CIRRELT-2026-09

Effective Policies for Addressing the Booking Problem in Intermodal Barge Transportation

Bita Payami
Ioana C. Bilegan
Teodor Gabriel Crainic
Walter Rei

March 2026

Bureau de Montréal
Université de Montréal
C.P. 6128, succ. Centre-Ville
Montréal (Québec) H3C 3J7
Tél : 1-514-343-7575
Télécopie : 1-514-343-7121

Bureau de Québec
Université Laval,
2325, rue de la Terrasse
Pavillon Palasis-Prince, local 2415
Québec (Québec) G1V 0A6
Tél : 1-418-656-2073
Télécopie : 1-418-656-2624

Effective Policies for Addressing the Booking Problem in Intermodal Barge Transportation

Bita Payami[1], Ioana C. Bilegan[2], Teodor Gabriel Crainic[1,*], Walter Rei[1]

[1] Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT) and School of Management, Université du Québec à Montréal

[2] Université Polytechnique Hauts-de-France, LAMIH, CNRS, UMR 8201 F-59313 Valenciennes, France

Abstract. This paper addresses a booking problem in intermodal barge transportation, where a carrier decides whether to accept or reject sequentially arriving shipment requests to maximize profit using fixed, vessel-supported services with predefined schedules and routes. In addition to these services, the carrier may rely on outsourced services with higher costs, provided that accepting the request remains profitable, i.e., cost-efficient relative to the revenue generated. The problem is inherently dynamic due to continuously arriving requests and is further complicated in inland waterway systems by fluctuating water levels that directly affect vessel capacity and service feasibility. We develop a bin-packing-based framework to model this booking problem under various decision policies that differ in timing and in their ability to anticipate future requests. In this context, we extend the classical bin packing problem to a time–space network setting, where “bins” represent scheduled, vessel-supported services defined by capacity and cost as well as route, departure time, and travel duration. “Items” represent shipment requests characterized by origin–destination, availability and due times, and economic value. Assignment decisions must satisfy spatial–temporal compatibility and profitability, whether using regular services or outsourcing alternatives. To capture the evolution of accepted requests and the operational impact of fluctuating water levels on vessel capacity, the model is embedded into a rolling-horizon framework. At each decision point, the model accounts for current requests, previously accepted requests, and, if applicable, predicted future requests and updated water-level forecasts. The performance of different booking policies is evaluated through extensive numerical experiments using a commercial solver.

Keywords: Booking problem; myopic policy; look-ahead policy; intermodal transportation; Item-to-bin assignment

Acknowledgements. While working on the project, the first author was Ph.D. candidate at the School of Management, Université du Québec à Montréal (UQAM), Canada, and member of CIRRELT. The third author held the UQAM Chair in Intelligent Logistics and Transportation Systems Planning and was an Adjunct Professor in the Department of Computer Science and Operations Research at the Université de Montréal, while the fourth author held the Canada Research Chair in Stochastic Optimization of Transport and Logistics Systems. We gratefully acknowledge the financial support provided by the Natural Sciences and Engineering Research Council of Canada (NSERC), through its Collaborative Research and Development, and Discovery grant programs. We also gratefully acknowledge the support of Fonds de recherche du Québec through their infrastructure grants. Finally, during the preparation of this work the authors used ChatGPT-4, under its most strict privacy enforcing settings, for text edit purposes exclusively. After using this tool, the authors reviewed and further edited the text as needed and take full responsibility for the content of the publication.

---

* Corresponding author: teodorgabriel.crainic@cirrelt.net

Dépôt légal – Bibliothèque et Archives nationales du Québec
Bibliothèque et Archives Canada, 2026

