Ladu, G. Feijoo, M.T. Moreira (2025): Key governance and sustainability indicators for certification systems: Bridging certification and policy frameworks in the bioeconomy. Sustainable Production and Consumption, v56, 156-181.
- Enayati M., P. Gudimetla, S. Arlikatti (2024): Blockchain technology as a tool to make supply chains more resilient and sustainable. Operations and Supply Chain Management, v17, no3, 165-178, https://doi.org/10.31387/oscm0580434.
- Gale F., D. Goodwin, H. Lovell, H. Murphy-Gregory, K. Beasy, M. Schoen (2024): Renewable hydrogen standards, certifications, and labels: A state-of-the-art review from a sustainability systems governance perspective. International Journal of Hydrogen Energy, v59, 654-667.
- Huang G.Q. (2022): Cyber-Physical Internet (CPI) for Cross-Border Logistics of Manufactured Products. TechTalk, Tam Wing Fan Innovation Wing, The University of Hong Kong, https://innowings.engg.hku.hk/cyber-physical-internet/
- Ibrahim H., P.U. Ahakonye, P.N. Ahakonye, W.J. Ryu, L.A.C. Ahakonye, J.M. Lee, D.-S. Kim (2026): Integrated Blockchain and Machine Learning Framework for Polystyrene Waste Traceability. 2026 IEEE International Conference on Consumer Electronics (ICCE), 1-6.
- Kshetri N. (2018): blockchain's roles in meeting key supply chain management objectives. International Journal of Information Management, v39, 80-89.
- Ouyang Z., Z. Yuan, Q. Chen, M. Li, Z. Zhao, G.Q. Huang (2025): Shipment Scheduling and Routing Protocols in Cyber-Physical Internet for Prefabricated Construction Modules Logistics. IEEE Transactions on Intelligent Transportation Systems, v26, no10, 16032-16046.
- Phan T.M., T.D.M. Le, L.T.H. Nguyen, H.X. Pham, N.-H. Duong (2026): Blockchain-enabled green traceability for sustainable organic food: Impacts on trust, value co-creation and willingness to pay a premium. Green Technologies and Sustainability, 100400.
- Qu X., M. Li, Z. Ouyang, C.-L. Ng, G.Q. Huang (2024): Routing protocols for B2B e-commerce logistics in cyber-physical internet (CPI). Computers & Industrial Engineering, v193, 110293.
- Saberi S., M. Kouhizadeh, J. Sarkis, L. Shen (2018): Blockchain technology and its relationships to sustainable supply chain management. International Journal of Production Research, v57, no7, 2117-2135.
- Westerkamp M., F. Victor, A. Küpper (2020): Tracing manufacturing processes using blockchain-based token compositions. Digital Communications and Networks, v6, 167-176.
- Xu X., Y. Lu, B. Vogel-Heuser, L. Wang (2021): Industry 4.0 and Industry 5.0—Inception, conception and perception. Journal of Manufacturing Systems, v61, 530-535.
- Yessenbayev O., D.C.D. Nguyen, T. Jeong, K.J. Kang, H.R. Kim, J. Ko, J.-Y. Park, M.-S. Roh, M. Comuzzi (2024): Combining blockchain and iot for safe and transparent nuclear waste management: A prototype implementation. Journal of Industrial Information Integration, v39, 100596.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE
MINES PARIS | PSL

Mapping the European Logistics Data Ecosystem for the Physical Internet: A Structural Analysis of Gaps and Interoperability

Yu LIU [1, 2] and Eric BALLOT [2]
[1] EPF Engineering School, 55 avenue du président Wilson, 94230 Cachan, France
[2] Mines Paris, PSL University, Centre for management science (CGS), i3 UMR9217 CNRS, 75006 Paris, France
Corresponding author: yu.liu@minesparis.psl.eu, yu.liu@epf.fr

Abstract: *As data become a strategic asset for digitally coordinated supply chains, logistics increasingly depends on inter-organizational data sharing. In Europe, this development is supported by a growing set of regulations, governance mechanisms, infrastructures, standards, and implementation initiatives. However, the overall structure and coherence of this emerging ecosystem remain insufficiently understood. This paper analyzes the European Logistics Data Ecosystem (ELDE) through a structured mapping and knowledge-graph-based approach. The study identifies key elements across policy, governance, data space infrastructures, standards, projects, and implementation communities, and examines their interrelations within a multi-layered framework. The results show that ELDE is organized as a layered but weakly coupled system characterized by four structural gaps: (1) misalignment between regulatory framing and operational adoption, (2) coexistence of EU regulation and industry standardization, (3) tensions arising from federated architectural choices, and (4) a project-based pattern of operationalization. These gaps reveal that increasing ecosystem activity does not necessarily lead to effective integration across layers. The findings provide a structured diagnostic perspective on the European data ecosystem and highlight key challenges for the development of data-driven coordination in the Physical Internet.*

Keywords: *European data ecosystem, Logistics data sharing, data governance, data spaces, interoperability, data-driven value creation, Cyber Physical Internet.*

Physical Internet (PI) Roadmap Fitness*: ☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories), ☒ PI Networks, ☒ System of Logistics Networks, ☒ Horizontal Supply Chain Alignment, ☒ Logistics/Commercial Data Platform, ☒ Access and Adoption, ☒ Governance.*

Targeted Delivery Modes*: ☒ Paper, ☒ In-Person presentation.*

1 Introduction

With the rise of AI and data-driven business models, data has become a strategic asset rather than a by-product of digitalization. This shift has moved attention from firm-level data use toward the governance of data sharing across organizational boundaries (Bechtsis et al., 2022; Otto et al., 2026). In logistics and supply chains, where operations are inherently distributed across multiple actors, value creation increasingly depends on inter-organizational data exchange rather than firm-internal optimization alone (Gelhaar & Otto, 2020; Marty & Ruel, 2024). In Europe, this transformation is accompanied by an evolving data strategy that seeks to establish a regulated and federated data economy. Rather than focusing solely on increasing data availability, the European approach emphasizes trusted, interoperable, and sovereign data sharing across sectors and borders (European Commission, 2020; Otto, 2022). For logistics, this shift is particularly significant, as coordination depends on the exchange of heterogeneous and often sensitive data across actors, systems, and transport modes.

Despite substantial policy activity and a growing number of initiatives, the European logistics data landscape remains fragmented. Data sharing is constrained by heterogeneous standards, uneven operational uptake, and partial alignment between regulatory frameworks, infrastructures, and operational systems (Delaplace, 2024). As a result, the presence of policies, standards, and technologies does not automatically translate into effective data-driven value creation. This limitation is also reflected in the literature. While prior research has examined data sharing from technical, organizational, or firm-level perspectives, less attention has been given to how broader institutional, infrastructural, and operational layers interact in shaping ecosystem-level outcomes (Otto et al., 2026; Wu et al., 2025). This gap is particularly relevant for the Physical Internet (PI), and more specifically for its cyber-physical extension, the Cyber-Physical Internet (CPI), which emphasizes the integration of digital and physical layers through data-driven coordination and real-time information exchange.

To address this gap, this paper analyzes the European Logistics Data Ecosystem (ELDE) as a multi-layered system composed of policies, governance structures, infrastructures, standards, projects, and implementation initiatives. Using a structured mapping and knowledge-graph-based approach, the study examines how these elements are related and identifies structural gaps that shape the weak coupling between policy ambition and operational value realization.

The paper addresses the following research questions:

- RQ1: What are the key elements of the European data ecosystem relevant to freight transport and logistics?
- RQ2: How are these elements related across policy, governance, infrastructures, standards, projects, and implementation initiatives?
- RQ3: What structural gaps emerge from these relationships, and what do they imply for data-driven value creation and the future development of the Physical Internet?

The paper makes three contributions. First, it provides a structured representation of ELDE across multiple interacting layers. Second, it develops a relational and knowledge-graph-based analysis to identify systemic gaps in the ecosystem. Third, it shows how these gaps shape the conditions under which the Physical Internet can evolve from conceptual vision to operational reality.

2 Related works

2.1 Data sharing and value creation in logistics

Digitalization has transformed products, processes, and business models in supply chains, but data do not directly generate value; rather, they function as a resource within innovation and decision-making processes (Saleem et al., 2020). In logistics, value creation increasingly depends on inter-organizational data sharing rather than firm-internal optimization, as operations are distributed across multiple actors, modes, and systems. Seamless data exchange enables functions such as real-time rescheduling, ETA calculation, AI/IoT applications, and emissions monitoring, while also supporting spatiotemporal traceability, resilient decision-making, and synchronized coordination across multi-tier and multimodal networks (Delaplace, 2024; Wu et al., 2025).

