IPIC 2025 11th International Physical Internet Conference
18-20 June 2025 | Hong Kong, China

Physical Internet Center

Hyperconnected Facility Location Contracting

Jiachen Shi, Santanu S. Dey, Benoit Montreuil
Physical Internet Center, Supply Chain and Logistics Institute
H. Milton Stewart School of Industrial and Systems Engineering
Georgia Institute of Technology, Atlanta, U.S.A.

Corresponding author: jshi81@gatech.edu

Abstract: *Facility location problems (FLP) in supply chain management involve determining the optimal placement of facilities to support operations of supply chain networks and meet customer demand. In a traditional FLP, decision-makers must weigh the costs of location and allocation, which typically include the construction and operational costs of the facilities, and transportation expenses to customer demand points. In light of Physical Internet (PI), logistics infrastructure is shifting from private networks to an open web of interrelated networks, each involving multiple organizations collaborating to share resources, optimize operations, and enhance service responsiveness. Physical Internet enabled hyperconnected fulfillment, where fulfillment operations are no longer isolated within a retailer's private network; instead, decision makers can access a wide range of open-source fulfillment facilities that are interconnected, openly available, and their service capacity available for contract on demand. This paper studies the Hyperconnected Facility Location Contracting problem that generalizes several FLPs. The framework integrates PI principles and facility location modeling. It consists of determining contracts for open-source facilities to fulfill e-commerce demand over a multi-period planning horizon.*

Keywords: *Physical Internet, Supply Chain Networks, Facility Location Problem, Hyperconnected Fulfillment, Service Contract, Optimization.*

Physical Internet (PI) Roadmap Fitness*:* ☒ *PI Nodes (Customer Interfaces, Logistic Hubs, Deployment Centers, Factories),* ☒ *PI Networks,* ☒ *System of Logistics Networks,* ☒ *Horizontal Supply Chain Alignment.*

Targeted Delivery Mode-s*:* ☒ *Paper*

1 Introduction

Traditionally, the vast majority of e-commerce retailers operate a limited number of proprietary fulfillment centers, constraining their ability to meet demand flexibly. Most warehouses or fulfillment centers in the market are dedicated to a single retailer, leading to a sparse distribution network. This lack of geographic density poses a significant challenge, particularly for customer segments that are highly sensitive to order-to-delivery time. When consumers expect rapid fulfillment, a substantial portion of demand may remain unservable unless they are willing to tolerate extended wait times. Expanding a private fulfillment network is often prohibitively expensive, and for many companies, financial constraints allow them to establish, at most, a single additional facility.

With the emergence of hyperconnected fulfillment, decision-makers can leverage a far more extensive range of fulfillment options beyond traditional in-house investments. Instead of incurring the capital and operational expenditures associated with building an entire fulfillment network from the ground up, firms can now access open-source PI fulfillment centers through service contracts. These contracts enable businesses to dynamically scale their fulfillment capabilities by securing warehouse space with predefined capacity, duration, and timing parameters—such as peak-season prioritization or advance reservation incentives. The shift toward hyperconnected facility location contracting offers an adaptable, cost-effective alternative that enhances supply chain responsiveness, reduces last-mile delivery costs, and improves service levels for geographically dispersed customers. By integrating on-demand warehousing solutions into their logistics strategy, retailers can achieve greater flexibility, minimize infrastructure investments, and maintain a competitive edge in an era of increasingly time-sensitive e-commerce fulfillment.

In the US, companies such as Flexe and ES3 are among the top players that offer hyperconnected fulfillment to the industry. Flexe operates one of the largest on-demand warehousing networks with over 2,000 facilities across U.S. and Canada. It offers capabilities like 2-day or next-day delivery by strategically positioning facilities close to demand points. Depending on the desired delivery speed, businesses can use anywhere from a few to 16 warehouses to cover the majority of the U.S. market ES3 uses a 300-acre facility to serve retailers in the consumer goods and food industry. It handles up to 70,000 pallets and ships 5 million cases every week. It helps retailers achieve significant cost savings, while enabling higher frequencies of deliveries.

