Integration of Hub Capacity Acquisition Decisions in the Scheduled Service Network Design Problem

Simon Belieres[a], Yannick Scherr[b], Mike Hewitt[c]

TBS Toulouse Business School[a], Toulouse, France, Loyola University Chicago[b], Chicago, USA, University of Vienna[c], Vienna, Austria

Extended abstract submitted for presentation at the 12th Triennial Symposium on Transportation Analysis conference (TRISTAN XII)
June 22-27, 2025, Okinawa, Japan

February 13, 2025

Keywords: Service Network Design; Hub capacity; Fourth-party logistics; IP-based heuristics

1 INTRODUCTION

Cost-effective middle-mile transportation relies on efficiency and high load consolidation. One can assist this tactical decision-making by solving the Scheduled Service Network Design Problem (SSNDP), which finds applications in various transportation sectors. In this study, we focus on less-than-truckload (LTL) transportation. The SSNDP has received increasing attention over the past decades, with significant contributions regarding solution strategies, as well as modeling more realistic variants that integrate practical operational constraints, such as terminal restrictions on the number of vehicles that can be loaded or unloaded per time period (Wu *et al.* (2023), He *et al.* (2022)). These terminal capacities pertain to facility location and sizing decisions, which typically fall under strategic planning due to the substantial investments and long-term effects involved. This is particularly true for third-party logistics (3PLs) providers who own their warehousing and transportation assets. Nevertheless, the growth of e-commerce applications has highlighted the need to improve the efficiency of a fragmented logistics industry (Jiang *et al.* (2024)), leading to increasing integration and coordination of assets among multiple 3PLs facilitated by fourth-party logistics (4PLs) providers. 4PLs are typically asset-less and focus on distributing freight through a middle-mile network using assets owned by 3PLs, including warehousing, vehicles, and loading/unloading capacity at terminal facilities, commonly known as "dock doors". In this context, 4PLs lease dock doors for relatively short durations, such that terminal capacity-related decisions can be made at the tactical planning stage rather than on a longer-term horizon. Our contributions are fourfold: (i) we introduce a mathematical formulation for the SSNDP with Hub Capacity Acquisition (SSNDP-HCA), (ii) we design an IP-based heuristic for this problem, (iii) we computationally demonstrate the effectiveness of the proposed approach in finding high-quality feasible solutions to the SSNDP-HCA in reasonable times, and (iv) we assess the value of integrating terminal capacity-related decisions into the SSNDP.

2 METHODOLOGY

We first provide an overview of the problem. We then outline the main components of our solution algorithm.

2.1 Problem Description

The 4PL aims to transport a set of commodities through the terminal network, with the requirement that commodities having the same origin and the same destination must follow the same physical path. The 4PL enforces the use of a consistent route for all shipments with the same origin and destination to reduce operational complexity and enhance service quality. The 4PL can allocate vehicles to point-to-point moves, thus incurring fixed costs. The allocation of a vehicle to a point-to-point move encompasses (i) a loading operation at the origin terminal, (ii) a traveling operation, and (iii) an unloading operation at the destination terminal. As such, the duration of these three operations must be observed. Additionally, a vehicle arriving at a terminal can wait before the start of its unloading operation, and we assume that terminals have unlimited parking and storage space. Without loss of generality, we assume that a vehicle departs as soon as it is loaded. Each terminal has a limited number of dock doors available, and its loading/unloading capacity is proportional to the number of inbound/outbound doors acquired. Acquiring a dock door incurs a fixed cost. The SSNDP-HCA addresses three levels of decisions, namely, (i) the acquisition of available dock doors at terminals, (ii) the selection of the services and the corresponding capacities to support the movements of freight through the terminal network, and (iii) the choice of the paths followed by the commodities from their origins to their destination as well as the timing of the operations.

2.2 IP-based local search heuristic

We develop an IP-based local search (IPBLS) approach to compute high-quality primal solutions in reasonable times. IPBLS approaches aim to improve a feasible solution by solving carefully chosen restricted MIPs. More precisely, the solution of these restricted MIPs reflects the exploration of a neighborhood and corresponds to the re-optimization of a subset of the decision variables. As these restricted MIPs are solved repeatedly throughout the process, their definition is crucial. Specifically, they need to be tractable while allowing the exploration of a sufficiently large neighborhood for the IPBLS to be effective.