This shift is closely aligned with the Cyber-Physical Internet (CPI), which extends the Physical Internet paradigm by integrating digital infrastructures, data sharing, and real-time coordination with physical logistics networks. Under such conditions, value creation is not generated within isolated firms, but through coordinated data exchange across interconnected actors and systems (Rachana Harish et al., 2025; Wu et al., 2025). As a result, data sharing becomes a structural

requirement for achieving system-level performance outcomes such as efficiency, resilience, and sustainability.

2.2 Data sharing as a socio-technical and governance challenge

Despite its importance, data sharing in supply chains remains difficult because it is not purely a technical problem, but a socio-technical coordination challenge shaped by the tension between sharing and control (Brewster et al., 2024; Gelhaar & Otto, 2020). Existing research shows that interoperability barriers, including heterogeneous data formats, system incompatibility, and semantic ambiguity, continue to constrain effective data exchange (Khan & Abonyi, 2022). At the same time, organizational and governance factors, such as trust, incentives, data sovereignty, and access rights, play a decisive role in shaping sharing behavior (Jussen, Möller, et al., 2024; Zheng et al., 2025).

These challenges are further amplified in logistics networks, where multiple stakeholders operate under different interests and constraints. Data sharing therefore requires not only technical solutions, but also governance mechanisms that define participation, access, and value distribution. Prior studies highlight that effective data-sharing systems must balance interoperability, trust, and incentives, while addressing issues such as data ownership, privacy, and coordination across actors (Rachana Harish et al., 2025; Wang et al., 2021).

2.3 European data ecosystem as a multi-layer institutional construct

In Europe, data sharing is not left to market dynamics, but is actively shaped through a coordinated institutional approach. The European Data Strategy aims to build a regulated and federated data economy in which data access, sharing, and reuse are governed across organizational and sectoral boundaries (European Commission, 2020; Gelhaar & Otto, 2020). This approach is based on principles such as data sovereignty, decentralized infrastructure, and public–private governance, implemented through a combination of regulations, governance frameworks, technical infrastructures, and sectoral initiatives (Nagel & Lycklama, 2022).

The European data ecosystem is inherently multi-layered, involving interactions between policy frameworks, governance bodies, technical architectures, standards, projects, and implementation initiatives. While these elements collectively support data sharing, their interactions remain only partially aligned, leading to fragmentation and coordination challenges across layers. Existing research has provided important insights into data sharing in supply chains, particularly regarding technical interoperability, governance mechanisms, and firm-level or dyadic collaboration. However, these studies typically address isolated aspects of the problem, either focusing on technological solutions or on specific organizational relationships. As a result, there is still limited understanding of how broader institutional, technical, and operational layers are structured and interact within an ecosystem-level perspective.

This limitation is particularly significant in the European context, where data sharing is being developed as a coordinated and policy-supported ecosystem rather than as a purely market-driven process. Yet, there is a lack of systematic analysis of how the European Logistics Data Ecosystem is organized across layers, how its elements are interrelated, and where structural gaps emerge. Addressing this gap, the present study analyzes ELDE through a multi-layered perspective and knowledge-graph-based approach.

3 Methodology

This study investigates the structure of the European Logistics Data Ecosystem (ELDE) through a combination of structured ecosystem mapping and knowledge-graph-based analysis. First, a structured mapping is conducted to identify the main elements relevant to data sharing in freight transport and logistics. These elements include policy and regulatory instruments, governance

bodies, data space infrastructures, standards, implementation projects, and operational initiatives. The selection of elements is based on their relevance to cross-organizational data exchange, interoperability, and logistics-specific applications within the European context. A multi-layer perspective is adopted, as prior research highlights that data ecosystems cannot be understood through isolated components, but require analysis across institutional, technical, and operational dimensions (Otto et al., 2026). Second, the mapped elements are represented through a knowledge-graph-based approach in order to analyze their interrelations. The objective is not to construct an exhaustive graph of all possible connections, but to identify structural patterns, dependencies, and discontinuities across layers. Relations are defined based on functional linkages such as regulatory influence, governance coordination, infrastructural enablement, standardization, and operational implementation. Third, qualitative content analysis is used to support the interpretation of these relationships. Selected initiatives, projects, and standards are examined to understand how data sharing is operationalized in practice and where mismatches occur between policy frameworks, infrastructures, and operational systems.

The result is a layered analytical representation of ELDE that enables the identification of structural gaps and the examination of their implications for data-driven value creation and the development of the Physical Internet.

4 Representative elements of the European Logistics Data Ecosystem

The European Logistics Data Ecosystem (ELDE) is composed of multiple layers that collectively shape data sharing in freight transport and logistics. Rather than constituting a single integrated system, these layers reflect different institutional logics and roles, ranging from regulatory frameworks to operational implementations. To support the subsequent structural analysis, this section identifies representative elements within each layer.

4.1 Policy and regulatory layer

At the policy level, the European Data Strategy provides the overarching vision for a single market for data, establishing the conditions for cross-sector and cross-border data sharing (European Commission, 2020). This strategic orientation is operationalized through key regulatory instruments such as the Data Governance Act and the Data Act, which define frameworks for trusted data sharing, access rights, and data use. In the logistics domain, the eFTI Regulation plays a more direct role by mandating the acceptance of electronic freight information, thereby creating a concrete regulatory driver for digital data exchange between economic operators and public authorities.

4.2 Governance and coordination layer

Governance and coordination are shaped by multi-stakeholder platforms and alliances that translate regulatory objectives into implementation frameworks. The Digital Transport and Logistics Forum (DTLF) is a central coordination body that develops interoperable and federated data-sharing concepts for logistics, including principles such as platform federation and trusted exchange. In parallel, initiatives such as the Data Spaces Business Alliance (DSBA) illustrate efforts to align key ecosystem actors and promote the adoption of data-space approaches across sectors.

4.3 Infrastructure and data space layer

The infrastructure layer is characterized by federated data-sharing architectures that enable secure and sovereign exchange across organizations. Key examples include the International Data Spaces (IDS) framework and GAIA-X, both of which provide mechanisms for interoperability, identity management, and usage control in decentralized environments. Implementation frameworks such as Eclipse Dataspace Components (EDC) translate these

concepts into deployable technical components. In the mobility domain, the European Mobility Data Space (EMDS) illustrates how these infrastructures are being applied to enable cross-sector and cross-border data exchange.

4.4 Project and implementation layer

Operationalization of data sharing is largely driven by EU-funded projects that test and validate architectures in real-world settings. Projects such as FENIX and FEDeRATED demonstrate the implementation of federated data exchange and semantic interoperability in logistics, while initiatives such as deployEMDS and eFTI4EU translate policy frameworks into operational environments. These projects are crucial to bridge the gap between conceptual architectures and practical deployment, although their impact often remains limited to pilot contexts.

4.5 Standards and domain-specific layer

At the operational level, interoperability is further shaped by industry-led standards that define data models and communication protocols within specific logistics domains. Examples include Digital container shipping Association (DCSA) in maritime transport, International Air Transport Association (IATA) ONE Record in air cargo, and GS1 standards for product and logistics identification. While these standards enable interoperability within individual domains, they remain fragmented across modes, highlighting the need to align cross-domain mechanisms.

4.6 Implementation communities and bridging initiatives

Between infrastructure and operations, implementation communities and industry-driven initiatives play a bridging role. Organizations such as the Open Logistics Foundation (OLF) translate regulatory frameworks and standards into deployable software components, while initiatives such as Shared Data Language (SDL) aim to address semantic interoperability across domains by introducing shared vocabularies. Similarly, projects such as iLEAP extend standardization into operational data exchange, particularly in areas such as emissions reporting.

These layers illustrate that ELDE is not a unified system, but a multi-layered and heterogeneous ecosystem in which different elements operate under distinct logics. Policies define conditions, governance bodies coordinate actors, infrastructures enable exchange, projects test solutions, standards structure operational data, and implementation communities translate concepts into practice. This layered structure provides the basis for analyzing how these elements interact and where structural gaps emerge.

5 Analysis of the overarching framework of data sharing in Europe

5.1 Policy stack and ecosystem structure

Figure 1 shows the European logistics data-sharing landscape. The logic of the stack is therefore sequential but also recursive. EU regulations create the legal and institutional conditions for data sharing; governance and coordination bodies translate these conditions into implementation frameworks; EU projects develop and test the necessary architectures and solutions; data space infrastructures provide the technical means for secure and sovereign exchange; and industry standards enable operational interoperability at the level of concrete business processes and applications. At the same time, feedback also moves upward: practical constraints encountered in projects, infrastructures, and industry standards can reveal gaps in governance design and regulatory assumptions. The stack should therefore not be seen as a one-way hierarchy, but as a multi-layered system of mutual shaping.