This paper leverages this new pan-industry business model and develops an optimization model for hyperconnected facility location contracting. To the best of our knowledge, this is the first systematic modeling framework to integrate PI concepts and facility location modeling. The main contribution of the paper lies in the uniqueness of the problem formulation, which dynamically constructs an entire fulfillment network from hyperconnected, contract-based facility networks while accounting for optimizing e-commerce fulfillment operations through the constructed network.

The rest of the paper is structured as the following: in section 2 of this paper, we review some of the relevant literature in Facility Location Problems. Section 3 provides a description of the overall problem settings and section 4 discusses detailed formulation of our proposed model. In Section 5, we present an experimental study and explore the model's performance and computational efficiency. Conclusions are summarized in section 6.

2 Literature Review

The Physical Internet (PI) is an innovative logistics transformation movement that aims to revolutionize global supply chains by shifting from closed, proprietary networks to open, shared systems. Montreuil (2011) defines PI as a "hyperconnected global logistics system enabling seamless open asset sharing and flow consolidation through standardized encapsulation, modularization, protocols, and interfaces" to enhance efficiency and sustainability in fulfilling physical object services. This paradigm encourages a rethinking of traditional logistics by promoting shared access to resources and open networks. The open asset utilization promoted by the PI creates opportunities for affordable access to decentralized networks, enabling companies of all sizes to leverage on-demand, open services regardless of their market share. It lays the foundation for hyperconnected logistics, characterized by multi-plane space structuring and meshed networks that interconnect hubs at multiple levels (Montreuil et al., 2018).

Hyperconnected distribution systems leverage a distributed web of openly accessible deployment centers such as long-stay warehouses, distribution centers (DCs), and fulfillment centers (FCs), owned and operated by multiple parties, offering storage and fulfillment services.

Hyperconnected distribution, in comparison to conventional dedicated and hub-and-spoke based systems, has been modeled an evaluated by Sohrabi et al. (2016) and Kim et al. (2021). Their studies demonstrated significant cost improvement solely by employing hyperconnected distribution while providing higher service levels. Under the context hyperconnected fulfillment, Pan et al. (2015) reported reductions in inventory level and total logistic costs in a multi-echelon supply chain utilizing PI hubs in a hyperconnected network, compared to a classical hierarchical supply chain, through better selection of storage locations as well as flexible and responsive replenishment plans. Although the literature on hyperconnected fulfillment is growing, most of the existing PI-based work emphasize the use of an open fulfillment network instead of how to construct it.

The decisions to be made to construct a hyperconnected fulfillment network through contracting is highly related to the field of Facility Location Problems (FLP). FLP seeks to determine the most cost-effective locations for a set of facilities to serve a set of clients through minimizing the total cost of opening facilities and serving client demands, with respect to certain constraints. Farahani and Hekmatfar (2009) presents a detailed summary on the application and classification, mathematical modeling, and solution techniques of different classes of location problems.

High costs associated with property acquisition and facility construction make facility location or relocation projects long-term investments. The simplest type of location problem is the Single Facility Location Problem, which locates a single new facility relative to a number of existing facilities. Multifacility Location Problem optimally locates more than one new facilities, where each new facility is linked to at least one other new facility (Ostresh, 1977). There are also other variations of the FLP, such as the Location-allocation (LA) problem that simultaneously determines the optimal locations for a set of new facilities and assigns demand to these facilities such that the transportation cost from facilities to customers is minimized or service efficiency maximized (Badri, 1999; Cooper, 1963). While another variation is the covering problem, where decision makers assume that each demand point can receive service by any facility if the distance them is within a threshold (Church and Velle, 1974).

The majority of models developed in the above classes of problems make locating decisions one time only. Decisions to be made to dynamically construct a network of hyperconnected facilities have a time component, as the need for the numbers, locations and capacities of contracted facilities vary over a planning horizon, and need to be made multiple times in accordance with demand fluctuation. Dynamic Facility Location Problem makes time dependent decisions, where decision makers find robust locations that serve changing demand over time, and sometimes even the timing (Daskin et al., 1992). However, the developed models are still limited to locating a fixed number of facilities over the entire planning horizon. Moreover, they can’t address unique features inherent in hyperconnected fulfillment that facilities can be contracted on a per-use basis, where contract costs are directly affected by factors such as contract length, capacity and timing.