We extend the recently proposed *consolidation-based* formulation for the SSNDP (Hewitt & Lehuédé, 2023) to define these restricted MIPs. Specifically, we add variables and constraints to model the workload caused by loading and unloading vehicles at terminals, thereby enforcing limitations on processing capacities. While *traditional formulations* rely on time-expanded networks to capture the synchronization of vehicles and shipments needed for consolidation, *consolidation-based* formulations capture these features by relying solely on the physical network. On the one hand, consolidation-based formulations enable breaking symmetries induced by time-expanded networks and provide tighter linear relaxations. On the other hand, these formulations require the enumeration of all the possible consolidations and may be intractable. To circumvent this issue, we propose a hierarchical strategy where we first select the physical path of all commodities, thus significantly reducing the number of potential consolidations. Based on those decisions, we can formulate a tractable consolidation-based formulation that optimizes the remaining decisions. The mathematical formulation we introduce does model terminal vehicle processing capacities, a feature that has not yet been proposed in the literature, to the best of our knowledge.

Our IPBLS approach first constructs a feasible solution S. It then iteratively applies a Variable Neighborhood Descent that consists of four local search operators, until a stopping criterion is met. Each local search operator modifies the path of one or multiple commodities. Based on those decisions, the network is partitioned into a first sub-network, which is impacted by the path modifications, and a second sub-network, where commodity paths remain unchanged. Specifically, the first sub-network includes all the nodes and arcs involved in the previous and/or the new path of a commodity, while the second sub-network includes all the remaining nodes and arcs. Then, a restricted MIP reflecting different re-optimization processes between both sub-networks is constructed and solved. On the first sub-network, all the decision variables are re-optimized, namely, the choice of the consolidations, the timing of their loading/unloading operations, the allocation of the vehicles, and the acquisition of inbound/outbound doors. On the second sub-network, the re-optimization process solely involves re-scheduling the loading and unloading operations associated with the selected consolidations in the current solution S.

The local search operator **Alleviate-Hubs** aims to reduce the door assignment cost. Specifically, it intends to reduce the number of required doors of a given terminal by changing the path of a minimum number of commodities. The local search operators **Attract-Arc** and **Alleviate-Arc** aim to reduce fixed service costs by improving the vehicle average fill rate. Specifically, they identify a physical arc with a low vehicle average fill rate, and they modify commodity paths such that more, or less, commodities flow along that arc, respectively. The local search operator **Individual-Rerouting** procedure is a classical removal-insertion procedure.

3 RESULTS

We conduct a series of experiments to evaluate the performance of the proposed algorithm. Our instances are derived from data provided by our industrial partner and depict three physical networks corresponding to distinct sub-parts of the carrier's complete network. Each network corresponds to a combination of US states. Network features are summarized in Table 1.

Network	States	Nodes	Arcs	Hubs	Spokes
ALGA	Alabama/Georgia	11	96	4	7
FLGA	Florida/Georgia	14	156	6	8
ALFLSC	Alabama/Florida/South Carolina	20	355	8	12

Table 1 – *Features of the physical networks*

The carrier provides us with certain values associated with the physical network, including fixed costs, variable costs, and travel times. Service capacities are also provided, and are uniform with a value of $22,000$ for all point-to-point moves. The number of available dock doors per terminal in the network is set to 16 units. Fixed costs incurred from door acquisition are generated using the parameter r, which provides a ratio between the average fixed service cost and the average door fixed cost. Ultimately, we define two bounds $[r_1, r_2]$ around r, with $r - r_1 = r_2 - r$. For each terminal, the door fixed cost is obtained by randomly drawing a value from a uniform distribution over the interval $[r_1, r_2]$ and multiplying this value by the average service fixed cost. Shipments are generated randomly, and their sizes are determined by randomly selecting a value from a normal distribution $\mathcal{N}(\mu, \sigma^2)$. We consider a time horizon of 24 hours, discretized into time periods of 30 min. The duration of a loading/unloading operation is set to 30 min, and the maximum waiting time at a terminal before unloading is set to 60 min.

Ultimately, we create instances based on the three physical networks. We use the value ranges $[1.9, 2.1]$, $[3.9, 4.1]$, and $[5.9, 6.1]$ for the bounds $[r_1, r_2]$. We use three normal distributions to generate the shipment sizes, with parameter values ($\mu = 5000, \sigma^2 = 2500$), ($\mu = 7000, \sigma^2 = 3500$), and ($\mu = 10000, \sigma^2 = 5000$). The sets of commodities range in cardinality from 50 to 200 in increments of 50. We also create instances with 400 commodities. For each combination of parameter values, we generate 3 instances with different seeds, resulting in 405 instances overall. Instances with 400 commodities are referred to as *large* instances, while the remaining are referred to as *small*.