A key structural feature of this landscape is that its layers are not governed by the same institutional logic. The lower layers, namely policy, governance, and much of the implementation architecture, are largely shaped by the European digital governance ecosystem.

They are concerned with sovereignty, trust, legal compliance, and the creation of common enabling conditions for data sharing. By contrast, the upper layers, including industry standards and operational applications, belong primarily to the global logistics ecosystem, where interoperability is driven by sectoral requirements, commercial practices, and internationally coordinated standards. This creates a structural interface problem: the EU can shape the institutional and infrastructural conditions of data sharing, but it does not directly control the standards through which logistics data is modeled and exchanged in routine operations. As a result, the effectiveness of the European logistics data-sharing landscape depends not only on regulatory coherence within the EU, but also on the extent to which EU-led infrastructures can connect with globally established domain standards.

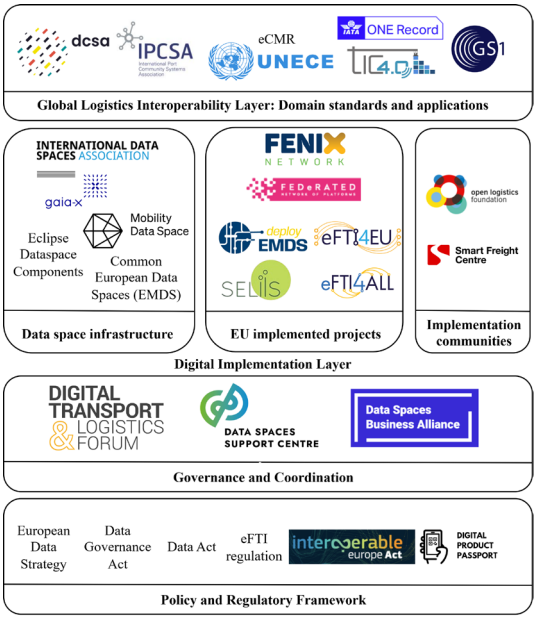


Figure 1. A layered framework of the European logistics data-sharing landscape.

5.2 Knowledge graph analysis of the European Data Ecosystem

Based on the Knowledge Graph (KG) analysis, Figure 2 does not attempt to represent all inter-element relations exhaustively. Instead, it restructures the complexity of the ecosystem into a five-layer analytical architecture and identifies four structural gaps that emerge across these layers. This abstraction serves two main purposes. First, it reduces the visual and analytical complexity of the knowledge graph by grouping heterogeneous elements according to their systemic roles. Second, it makes visible the key discontinuities that help explain why the ecosystem remains fragmented despite the increasing number of policies, governance arrangements, technical enablers, standards, projects, and communities. The detailed analysis and the insights derived from it are presented in the following subsection.

5.3 Knowledge graph analysis and insights

5.3.1 Regulatory framing vs operational adoption gap

The KG reveals a structural gap between regulatory framing and operational adoption. Rather than forming a linear translation chain, EU-level regulations are implemented through multiple indirect and fragmented pathways. While some instruments, such as the eFTI Regulation, connect relatively directly to implementation projects, most regulatory elements are translated through governance bodies, infrastructures, and pilot initiatives before reaching operational environments. This results in a weak coupling between data-space infrastructures and real operational systems. Initiatives such as EMDS, FENIX, and FEDeRATED enable secure and federated data exchange, but do not directly provide executable data models, workflow

integration, or incentives required for routine logistics operations. As a result, the transition from infrastructural capability to operational use remains uneven and path-dependent.

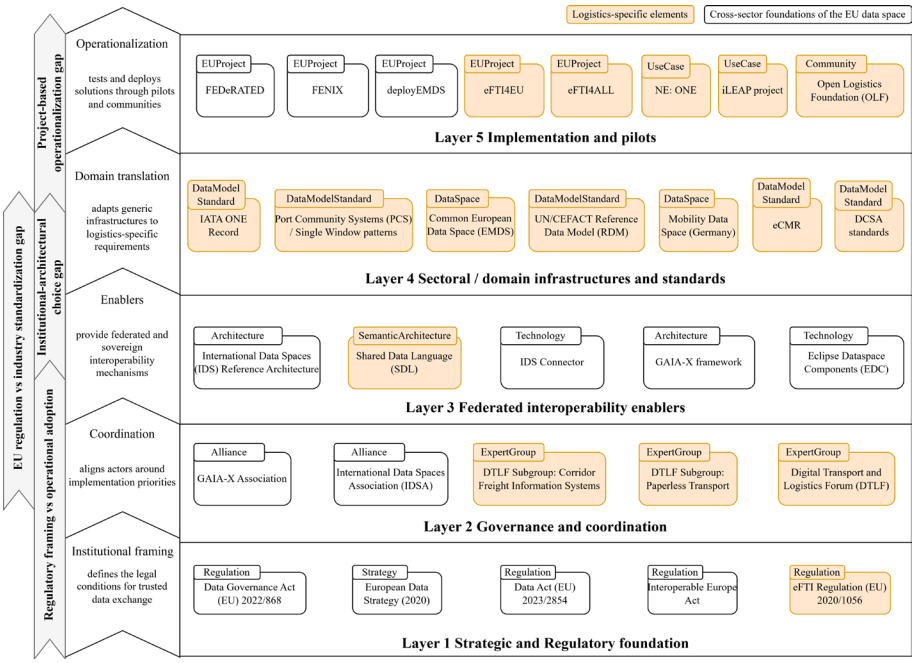


Figure 2. Layered architecture and structural gaps in the European Logistics Data Ecosystem (ELDE).

5.3.2 EU regulation vs industry standardization

The ecosystem is structured by a dual governance logic combining EU regulatory frameworks and industry-led standards. EU institutions define legal conditions for data sharing, while domain-specific standards such as DCSA, IATA ONE Record, and GS1 govern operational data exchange within logistics systems. These two logics differ in scope, objectives, and pace of development. Regulatory frameworks emphasize sovereignty, compliance, and cross-border coordination, whereas industry standards prioritize operational efficiency and domain-specific implementation. As a result, interoperability is not governed through a unified architecture, but through partially aligned systems that require continuous mediation. This creates structural frictions, including mismatches between compliance requirements and operational practices, as well as increased coordination costs across domains and actors.

5.3.3 Institutional-architectural choice gap

The ELDE is structured around a federated architecture that reflects a deliberate institutional choice emphasizing data sovereignty, decentralization, and regulated trust. Frameworks such as GAIA-X, IDS, and EDC embody this logic by enabling controlled and distributed data exchange rather than centralized platforms. While this design strengthens trust and compliance, it also increases system complexity and coordination requirements. Compared to centralized models, federated architectures demand higher levels of alignment across actors, infrastructures, and standards, which may slow adoption and increase entry barriers, particularly for smaller participants. In addition, the ecosystem exhibits an imbalance between regulatory push and market pull. Much of the development is driven by policy initiatives and public funding, while evidence of widespread bottom-up adoption remains limited.

5.3.4 Project-based operationalization gap

Implementation within ELDE is strongly driven by EU-funded projects, which act as primary vehicles for experimentation and validation. Initiatives such as FENIX, FEDeRATED, and deployEMDS create environments for testing interoperability and data-sharing concepts across actors and contexts. However, this results in a project-centric operationalization pattern. While projects are effective in generating pilot results, they provide limited evidence of long-term

continuity, institutional embedding, and large-scale adoption beyond their lifecycle. At the same time, operationalization remains fragmented across domains, with standards and implementations evolving within specific transport modes rather than converging toward integrated, cross-modal solutions. The current pattern remains predominantly project-centric and fragmented across domains. Without mechanisms to ensure continuity, standard consolidation, and cross-project integration, these initiatives risk remaining isolated rather than evolving into stable, reusable infrastructures.