This paper studies the Hyperconnected Facility Location Contracting problem that generalizes several FLPs and adapts to hyperconnected, contract-based networks. The proposed model accounts for multiple interconnected facilities like the Multifacility Location model, but allows for more flexibility on the number of facilities instead of a fixed number. The decision maker simultaneously optimizes the locations of facilities and assigns demand to these facilities as in the Location Allocation Problem. Our proposed model also assumes fulfillment flexibility in that each demand point can receive service by any facility if the distance is within a threshold. But unlike the general covering problem where the capacity of facilities is unlimited, our model allows capacity dimensioning at facilities (Shulman, 1991) through contracts. Inspired by key concepts of hyperconnectivity and various Facility Location Problems, our problem consists of determining contracts for open-source facilities to meet e-commerce demand over a multi-period planning horizon for a portfolio of multiple products.

3 Problem Setting

In light of hyperconnected fulfillment, e-commerce retailers can access a wide range of open-source fulfillment centers in the form of service contracts. Through these contracts, decision makers dynamically adjust the durations and capacities at which contracted facilities are operated in accordance with demand trends. This shifts fundamentally away from the traditional business model of owning a private network or long-term leasing. With these services, even a small business can leverage a dense fulfillment network, such as illustrated in Figure 1. As demand patterns shift, retailers can construct entirely different fulfillment networks with very short lead time and without the burden of property acquisition and facility construction, for example, from Figure 1 center to Figure 1 right.

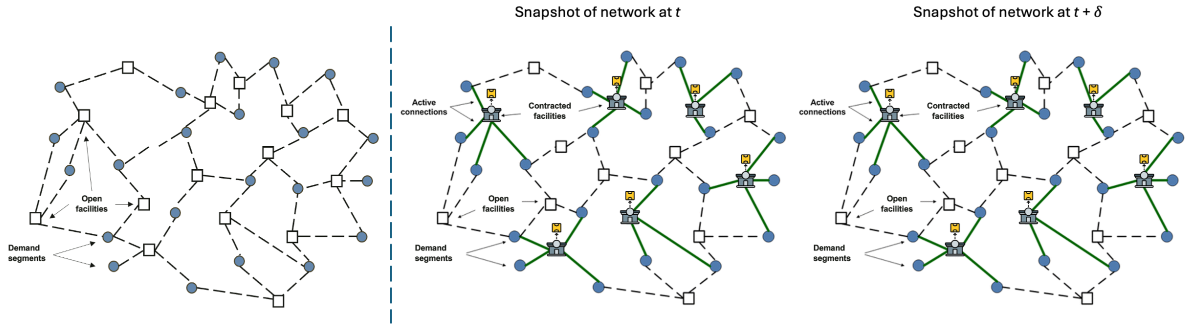

Figure 1: Hyperconnected PI facilities (left) and potential fulfillment networks that can be achieved through contracting (center and right)

We focus on a single e-commerce retailer who utilizes the PI fulfillment web through service contracts for this paper. The decision maker needs to determine a series of contracts for the interconnected PI facilities over a multi-period planning horizon to satisfy e-commerce demand. Each PI facility is subject to capacity dimensioning, where decision maker can determine different capacities to be installed with different contracts. Through the contracted network, the e-commerce retailer routes multiple products to multiple customer regions.

We consider a planning horizon that is much longer than demand trend cycles, hence there is need for multiple chances to make contract decisions. The planning horizon is therefore divided into multiple periods denoted by $\mathcal{T}$, with the following dynamic in each period $t \in [\mathcal{T}]$: at the beginning of each period and before demand realization, the decision maker has the option to enter into a contract for each of the inactive facilities (facilities not under active contract); for active facilities, each contract has to be executed until completion without breaching or amending; if any adjustments to the duration or utility is required, a new contract has to be created.

We identify four key characteristics that define a contract, namely the length of the contract; the amount of service capacities, be it for storage, processing or loading; the timing of the contract, whether it overlaps peak seasons such as Black Friday; and how much a contract is signed in advance to account for advance reservation incentives. Capacity levels are commonly defined in discrete intervals, like in a staircase function. Models that consider capacity levels as such represent capacities as modular structures (Antunes and Peeters, 2001). A module is a block of capacity associated with a unique operating cost. As a result, an example of a set of contracts to construct a hyperconnected network is illustrated in Figure 2, with multiple contracts for each facility with different utility modules, and with contracts for multiple facilities at each period over a multi-period planning horizon.