Two methods are used to solve the problem. The benchmark method, *MIP*, involves solving the SSNDP-HCA mathematical formulation with an industrial solver. The second method, *IPBLS*, involves applying our algorithmic strategy. All MIPs are solved using CPLEX v22.1 with a 1% MIP tolerance and default parameter values. The algorithms and mathematical models are coded in C++. The time limit is set to 1 hour for *MIP* and 15 minutes for *IPBLS*.

We investigate the ability of *IPBLS* to generate high-quality primal solutions. Thus, for each instance, we measure a primal gap between the solution produced by *MIP* and that produced by *IPBLS*: $UB_{gap} = \frac{UB_{IPBLS} - UB_{MIP}}{UB_{IPBLS}}$. A negative ratio indicates that the solution returned by

IPBLS has a lower cost than that of *MIP*. For each instance, we also measure the ratio between the computational times of *IPBLS* and *MIP*: $\rho_{time} = \frac{time_{IPBLS}}{time_{MIP}}$.

We first discuss the results obtained for the *small* instances, amongst which 159 were solved within the time limit by *MIP*. For the 165 remaining *small* instances, *MIP* failed to find bounds satisfying the 1% optimality tolerance within the time limit. Figure 1 provides the distribution of the instances according to their primal gap, discriminating solved and open instances.

UB_{gap}	0%	$]0\%, 1\%]$	$]1\%, 2\%]$	$]2\%, 3\%]$	+3%
# inst.	31	80	35	8	5

(a) *Solved instances*

UB_{gap}	0%	$]0\%, 1\%]$	$]1\%, 2\%]$	$]2\%, 3\%]$	+3%
# inst.	3	44	46	36	36

(b) *Open instances*

Figure 1 – *Distribution of the instances based on the obtained primal gap*

IPBLS yields a primal gap lower than 2% for 90% of the solved instances. These close-to-optimum solutions are found in 90 seconds on average for *IBPLS*, against 321 seconds for *MIP*. For the open instances, *IPBLS* yields a primal gap lower than 3% in 75% of the cases. For these instances, the computation time of *IPBLS* is significantly smaller than that of *MIP*, as the time ratio ρ_{time} averages a value of 7.5%. Results for the *large* instances are summarized in Table 2.

Network (μ, σ^2)	r	ALGA UB_{gap}	FLGA UB_{gap}	ALFLSC UB_{gap}
(5000,2500)	2	2.05%	-6.92%	-3.96%
	4	-3.54%	-3.86%	-8.98%
	6	0.56%	-5.43%	-3.49%
(7000,3500)	2	-4.09%	-8.04%	-11.20%
	4	-4.05%	-9.67%	-13.37%
	6	-8.43%	-5.65%	-21.40%
(10000,5000)	2	-1.99%	-9.31%	-8.84%
	4	-2.70%	-3.41%	-12.90%
	6	-3.74%	-2.20%	-9.75%

Table 2 – *Performance of IPBLS for the large instances*

Overall, the *IPBLS* provides solutions that are 6.72% better than those obtained from *MIP* and these solutions are achieved in no more than a quarter of the time required by *MIP*, as *MIP* reaches the 1-hour time limit for all of the large instances. Notably, the heuristic outperforms the commercial solver for 25 out of 27 classes of instances. To understand the benefits of incorporating terminal capacity-related decisions into the SSNDP, we compare the solutions produced by *MIP* those generated by a practice-oriented sequential decision-making process. For the small instances that are solved to optimality by *MIP*, the value of integration reaches up to 10%. A qualitative analysis of the solutions shows that the routing decisions differ between the integrated approach and the sequential method, which helps reducing the overall costs.

References

He, Edward, Boland, Natashia, Nemhauser, George, & Savelsbergh, Martin. 2022. An exact algorithm for the service network design problem with hub capacity constraints. *Networks*, **80**(4), 572–596.

Hewitt, Mike, & Lehuédé, Fabien. 2023. New Formulations for the Scheduled Service Network Design Problem. *Transportation Research Part B: Methodological*, **172**, 117–133.

Jiang, Songchen, Huang, Min, Zhang, Yuxin, Wang, Xingwei, & Fang, Shu-Cherng. 2024. Fourth-party logistics network design with demand surge: A greedy scenario-reduction and scenario-price based decomposition algorithm. *International Journal of Production Economics*, **269**, 109135.

Wu, Haotian, Herszterg, Ian, Savelsbergh, Martin, & Huang, Yixiao. 2023. Service network design for same-day delivery with hub capacity constraints. *Transportation Science*, **57**(1), 273–287.