6 Implications for the Physical Internet and its Cyber-Physical Extension

6.1 ELDE as a digital backbone for the Physical Internet

The Physical Internet (PI) envisions a highly interconnected, open, and collaborative logistics system. Its cyber-physical extension, the Cyber-Physical Internet (CPI), further emphasizes the integration of digital infrastructures, real-time data exchange, and intelligent coordination mechanisms that enable the operational realization of this vision. In this perspective, ELDE can be interpreted not only as a digital backbone for PI, but also as a foundational layer enabling the emergence of CPI. The layered structure of ELDE reflects key enabling dimensions of PI. Regulatory frameworks define the legal conditions for data access and exchange; governance bodies coordinate actors and establish shared rules; data-space infrastructures enable secure and sovereign data sharing; standards structure operational data; and implementation initiatives translate these elements into practical applications. Together, these layers create the institutional, technical, and operational environment required for data-driven coordination in logistics networks. However, the analysis also shows that ELDE is not yet a fully integrated system. Rather than functioning as a seamless backbone, it currently operates as a layered but weakly coupled architecture. As a result, while the necessary components for PI are emerging, their alignment remains incomplete.

6.2 Structural gaps as constraints on PI realization

The four structural gaps identified in this study represent not isolated issues, but systemic constraints that directly affect the feasibility of PI implementation. First, the gap between regulatory framing and operational adoption limits the integration of data-sharing capabilities into real logistics processes. While infrastructures for data exchange are increasingly available, they are not yet systematically embedded in operational systems where routing, scheduling, and coordination decisions are made. This restricts the ability to achieve real-time synchronization and adaptive network behavior, which are core requirements of PI. Second, the misalignment between EU regulatory frameworks and industry standards constrains end-to-end interoperability. PI depends on consistent data exchange across modes, networks, and actors, yet current interoperability remains fragmented along domain boundaries. Without stronger alignment between institutional frameworks and operational standards, cross-modal coordination and network integration remain limited. Third, the institutional–architectural choice of federated and sovereignty-oriented infrastructures introduces both enabling conditions and constraints. While these architectures support trust and control, they also increase complexity and coordination requirements. This may slow adoption and limit network effects, particularly in a system such as PI that relies on broad participation and scalable connectivity. Fourth, the project-based pattern of operationalization reflects an early-stage innovation dynamic, in which experimentation and validation are prioritized. While such projects are essential for developing and testing PI-relevant solutions, their current fragmentation and limited continuity constrain the transition toward stable, system-wide infrastructures. These gaps indicate that the challenge for PI is not the absence of technological or institutional elements, but the insufficient alignment between them.

6.3 From ELDE today to PI in the near future

The analysis suggests that ELDE is currently in a transitional phase between fragmented experimentation and more integrated ecosystem development. This phase is characterized by strong regulatory activity, growing infrastructure capabilities, and active pilot implementations, but limited system-wide integration and operational embedding. In the near future, the development of PI will depend less on the introduction of additional initiatives and more on the strengthening of alignment across existing layers. Three directions appear particularly critical. First, greater integration between data-space infrastructures and operational systems is required to enable the execution of data-driven coordination in real logistics environments. This includes the embedding of data-sharing mechanisms into transport management systems, platforms, and logistics hubs. Second, improved alignment between regulatory frameworks and industry standards to support consistent and scalable interoperability across modes and domains. This may involve stronger coordination between institutional actors and standardization bodies, as well as the development of bridging mechanisms such as shared semantic layers. Third, the transition from project-based experimentation to durable ecosystem capabilities must be supported. This includes ensuring continuity beyond pilot projects, consolidating successful solutions, and enabling their adoption at scale across logistics networks.

From a PI perspective, ELDE should therefore be understood not only as an enabling infrastructure, but also as a system whose current structure shapes the trajectory and speed of PI realization. The extent to which these structural gaps are addressed will determine whether PI remains a conceptual vision or evolves into an operational reality in the coming years.

7 Conclusion

This paper analyzed the European Logistics Data Ecosystem (ELDE) to better understand the conditions under which data-driven value creation can emerge in freight transport and logistics. By combining structured ecosystem mapping with a knowledge-graph-based approach, the study develops a relational perspective that captures how different elements interact at the system level. The main contribution lies in shifting the analytical focus from individual initiatives to the structural organization of the ecosystem. By identifying patterns of alignment and misalignment across layers, the paper provides a diagnostic framework for understanding why the increasing number of policies, infrastructures, and initiatives does not automatically translate into effective data-driven coordination. This perspective contributes to ongoing discussions on data ecosystems by emphasizing the importance of cross-layer coherence rather than isolated technological or regulatory solutions. For the development of the PI, or more specifically, the CPI, the analysis shows that current ELDE provide important enabling conditions for digitally coordinated logistics networks. However, the transition from fragmented data-sharing initiatives toward system-wide coordination depends on stronger alignment between regulatory frameworks, data-space infrastructures, domain standards, and operational systems. In this sense, ELDE shapes not only the current state of logistics data sharing, but also the trajectory through which PI can evolve toward cyber-physical realization.

This study has several limitations. As an exploratory and analytical approach, it abstracts from the full complexity of real-world interactions and does not assess adoption levels, performance outcomes, or stakeholder-specific dynamics. Future research could complement this work through empirical case studies, stakeholder-based analyses, or comparative studies across regions, with particular attention to how ecosystem-level alignment can be achieved and sustained in operational environments.

References

Ballot, E., Montreuil, B., & Meller, R. (2014). *The Physical Internet: The Network of Logistics Networks*. La Documentation française. https://books.google.fr/books?id=iX2ZAQAACAAJ

Bechtsis, D., Tsolakis, N., Iakovou, E., & Vlachos, D. (2022). Data-driven secure, resilient and sustainable supply chains: gaps, opportunities, and a new generalised data sharing and data monetisation framework. *International Journal of Production Research*, *60*(14), 4397–4417.

Brewster, C., Kalatzis, N., Nouwt, B., Kruiger, H., & Verhoosel, J. (2024). Data sharing in agricultural supply chains: Using semantics to enable sustainable food systems. *Semantic Web*, *15*(4),

Delaplace, J. (2024). *Digitalization of the logistics initiatives in Europe and recommendations for the introduction of the " Shared Data Language " developed by Geodis into this landscape* (Issue October).

European Commission. (2020). *A European strategy for data*. https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX%3A52020DC0066

Gelhaar, J., & Otto, B. (2020). Challenges in the emergence of data ecosystems. *Proceedings of the 24th Pacific Asia Conference on Information Systems: Information Systems (IS) for the Future, PACIS 2020*, *2019*, 1–14.

Jussen, I., Fassnacht, M., Schweihoff, J., & Möller, F. (2024). Reaching for the Stars: Exploring Value Constellations in Inter-Organizational Data Sharing. *European Conference on Information Systems*, *June*, 1–16.

Jussen, I., Möller, F., Schweihoff, J., Gieß, A., Giussani, G., & Otto, B. (2024). Issues in inter-organizational data sharing: Findings from practice and research challenges. *Data and Knowledge Engineering*, *150*(January). https://doi.org/10.1016/j.datak.2024.102280

Khan, A. A., & Abonyi, J. (2022). Information sharing in supply chains – Interoperability in an era of circular economy. *Cleaner Logistics and Supply Chain*, *5*(February), 100074.

Marty, J., & Ruel, S. (2024). Why is "supply chain collaboration" still a hot topic? A review of decades of research and a comprehensive framework proposal. *International Journal of Production Economics*, *273*(February). https://doi.org/10.1016/j.ijpe.2024.109259

Nagel, L., & Lycklama, D. (2022). How to build, run, and govern data spaces. In *Designing data spaces: The ecosystem approach to competitive advantage* (pp. 17–28). Springer.

Otto, B. (2022). The Evolution of Data Spaces. In B. Otto, M. ten Hompel, & S. Wrobel (Eds.), *Designing Data Spaces : The Ecosystem Approach to Competitive Advantage* (pp. 3–15). Springer International Publishing. https://doi.org/10.1007/978-3-030-93975-5_1

Otto, B., Jussen-Lengersdorf, I., Emons, S., Parker, G. G., & Chang, B. (2026). International Strategies Toward the Data Economy - A Comparative Analysis. *Software and Systems Engineering ISST*.

Rachana Harish, A., Liu, X., Li, M., Zhong, R. Y., & Huang, G. Q. (2025). The new supply chain information sharing Renaissance through crypto valuation mechanism of digital assets. *Transportation Research Part E: Logistics and Transportation Review*, *195*(July 2024), 103962.

Saleem, H., Li, Y., Ali, Z., Ayyoub, M., Wang, Y., & Mehreen, A. (2020). Big data use and its outcomes in supply chain context: the roles of information sharing and technological innovation. *Journal of Enterprise Information Management*, *34*(4), 1121–1143.