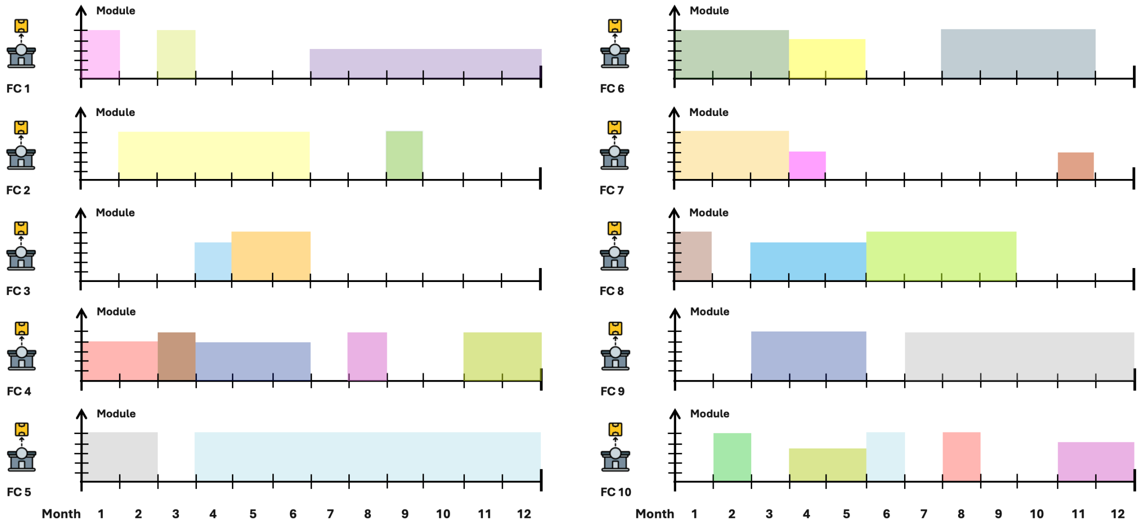

Figure 2: Example of contracts that form a hyperconnected fulfillment network over a multi-period planning horizon

Order fulfillment decisions are considered in our problem periodically with the network of contracted FCs of each period. Order fulfillment is flexible in that customer orders may be fulfilled through multiple channels. So, the problem in each specific period is effectively analogous to a covering problem, where services that customers can receive service by each facility if the distance between the customer and facility is equal or less than a predefined number. To be more specific, an order can be fulfilled by any contracted FC within a predefined radius. Alternatively, a contracted FC can fulfill e-commerce demand from any customer region within that radius.

4 Model Formulation

In this section, we formally introduce our optimization model for decision making to construct a hyperconnected fulfillment network of open-source PI facilities through contracting. The model determines the start, the end, and the module of utility contracted of each contract in optimal quantities, given different costs. The model also accounts for demand fulfillment quantities for different markets with the fulfillment network constructed with the contracted open-source FCs. We formulate the problem as a mixed integer programming (MIP) model.

In this MIP model, we introduce two types of constraints. We use constraints that regulate the timing of the contracts to ensure their legitimacy and feasibility over the entire planning horizon. We then enforce demand fulfillment constraints for each period to model the fulfillment behaviors of contracted facilities with the chosen utility modules, so that the model constructs a hyperconnected network that best meets e-commerce demand for a portfolio of multiple products from various demand sections.

The problem of determining optimal contracts of open-source fulfillment centers to form a hyperconnected network is modeled with the following parameters and variables.

Sets:

$\mathcal{F}$	set of hyperconnected fulfillment centers (FC), index by f
$\mathcal{P}$	set of unique products, index by p
$\mathcal{T}$	set of time periods, index by t
$\mathcal{K}$	set of markets, index by k
$\mathcal{L}$	set of utility modules, index by l
$\mathcal{K}_f$	subset of markets that can be served by fulfillment center f
$\mathcal{F}_k$	subset of fulfillment centers that can serve market k

Decision Variables:

$x^t_{p,f,k}$	proportion of demand for product p from market k served by FC f in period t
$y_{t_1,t_2,f,l}$	binary variable indicating a contract is bound for FC f from period t_1 to period t_2 with module l
$z^t_{f,l}$	binary variable indicating if FC f is active or under contract in period t with module l

Parameters:

$u_{f,l}$	utility of FC f associated with operating module l
$c_{t_1,t_2,f,l}$	contract cost function of a PI facility f that begins in period t_1, ends in period t_2 and contracted with utility module l

The objective of our model is to maximize the total profit across the constructed hyperconnected fulfillment network, the product portfolio and the planning horizon. The total profit of the system is defined as gross sales revenue minus contract costs. As shown in (1), the objective function accounts for gross revenue from demand fulfillment, while contracts are determined by the start, end, the contracted module of each open-source fulfillment centers:

$$\underbrace{\sum_p \sum_f \sum_k \sum_t r_p \cdot x^t_{p,f,k}}_{\text{revenue for fulfilled demand}} \quad - \quad \underbrace{\sum_{t_1} \sum_{t_2} \sum_f \sum_l c_{t_1,t_2,f,l} \cdot y_{t_1,t_2,f,l}}_{\text{contract cost}} \tag{1}$$

The constraints in this optimization model can be presented as follows: Constraint (2) is the time feasibility constraint, which defines that a contract is only feasible when it ends after it has started. Constraint (3) ensures that at most one contract can be active for any potential FC f at any given time t. Constraints (4) – (5) link the binary variable $z^t_{f,l}$, which indicates if there is an active contract at time t for facility f, with the binary variable $y_{t_1,t_2,f,l}$, which represents the begin, end, and module of a contract. Constraint (6) represents that the entire contracted hyperconnected fulfillment network combined doesn't fulfill more than there is demand. While constraint (7) stipulates that each of the contracted fulfillment centers operates within its contracted module's utility $u_{f,l}$.:

$$y_{t_1,t_2,f,l} = 0,\ \ \forall t_1 > t_2,\ \forall f \in \mathcal{F},\ \forall l \in \mathcal{L}. \tag{2}$$

$$\sum_{t_1 \le t \le t_2} \sum_l y_{t_1,t_2,f,l} \le 1,\ \ \forall f \in \mathcal{F},\ \forall t \in \mathcal{T} \tag{3}$$

$$z^t_{f,l} \ge y_{t_1,t_2,f,l},\ \ \forall t_1 \le t \le t_2,\ \forall f \in \mathcal{F},\ \forall l \in \mathcal{L} \tag{4}$$

$$z^t_{f,l} \le \sum_{t_1 \le t \le t_2} y_{t_1,t_2,f,l}\,,\ \ \forall t \in \mathcal{T},\ \forall f \in \mathcal{F},\ \forall l \in \mathcal{L} \tag{5}$$

$$\sum_f x_{p,f,k}^t \leq d_{p,k}^t, \ \ \forall p \in \mathcal{P}, \forall k \in \mathcal{K}, \forall t \in \mathcal{T} \tag{6}$$

$$\sum_p \sum_k x_{p,f,k}^t \leq \sum_l u_{f,l} \cdot z_{f,l}^t, \ \ \forall f \in \mathcal{F}, \forall t \in \mathcal{T} \tag{7}$$

$$x_{p,f,k}^t \geq 0, \ y_{t_1,t_2,f,l} \in B, \ z_{f,l}^t \in B, \forall\, p \in \mathcal{P}, \ \forall f \in \mathcal{F}, \ \ \forall l \in \mathcal{L}, \forall t_1, t_2 \in \mathcal{T} \tag{8}$$

5 Computational Experiment

In this section, we develop an experimental framework to study and evaluate the performance of the proposed hyperconnected facility location contracting model. The case study considers a single e-commerce retailer with access to a number of interconnected, open-source fulfillment centers located over the United States. The e-commerce retailer routes a portfolio of 44 products jointly to 98 demand segments represented by the 3-digit zip codes of US through the network of FCs. We simulated facility contracting and fulfillment operations over a 12-month period (2020/01/01 – 2020/12/31) with one-month intervals. Real data for sales, prices, costs and product information from an online e-commerce company were used to conduct the computational experiments.