Wang, Z., Zheng, Z., Jiang, W., & Tang, S. (2021). Blockchain-enabled data sharing in supply chains: Model, operationalization, and tutorial. *Production and Operations Management*, *30*(7),

Wu, H., Li, M., Yu, C., Ouyang, Z., Lai, K. hung, Zhao, Z., Pan, S., Wang, S., Zhong, R. Y., Kuo, Y. H., Zhang, F., Huang, W., Shen, Z. J. M., Ballot, E., & Huang, G. Q. (2025). Towards cyber-physical internet: A systematic review, fundamental model and future perspectives. *Transportation Research Part E: Logistics and Transportation Review*, *197*(February), 104051.

Zheng, G., Ivanov, D., & Brintrup, A. (2025). An adaptive federated learning system for information sharing in supply chains. *International Journal of Production Research*, *63*(11), 3938–3960.

IPIC 2026 12th International Physical Internet Conference June 8-10, 2026 | Bordeaux, FRANCE

Place your logo here

Toward an Aviation-Enabled Physical Internet: Designing High-Speed Layers for African Multimodal Logistics Networks

Abiy Asrat Jiru
Addis Ababa University, Addis Ababa, Ethiopia
African Network for Innovation, Addis Ababa, Ethiopia
abiyasratj@gmail.com

Abstract: *The Physical Internet (PI) has emerged as a transformative logistics paradigm aimed at developing highly interconnected, modular, interoperable, and sustainability-oriented freight systems inspired by the architecture of the digital internet. Despite significant advances in PI research, existing literature remains predominantly focused on road-based freight systems, urban logistics, and rail-oriented multimodal transport structures, while aviation-centered logistics systems remain comparatively underexplored. This gap is particularly important in perishability-sensitive export environments in which air cargo functions not merely as a transportation mode, but as a critical coordination-intensive logistics layer.*

This paper proposes an airport-centric Physical Internet framework that conceptualizes aviation as a selective high-speed coordination layer within PI-enabled logistics systems. The framework introduces shipment-level coordination intelligence through an H-index intervention structure that dynamically evaluates shipment vulnerability, disruption exposure, perishability sensitivity, and coordination criticality in order to activate adaptive logistics interventions.

The study is empirically grounded using shipment-level operational data obtained from Ethiopian Airlines Cargo operations. A total of 5,165 valid export shipment observations were analyzed using operational timestamp structures derived from cargo acceptance and departure processes. Simulation-supported comparative analysis was implemented in R to evaluate operational performance under both conventional export coordination conditions and a modeled PI-enabled airport-centric coordination environment.

The findings indicate that the PI-enabled coordination environment reduced mean acceptance-to-departure time from 11.11 hours to 8.94 hours while improving modeled shipment reliability from 70.84% to 81.71%. Sensitivity analysis further demonstrated that the comparative performance of the PI-enabled coordination environment remained stable across moderate variation in intervention thresholds. The findings suggest that airport-centric coordination intelligence may significantly strengthen resilience, reduce spoilage exposure, and improve synchronization performance in perishability-sensitive export systems.

The paper contributes theoretically by extending PI research toward airport-centric logistics orchestration and shipment-level coordination intelligence beyond predominantly terrestrial PI architectures (Crainic et al., 2020; Treiblmaier et al., 2020). Methodologically, the paper demonstrates how operational air-cargo shipment data may support simulation-oriented PI coordination modeling. Practically, the framework provides strategic insights for airport operators, airlines, logistics planners, and policymakers seeking to improve resilience-oriented export logistics systems in emerging economies.

Keywords: *Airport-centric, Physical Internet, Multimodal Logistics, Emerging Economy, Mega Airport, Aviation-enabled Physical Internet, High-speed layer, Physical Internet, airport-centric logistics, air cargo, perishability-sensitive exports, coordination intelligence, resilience modeling, Ethiopian Airlines Cargo, shipment-level logistics, PI-enabled coordination.*

Physical Internet (PI) Roadmap Fitness*: Select the most relevant area(s) for your paper according to the PI roadmaps adopted in Europe and Japan: ☒ PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories), ☐ Transportation Equipment, ☒ PI Networks, ☐ System of Logistics Networks, ☐ Vertical Supply Consolidation, ☐ Horizontal Supply Chain Alignment, ☐ Logistics/Commercial Data Platform, ☐ Access and Adoption, ☐ Governance.*

Targeted Delivery Mode-s*: ☒ Paper, ☐ Poster, ☐ Flash Video, ☒ In-Person presentation*

1 Introduction

The growing complexity, vulnerability, and fragmentation of global logistics systems have intensified scholarly and industrial interest in resilient, interoperable, and digitally coordinated freight architectures. Within this context, the Physical Internet (PI) has emerged as a transformative logistics paradigm aimed at redesigning freight transportation systems through modularization, open interconnectivity, standardized logistics interfaces, and synchronized logistics coordination (Montreuil, 2011; Ballot et al., 2014). Inspired by the architecture of the digital internet, PI seeks to establish a globally interconnected network-of-networks in which freight flows move through intelligent, adaptive, and highly interoperable logistics structures.

Despite substantial advances in PI research, existing literature remains heavily concentrated on road freight systems, urban logistics, rail-oriented multimodal transport, and containerized terrestrial freight architectures (Treiblmaier et al., 2020; Pan et al., 2021). In comparison, aviation-centered logistics systems remain significantly underrepresented within PI scholarship. This limitation is increasingly important because air cargo plays a critical role in global high-value and perishability-sensitive supply chains.

The underrepresentation of aviation within PI research becomes particularly significant in emerging export economies characterized by perishability-sensitive logistics systems. In such contexts, aviation frequently functions not merely as a transportation mode, but as a strategic resilience layer that supports time-sensitive exports, export-value preservation, cold-chain continuity, and global market connectivity.

Ethiopia provides a particularly relevant context for examining this challenge. Ethiopian Airlines Cargo has evolved into one of Africa's largest air-cargo logistics systems, supporting substantial volumes of horticulture, perishables, pharmaceuticals, and time-sensitive exports. The operational structure of such export systems requires highly synchronized coordination between exporters, consolidation centers, cold-chain facilities, airport logistics operations, airlines, and international gateway networks.

Conventional export coordination systems often exhibit fragmented visibility, reactive disruption management, heterogeneous cargo prioritization, and operational coordination delays. Such conditions may increase lead-time variability, spoilage exposure, reliability deterioration, and export-value losses within perishability-sensitive logistics environments.

This paper argues that airport-centric coordination intelligence may provide a strategically important extension of PI-enabled logistics systems. Rather than conceptualizing aviation merely as a transportation mode, the paper proposes that airports can function as intelligent PI coordination nodes capable of orchestrating shipment-level synchronization, adaptive intervention, resilience-oriented coordination, and disruption-sensitive logistics management.

The paper introduces an airport-centric PI framework that operationalizes shipment-level coordination intelligence through an H-index intervention structure. The framework dynamically evaluates shipment vulnerability, perishability sensitivity, disruption exposure, and coordination criticality in order to activate adaptive coordination mechanisms within the airport logistics environment.

To empirically ground the framework, the study utilizes shipment-level operational data derived from Ethiopian Airlines Cargo operations. Comparative simulation-supported analysis was implemented in R to evaluate operational performance under both conventional export coordination conditions and a modeled PI-enabled airport-centric coordination environment.

The paper addresses the following research questions:

1. How can airport-centric coordination intelligence extend Physical Internet architectures for perishability-sensitive export systems?
2. How may shipment-level operational vulnerability structures support adaptive PI-enabled coordination mechanisms?
3. To what extent can simulation-supported airport-centric coordination improve operational resilience, reliability, and lead-time performance?

The paper contributes theoretically by extending PI research toward airport-centric logistics orchestration and shipment-level coordination intelligence. Methodologically, the paper demonstrates how operational shipment-level air-cargo data may support simulation-oriented PI coordination modeling. Practically, the framework provides insights for airport operators, airlines, logistics planners, and policymakers seeking to strengthen resilience-oriented export logistics systems..

2 Literature Review

2.1. The Physical Internet Paradigm

The Physical Internet (PI) was originally introduced as a transformative logistics vision aimed at redesigning freight transportation systems through open interconnectivity, modular logistics units, standardized interfaces, intelligent synchronization, and highly coordinated logistics networks (Montreuil, 2011; Ballot et al., 2014). PI seeks to overcome inefficiencies associated with fragmented transportation systems by enabling freight flows to move dynamically through interconnected logistics infrastructures.

Existing PI scholarship has made significant contributions in areas such as: modular freight systems, hyperconnected logistics networks, open logistics standards, collaborative logistics architectures, synchromodal transport systems, and sustainable logistics integration.

However, most PI research remains concentrated on: road freight systems, urban logistics, inland freight corridors, and rail-road multimodal architectures.