We develop three scenarios of facility density over the US markets: low density, moderate density and high density. We identified the top 50 metropolitan statistical areas based on total population (U.S. Census Bureau, 2024) and divided them into three tiers: tier I metro areas (population > 5,000,000); tier II metro areas (population > 2,000,000); and tier III consisting of the rest. This gave us a total of 10 tier I areas, 25 tier II areas and 15 tier III areas, and we construct facility density scenarios in proportions to population densities. For the low facility density case, one open-source facility is allocated to each of the tier I metro areas, giving a network of 10 facilities. Two facilities are allocated to each of the tier I metro areas to create the moderate facility density case, in addition to one facility to each of the tier II metro areas, which gives 45 facilities combined. Finally, five facilities are allocated to each of the tier I areas, two to each of the tier II areas and one to each of the tier III areas, which constructs the scenario of high facility density with a total of 115 hyperconnected open facilities.

The results of with the three different facility density scenarios are illustrated below in Figure 3, with the color of the circle interior representing the proportion of time a facility is contracted, the color of the circle border representing the number of contracts as an indication of the complexity of contracting at a specific facility, and the size of the circles representing the amount of flow through each contracted facility. Furthermore, a detailed representation of contracts spanning over the multi-period planning horizon for the 10 facilities case can be found in Figure 2.

5.1 Model Performance Results

The goal of this study is to examine the benefits of leveraging open PI fulfillment facilities to dynamically construct a fulfillment network through contracting. We focus on several key performance metrics including total profit, contract cost, and service level. The experiment compares the proposed hyperconnected facility contracting approach with a static benchmark in accordance with common industry practice. The benchmark is analogous to a traditional multifacility location problem, where the model selects a number of facilities from the set of available open facilities to construct the fulfillment network at the beginning of the planning horizon. The facilities can only be leased at their full capacity and for the entire horizon, as a result, the fulfillment network remains unchanged throughout.

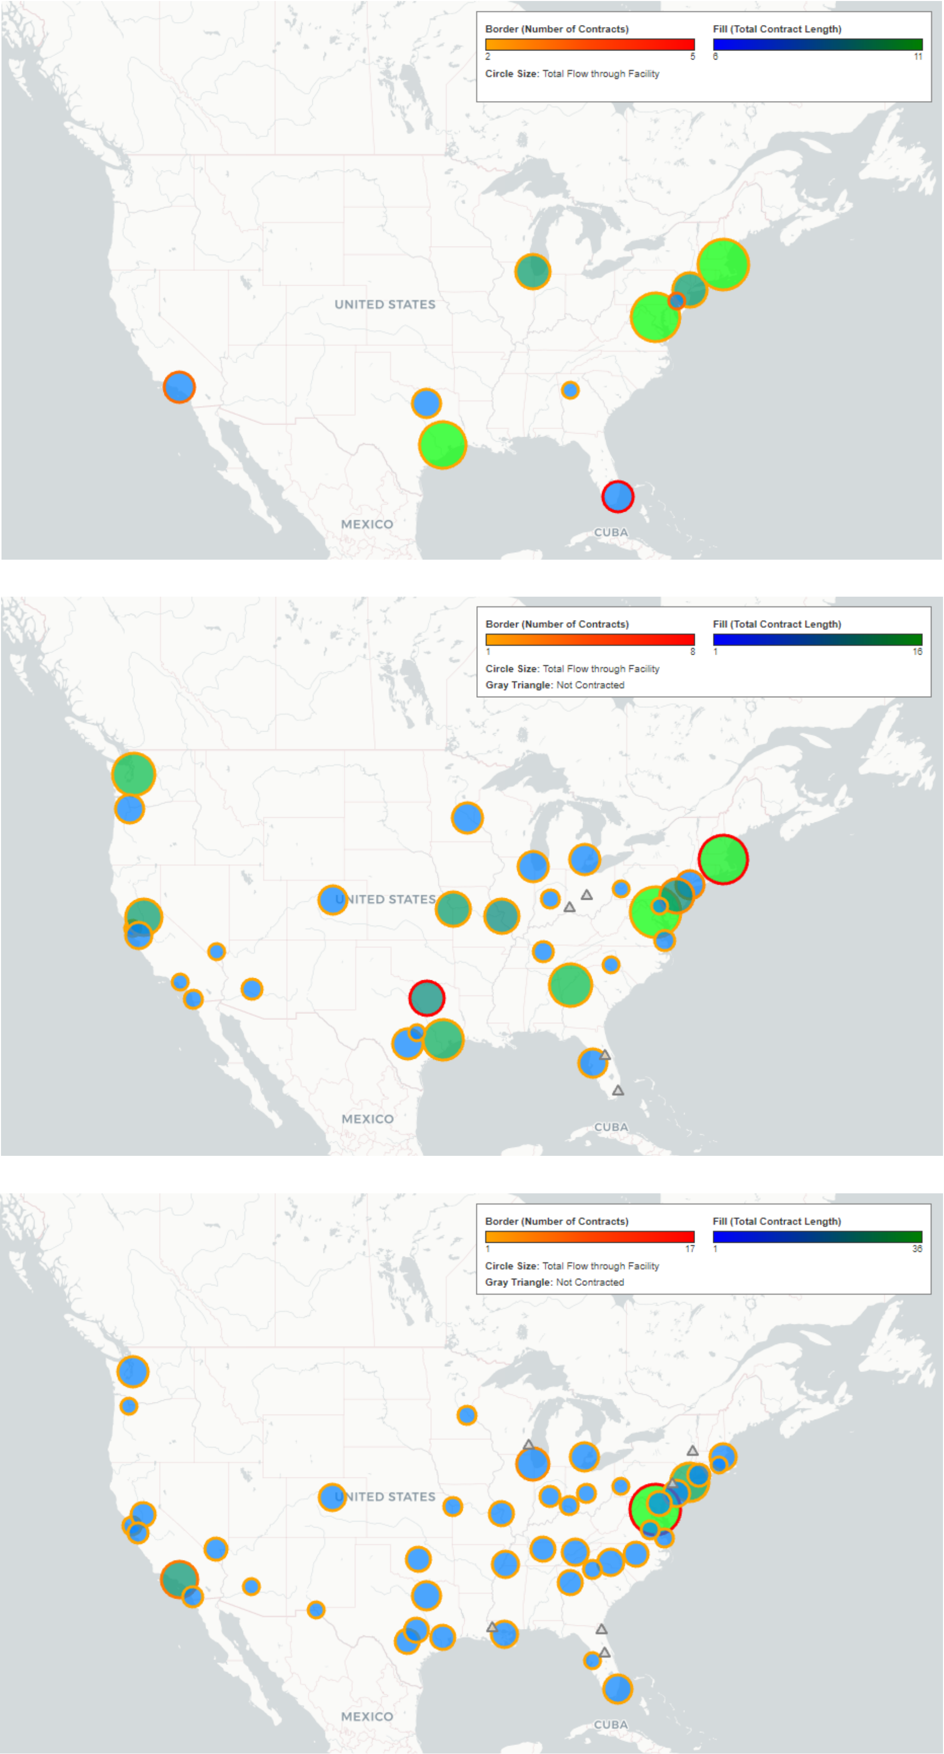

Figure 3: Results with low (top), medium (middle), and high (bottom) facility density scenarios

Table 1: Model performance compared to benchmark

	Benchmark	Proposed Model
Total Revenue	844486	915235
Contract Cost	500520	429191
Total Profit	343966	486044

The results of this case study are presented in Table 1. We observed that our proposed model reduced total contract cost significantly, by 14.2%, compared to the benchmark model. This was achieved through the flexibility of dynamically adjusting service contracts. It is worth noticing that our model fulfilled more demand and achieved a higher total revenue, 8.4% more compared to benchmark, while saving on contract costs. As a result, our proposed model increased total profit by 29.2%.

5.2 Computational Efficiency

We compare the computation time of our proposed model under different problem instances to demonstrate how the MILP scales as the network of open-source facilities grows. To limit the source of error, we run the different instances using the exact same initial conditions. We repeat the experiment for all settings 25 times, where our model is tested with 10, 45 and 115 open FCs respectively. Each run is timed and recorded, and the runtimes are presented in a table.