Comparatively limited attention has been directed toward airport-centric logistics systems despite the strategic role of aviation within time-sensitive global supply chains.

2.2. Aviation and Airport-Centric Logistics Systems

Air cargo functions as a strategic logistics infrastructure for high-value, time-sensitive, and perishability-sensitive supply chains (Merkert & Alexander, 2018).

Air cargo functions as a strategic logistics infrastructure for: perishability-sensitive exports, pharmaceuticals, high-value products, emergency logistics, and time-critical supply chains.
Unlike conventional freight systems, aviation logistics environments exhibit extremely high coordination intensity because shipment synchronization depends on: airport handling processes, cargo readiness timing, customs coordination, cold-chain continuity, aircraft scheduling, and disruption-sensitive operational sequencing.
Airport logistics systems therefore function not merely as transportation gateways, but as highly interconnected coordination ecosystems. Despite this, airport-centered coordination structures remain insufficiently integrated within mainstream PI scholarship.

2.3. Synchromodality, Resilience, and Coordination Intelligence

Recent logistics scholarship increasingly emphasizes adaptive coordination, resilience-oriented logistics systems, and digitally synchronized freight architectures (Christopher & Peck, 2004; Pfoser et al., 2021). Recent logistics scholarship increasingly emphasizes: adaptive coordination, resilience-oriented logistics systems, disruption management, digital visibility, and intelligent logistics synchronization.
Synchromodal logistics systems particularly emphasize dynamic modal coordination, operational flexibility, and adaptive routing structures (Ambra et al., 2019; Giusti et al., 2019). However, existing synchromodal literature still remains predominantly terrestrial in orientation.
The concept of coordination intelligence has also gained importance within resilience-oriented logistics systems. Coordination intelligence refers to the ability of logistics systems to dynamically identify operational vulnerability and activate adaptive coordination mechanisms in response to disruption-sensitive conditions (Ponomarov & Holcomb, 2009; Sheffi, 2005).
Within airport-centric logistics systems, coordination intelligence becomes particularly important because: perishability-sensitive cargo, time-critical export flows, and disruption-sensitive operational environments require differentiated coordination intervention.

2.4. Research Gap

The literature review identifies four principal research gaps.
First, aviation-centered logistics systems remain underrepresented within Physical Internet scholarship.
Second, existing PI frameworks rarely conceptualize airports as intelligent coordination nodes capable of orchestrating adaptive logistics intervention.
Third, shipment-level coordination intelligence mechanisms remain insufficiently developed within PI research.
Fourth, empirical operational shipment-level evidence remains limited within airport-centered PI research.
This paper addresses these gaps by proposing an airport-centric PI framework grounded in shipment-level operational analysis and simulation-supported coordination modeling.

3 Conceptual Framework

3.1. Research Gap

As illustrated in Figure 1 below, the Airport-Centric Physical Internet Coordination Framework conceptualizes airports as intelligent PI coordination nodes embedded within wider multimodal logistics architectures. Within this structure, airports function as: synchronization centers, resilience-oriented coordination hubs, adaptive intervention environments, and interoperability gateways. The framework consists of five principal layers: Export production layer; Logistics

and cold-chain layer; Airport-centric PI coordination layer; Airline and global gateway layer; and International distribution layer.

Across these layers operate cross-cutting enabling structures including: digital visibility, interoperability standards, disruption management, sustainability coordination, and institutional governance.

The below figure portrays the discussion above – as a representative conceptual framework.

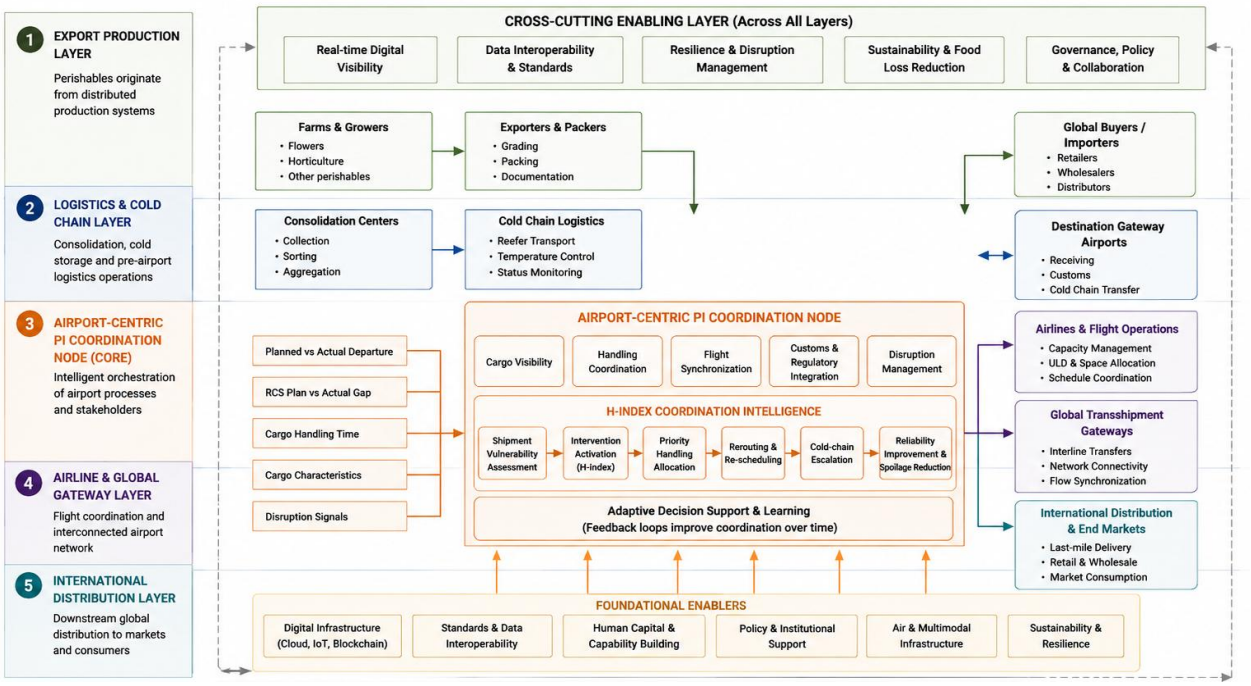


Figure 1. Airport-Centric Physical Internet Coordination Framework

3.2. H-index Coordination Intelligence

The framework introduces a shipment-level coordination intelligence structure through an H-index intervention mechanism.

The H-index operationalizes shipment vulnerability through: time sensitivity, perishability, reliability criticality, and disruption exposure.

The H-index is represented as:

$$H = 0.25(T) + 0.25(P) + 0.25(R) + 0.25(D) \quad (1)$$

Where:

T = time sensitivity,

P = perishability,

R = reliability criticality,

D = disruption exposure.

The equal-weight structure was analytically assigned to maintain balanced representation across operational vulnerability dimensions during the exploratory simulation stage.

When shipment coordination sensitivity exceeds the intervention threshold, the framework activates: priority handling, adaptive rerouting, synchronization escalation, and resilience-oriented coordination mechanisms.

3.3 PI-Enabled Coordination Logic

The framework proposes that airport-centric PI coordination systems may dynamically differentiate shipment handling according to: disruption sensitivity, perishability exposure, operational urgency, and coordination criticality.

The proposed coordination logic therefore shifts airport logistics systems from reactive coordination, towards adaptive coordination intelligence

4 Methodology

4.1 Research Design

The study adopted a simulation-supported operational logistics analysis approach grounded in shipment-level air-cargo operational data. The methodology combined:

- operational timestamp analysis,
- coordination-variable construction,
- shipment-level vulnerability modeling,
- and comparative PI-enabled simulation.

The study does not attempt to estimate causal operational implementation effects. Rather, the analysis evaluates analytically modeled coordination outcomes under a PI-enabled airport-centric coordination environment.

4.2 Data Source

The study utilized shipment-level operational data obtained from Ethiopian Airlines Cargo operations. The original dataset contained 5,756 export shipment observations.

Following data cleaning, timestamp validation, and operational filtering, 5,165 valid shipment observations were retained for analysis.

The dataset included: airway bill identifiers, origin and destination information, export flight identifiers, planned departure times, actual departure times, cargo readiness timestamps, and special handling codes.

4.3 Variable Construction

The study constructed operational coordination variables from shipment timestamp structures.

The core Operational Variables used are: Acceptance-to-departure time; Planned versus actual departure delay; RCS planning gap.

Whereas, the Coordination Intelligence Variables that the study derived are: disruption exposure, perishability sensitivity, reliability criticality, and the H-index coordination structure.