Table 2: Comparative analysis on model computation times

FC Density Scenario	Runtime Mean (sec)	Runtime Standard Dev (sec)
Low (10)	600.13	163.77
Medium (45)	2656.24	987.54
High (115)	9522.69	2677.34

The results of this study are presented in Table 2. On average, we observed that as the number of open-source FCs increased, the runtime grew more than linearly with problem scale. Moreover, the variance of the computation times also increased with the number of open facilities, i.e., as the size of the problem instances increased. These observations demonstrate that our original proposed formulation may face challenges from the standpoint of computational efficiency and problem scalability.

6 Conclusion

We proposed an innovative mixed integer programming model to solve the problem of hyperconnected facility location contracting, integrating Physical Internet concepts into facility location modeling. The model leverages service contracts to dynamically utilize open-source PI facilities and optimizes e-commerce fulfillment through the constructed hyperconnected, contract-based facility networks. The proposed framework shifts fundamentally away from stylized assumptions and structures such as private supply chain networks, and offers promising insights into how a new business model like hyperconnected fulfillment can reshape e-commerce logistics, paving the way for more flexible and responsive supply chain networks.

Application of the proposed framework has been demonstrated through a real-world e-commerce retailer. Experiments were conducted to examine the proposed approach under difference scenarios of PI facility density. The proposed approach achieved improvements both in terms of saving costs and increasing service level and profit compared to a benchmark approach in accordance with common industry practice. Scalability issues of the model arise as the number of open facilities increases. Future research will explore various computational and analytical approaches to effectively manage and mitigate these challenges. Algorithms like Benders decomposition and analytical solution methodologies like product clustering hold promise for advancing the theoretical robustness and practical applicability of our framework.

References

- Antunes A., D. Peeters (2001): On solving complex multi-period location models using simulated annealing. European Journal of Operational Research, v130, no1, 190–201.
- Badri M.A. (1999): Combining the analytic hierarchy process and goal programming for global facility location-allocation problem. International J. of Production Economics, v62, no3, 237-248.
- Church R., C.R. Velle (1974): The maximal covering location problem. Papers in regional science, v32, no1, 101-118.
- Cooper L. (1963): Location-allocation problems. Operations research, v11, no3, 331-343.
- Daskin M.S., W.J. Hopp, B. Medina (1992): Forecast horizons and dynamic facility location planning. Annals of Operations Research, v40, 125-151.
- Farahani R.Z., M. Hekmatfar (2009): Facility location: concepts, models, algorithms and case studies. Springer Science & Business Media.
- Kim, N., B. Montreuil, W. Klibi, N. Kholgade (2021). Hyperconnected urban fulfillment and delivery. Transportation Research Part E: Logistics and Transportation Review, 145, 102104.
- Montreuil, B. (2011): Toward a Physical Internet: meeting the global logistics sustainability grand challenge. Logistics Research, v3, 71-87.
- Montreuil, B., S. M. Buckley, L. Faugère, R. Khir, & S. Derhami (2018): Urban Parcel Logistics Hub and Network Design: The Impact of Modularity and Hyperconnectivity.
- Ostresh L.M. (1977): The multifacility location problem: Applications and descent theorems. Journal of Regional Science, v17, no3.
- Pan, S., M. Nigrelli, E. Ballot, Y. Yang (2015): Perspectives of inventory control models in the Physical Internet: A simulation study. Computers & Industrial Engineering, v84, 122–132.
- Sohrabi, H., B. Montreuil, W. Klibi (2016): On Comparing Dedicated and Hyperconnected Distribution Systems: An Optimization-Based Approach. Information Systems Logistics and Supply Chain 6th International Conference.
- Shulman A. (1991): An algorithm for solving dynamic capacitated plant location problems with discrete expansion sizes. Operations research, v39, no3, 423-436.
- U.S. Census Bureau (2024): Annual Estimates of the Resident Population for Metropolitan and Micropolitan Statistical Areas: April 1, 2020 to July 1, 2024, https://www.census.gov/data/tables/time-series/demo/popest/2020s-total-metro-and-micro-statistical-areas.html, 2025/04/28.