Furthermore, the PI-Enabled Scenario Variables that are modeled in the coordination environment are: adaptive intervention activation, priority handling, synchronization improvement, and reduced coordination delay.

4.4 Simulation Architecture

The simulation architecture was implemented in R using: tidyverse, ggplot2, ggplotly, dplyr, and shipment-level comparative analytical structures. The simulation evaluated two comparative environments: firstly, the conventional export coordination environment; and then the PI-enabled airport-centric coordination environment.

The same shipment structures were evaluated under both coordination environments.

4.5 Intervention Logic

The intervention threshold was analytically specified to distinguish ordinary coordination conditions from elevated disruption-sensitive shipment environments.

When the H-index exceeded the intervention threshold, the modeled PI-enabled coordination system activated: priority handling, synchronization escalation, adaptive rerouting, and resilience-oriented intervention.
The intervention threshold should therefore be interpreted as a structured analytical parameter within an exploratory simulation-supported coordination framework rather than a fixed operational constant. Similarly, the modeled PI-enabled coordination reduction factor represents an analytically specified coordination-improvement parameter designed to simulate the potential effects of enhanced synchronization, adaptive intervention, and airport-centric coordination intelligence within the exploratory PI-enabled environment. The simulation therefore evaluates directional comparative behavior rather than precise operational implementation outcomes.

4.6 Sensitivity Analysis

Sensitivity analysis was conducted by varying the intervention threshold between 0.60 and 0.80. The objective of the sensitivity analysis was to evaluate whether comparative PI-enabled coordination outcomes remained stable across moderate variation in intervention activation conditions.

5 Empirical Findings

5.1 Descriptive Statistics

Table 1 presents the principal descriptive statistics.

Table 1. Descriptive Statistics

Variable	Mean	SD	Median
Acceptance-to-departure time (hrs)	11.11	9.31	8.50
Departure delay (hrs)	2.29	15.75	0.42
RCS planning gap (hrs)	-8.48	15.54	-7.63
H-index	0.748	0.161	0.813

Shipment-level operational analysis revealed substantial variability within airport coordination performance. The acceptance-to-departure distribution exhibited significant right-skewness, indicating the presence of disruption-sensitive operational behavior affecting a subset of shipments.
The departure-delay distribution similarly demonstrated elevated operational variability, suggesting that severe disruption conditions affect a subset of export operations despite relatively moderate median delay levels.
The RCS planning gap results indicate that shipments are generally delivered substantially earlier than planned departure times, suggesting precautionary coordination behavior within perishability-sensitive export systems.
The H-index distribution further demonstrated that a substantial proportion of shipments exhibit elevated coordination sensitivity within the airport logistics environment.

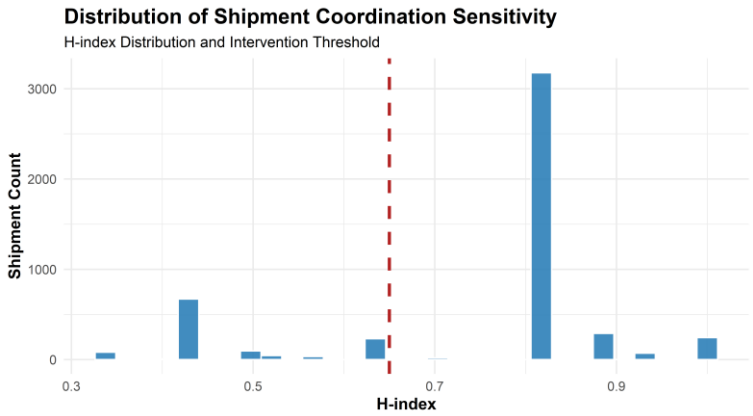


Figure 2. Distribution of Shipment Coordination Sensitivity (H-index Histogram)

The above histogram visually demonstrates the concentration of coordination-sensitive shipments and the intervention-threshold structure used within the PI-enabled coordination framework

5.2 Correlation Analysis

Table 2 presents the principal correlation results.

Table 2. Correlation Matrix

Variable	Acceptance-to-departure	Departure delay	RCS planning gap	H-index
Acceptance-to-departure	1.000	0.319	-0.276	-0.125
Departure delay	0.319	1.000	0.823	0.059
RCS planning gap	-0.276	0.823	1.000	0.134
H-index	-0.125	0.059	0.134	1.000

The results indicate that larger departure disruptions are associated with increased coordination time, while earlier cargo readiness is associated with reduced coordination delay.

These findings support the proposition that operational synchronization significantly affects airport coordination performance.

5.3 Regression Analysis

A regression analysis was implemented to evaluate operational associations between shipment coordination variables and acceptance-to-departure time.

The regression analysis should be interpreted as an operational association analysis rather than a causal predictive model because several timing variables are structurally related through shared operational timestamps.

The regression results demonstrated statistically significant associations between: operational disruption, cargo readiness timing, and airport coordination performance.

The H-index coefficient was statistically significant, suggesting that coordination-sensitive shipment structures are associated with differentiated operational coordination behavior.

5.4 Comparative PI-Enabled Simulation Results

Table 3 presents the comparative simulation-supported findings.

Table 3. Comparative Performance Between Conventional and PI-Enabled Coordination Systems

Metric	Conventional	PI-enabled
Mean lead time (hrs)	11.11	8.94
Reliability (%)	70.84	81.71
Spoilage exposure	0.366	0.295

The PI-enabled coordination environment reduced mean acceptance-to-departure time from 11.11 hours to 8.94 hours. The modeled PI-enabled environment additionally improved shipment reliability from 70.84% to 81.71%. The simulation further demonstrated reduced spoilage exposure under the PI-enabled coordination scenario.
The density-distribution comparison indicated that the PI-enabled coordination environment produced:

- reduced delay concentration,
- lower operational variability,
- and improved synchronization performance.

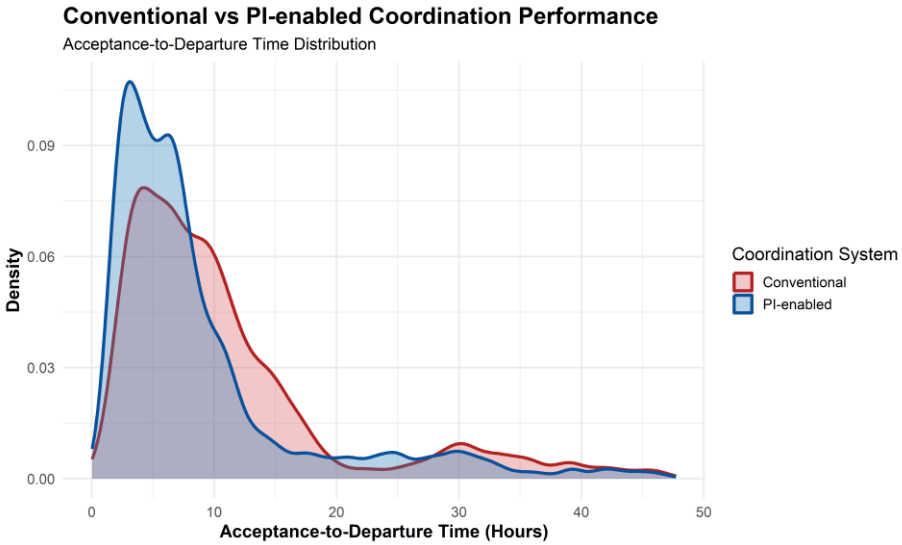


Figure 3. Conventional versus PI-enabled Coordination Performance (Density Plot)

The density-distribution comparison visually demonstrates the leftward shift of the PI-enabled coordination environment, indicating reduced coordination delay and improved synchronization performance.



Figure 4. Airport Coordination Time by Shipment Priority (Boxplot)

The boxplot illustrates differentiated operational coordination behavior across shipment-priority categories within the airport logistics environment.

5.5 Sensitivity Analysis

Sensitivity analysis demonstrated that the comparative performance of the modeled PI-enabled coordination environment remained stable across moderate variation in intervention thresholds.

Table 4. Sensitivity Analysis Results

Threshold	Intervention Count	Mean PI Lead Time	PI Reliability	PI Spoilage Exposure
0.60	4038	8.59	82.19%	0.285
0.65	3806	8.94	81.71%	0.295
0.70	3806	8.94	81.71%	0.295
0.75	3786	8.95	81.67%	0.295
0.80	3784	8.95	81.67%	0.295

The sensitivity results indicate that the comparative behavior of the PI-enabled coordination framework remained directionally stable across moderate variation in intervention thresholds.

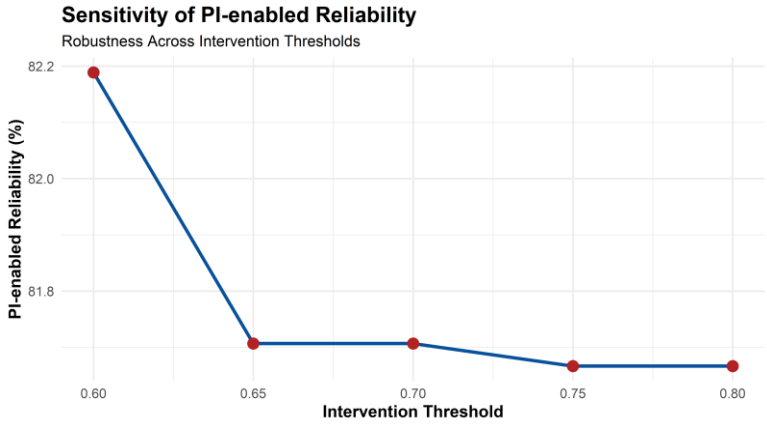


Figure 5. Sensitivity of PI-enabled Reliability Across Intervention Thresholds

The sensitivity plot demonstrates that the modeled PI-enabled coordination environment remained comparatively stable across moderate variation in intervention-threshold conditions

6 Discussion

6.1 Theoretical Contributions

The paper contributes theoretically by extending Physical Internet research toward airport-centric logistics orchestration. Existing PI literature remains heavily concentrated on terrestrial freight systems, while airport-centered logistics coordination remains comparatively underdeveloped.

The proposed framework reconceptualizes airports not merely as transportation gateways, but as intelligent PI coordination nodes capable of orchestrating adaptive synchronization,

resilience-oriented coordination, shipment-level intervention, and disruption-sensitive logistics management. The central conceptual contribution of the paper therefore lies in positioning airport-centric orchestration as a strategic extension of Physical Internet systems within perishability-sensitive export environments.
The framework additionally contributes to resilience and synchromodal logistics scholarship through the integration of: coordination intelligence, adaptive intervention, and shipment-level vulnerability assessment.

6.2. Methodological Contributions

Methodologically, the paper demonstrates how shipment-level operational air-cargo data may support simulation-oriented PI coordination modeling.
The study additionally demonstrates how airport coordination behavior may be operationalized through: timestamp-derived coordination variables, disruption-sensitive operational structures, and intervention-oriented coordination intelligence mechanisms.
The incorporation of robustness-oriented sensitivity analysis further strengthens the methodological maturity of the framework.

6.3. Practical Contributions

The findings suggest that airport-centric coordination intelligence may substantially strengthen: export resilience, cold-chain continuity, synchronization performance, and spoilage reduction within perishability-sensitive export systems.
The framework therefore provides practical implications for: airport operators, airlines, logistics planners, policymakers, and export coordination authorities.
The findings are particularly relevant for emerging economies characterized by perishability-sensitive exports, air-cargo dependence, and disruption-sensitive logistics systems. In the Ethiopian context, horticulture and time-sensitive export sectors rely heavily on coordinated airport logistics performance, making airport-centric synchronization and resilience capabilities strategically important for export competitiveness and spoilage reduction.

7 Conclusion

This paper proposed an airport-centric Physical Internet framework grounded in shipment-level coordination intelligence and resilience-oriented logistics modeling.
The framework conceptualized airports as intelligent PI coordination nodes capable of dynamically orchestrating: adaptive synchronization, shipment-level intervention, disruption-sensitive coordination, and resilience-oriented logistics management.
Using shipment-level operational data derived from Ethiopian Airlines Cargo operations, the study implemented comparative simulation-supported analysis to evaluate operational performance under both conventional and PI-enabled coordination environments.
The findings demonstrated that the PI-enabled coordination environment improved: operational lead-time performance, shipment reliability, and spoilage reduction.
The robustness-oriented sensitivity analysis further indicated that the comparative behavior of the framework remained stable across moderate variation in intervention thresholds.
The paper therefore contributes to the growing evolution of Physical Internet scholarship by extending PI research toward airport-centric coordination intelligence and perishability-sensitive export logistics systems.

8 Limitations and Future Research

This study has several limitations. First, the PI-enabled coordination environment represents a simulation-supported analytical framework rather than an empirically implemented operational system.

Second, the intervention thresholds and weighting structures were analytically assigned during the exploratory modeling stage. Third, the study focused primarily on airport-centered coordination performance and did not model wider end-to-end multimodal optimization structures.

Future research may extend the framework through: digital twin modeling, optimization-oriented intervention calibration, machine learning-supported coordination intelligence, and wider multimodal PI integration.

Future studies may additionally evaluate: airport interoperability standards, governance architectures, and large-scale cross-border PI-enabled logistics ecosystems.

References

- Aloui A., S. Hamdan, I. Tlili (2023): Digital Transformation and Airport Logistics Resilience in Air-Cargo Systems.
- Ambra T., A. Caris, C. Macharis (2019): Towards Freight Transport System Unification: Reviewing and Combining the Advancements in the Physical Internet and Synchromodal Transport Research, *International Journal of Production Research*, v57, no6, 1606-1623.
- Ballot E., F. Fontane (2008): Rendement et efficience du transport: un nouvel indicateur de performance, *Revue Française de Gestion Industrielle*, v27, no2, 41-55.
- Ballot E., B. Montreuil, R.D. Meller (2014): *The Physical Internet: The Network of Logistics Networks*, La Documentation Française.
- Christopher M., H. Peck (2004): Building the Resilient Supply Chain, *International Journal of Logistics Management*, v15, no2, 1-14.
- Crainic T.G., B. Montreuil, E. Ballot (2020): Physical Internet Enabled Hyperconnected Logistics Systems, *Progress in Material Handling Research*, 1-24.
- Giusti R., D. Manerba, R. Tadei, G. Perboli (2019): Synchromodal Logistics and Integrated Freight Transportation Systems, *Transportation Research Procedia*, v39, 105-124.
- Klibi W., A. Martel (2012): Modeling Approaches for the Design of Resilient Supply Networks under Disruptions, *International Journal of Production Economics*, v135, no2, 882-898.
- Labarthe O., W. Klibi, B. Montreuil (2023): Urban Synchromodality: Synergies between Freight and People Mobility in Urban Areas, *Proceedings of IPIC 2023 International Physical Internet Conference*, Athens, Greece, 10 p., www.pi.events
- Meller R.D., K.P. Ellis, B. Loftis (2012): From Horizontal Collaboration to the Physical Internet: Quantifying the Effects on Sustainability and Profits When Shifting to Interconnected Logistics Systems, Final Research Report of the CELDi Physical Internet Project, Phase I, http://faculty.ineg.uark.edu/rmeller/web/CELDi-PI/Final%20Report%20for%20Phase%20I.pdf, 2014/02/20.
- Merkert R., D. Alexander (2018): Air Freight Logistics and Supply Chain Management, *Journal of Air Transport Management*, v68, 1-2.
- Montreuil B. (2011): Towards a Physical Internet: Meeting the Global Logistics Sustainability Grand Challenge, *Logistics Research*, v3, no2-3, 71-87, https://doi.org/10.1007/s12159-011-0045-x
- Montreuil B. (2013): Physical Internet Manifesto, v1.10, (Original v1.0, 2009), www.physicalinternetinitiative.org, 2014/02/20.
- Pan S., E. Ballot, F. Fontane (2021): Physical Internet and Future Logistics Systems, *Sustainability*, v13, no4, 2175.

- Pfoser S., H. Treiblmaier, O. Schauer (2021): Synchromodality and Intelligent Freight Coordination Systems, *Logistics*, v5, no2, 31.
- Ponomarov S.Y., M.C. Holcomb (2009): Understanding the Concept of Supply Chain Resilience, *International Journal of Logistics Management*, v20, no1, 124-143.
- Sheffi Y. (2005): *The Resilient Enterprise: Overcoming Vulnerability for Competitive Advantage*, MIT Press.
- SteadieSeifi M., N. Dellaert, W. Nuijten, T. Van Woensel, R. Raoufi (2014): Multimodal Freight Transportation Planning: A Literature Review, *European Journal of Operational Research*, v233, no1, 1-15.
- Treiblmaier H., K. Mirkovski, P.B. Lowry, Z.G. Zacharia (2020): The Physical Internet as a New Supply Chain Paradigm, *International Journal of Logistics Management*, v31, no2, 239-258.