1 Introduction

Intermodal freight transportation refers to the movement of freight using a combination of at least two distinct transportation modes, with transfers between modes occurring at designated intermodal terminals (such as seaports, inland ports, or rail yards) without the need to handle the freight itself (Bektaş and Crainic, 2008; SteadieSeifi et al., 2014). Throughout the journey, each shipment typically remains within a single loading unit, most commonly a standardized container, to ensure continuity, reduce transshipment costs, and facilitate interoperability across modes. Intermodal freight transportation systems are designed to capitalize on the unique advantages of each mode along different segments of the network: for example, road transport offers accessibility and flexibility for local pickup and delivery; rail provides cost efficiency and scalability for long-haul transport; and inland waterway transport offers environmental benefits and efficiency for medium-distance segments between inland terminals. A significant share of inland and medium- to long-distance intermodal freight is managed by consolidation-based carriers, which group shipments from multiple shippers into shared transportation services. This consolidation strategy enables carriers to leverage economies of scale, optimize capacity utilization, and reduce per-unit transportation costs. These carriers typically operate along fixed service routes with regular schedules defined over short-term planning cycles (e.g., weekly), which are repeated consistently over a longer planning horizon (e.g., a season) (Crainic and Rei, 2024). For instance, a barge-based consolidation carrier may offer weekly round-trip services connecting inland terminals such as Antwerp, Rotterdam, and Strasbourg, with predefined departure and arrival times at each terminal that remain fixed and are repeated weekly throughout the season.

Shippers requesting these transportation services can generally be classified into two categories. The first category comprises contractual shippers, who sign formal agreements with the carrier prior to the start of the season. These contracts define the key characteristics of the shipments in advance, including the origin and destination terminals, the availability and due times, and the shipment size. They also define the obligation for the carrier to fulfill their transportation requests throughout the season. The second category consists of non-contractual shippers, who do not have such pre-established agreements but may submit shipment requests during the season, typically several days before the shipment becomes available at the origin terminal. Each non-contract-shipment request specifies the origin–destination pair, the shipment size, and the associated availability and due time, all within the same season. Non-contract-shipment requests can be submitted through various channels, including digital platforms, phone calls, or email, and are evaluated by a general booking system, which determines whether to accept or reject them. Importantly, booking decisions must be made before the shipment becomes physically available, often immediately upon request submission or within a short, controlled delay upon receiving non-contract-shipment requests. Early decisions are essential to enable shippers with accepted requests to prepare their shipment and allow shippers with rejected requests sufficient time to seek alternative carriers.

However, because shipments are not yet physically available at the time of booking, they cannot be directly assigned to specific scheduled services. In other words, an optimal shipment itinerary cannot be determined, nor can actual service capacity be allocated, as would occur during operational planning once shipment availability information is fully revealed (i.e., when shipments are available). Consequently, the booking system must make accept/reject decisions before operational planning can take place, relying only on an estimation of the shipment's potential operational plan (i.e., the likely shipment itinerary). This requires both a conceptual assessment of whether the shipment can be accommodated within the scheduled service network and a quantitative estimation of its contribution to overall profitability, including transportation costs and potential outsourcing costs, while remaining committed to all seasonal obligations for contract-shipment requests. Such evaluations provide the foundation for making informed and economically rational accept/reject decisions for non-contract-shipment requests.

One of the key challenges faced by carriers, yet insufficiently addressed in the existing literature, is the lack of effective decision-support tools to guide sequential acceptance decisions in dynamic booking environments. Existing methodologies for determining optimal shipment itineraries, most often used in operational planning, such as multi-commodity network flow models, are not well suited to this context, as they require detailed shipment itineraries. These models must specify exactly which terminals each shipment passes through, including the departure and arrival times at each visited terminal, the specific times at which loading and unloading occur, and the consolidation of shipments along the journey. Time–space networks are the modelling instruments typically used to capture this level of detail (Crainic and Hewitt, 2021). To represent the temporal dimension, the network is divided into many small time intervals, creating a separate copy of each terminal for every discretized time step. This approach causes the model to become extremely large, often containing thousands of nodes and hundreds of thousands of arcs (i.e., arcs between different terminals representing movement over time and arcs between the same terminal representing loading, unloading, and holding activities over time). In large operational networks, adding separate decision variables for the routing of both contract-shipment requests (which is necessary because, when accepting non-contract-shipment requests, the carrier must remain committed to fulfilling all seasonal obligations for contract-shipment requests) and non-contract-shipment requests further increases the size, resulting in mixed-integer programming problems that are too large to solve within reasonable time. This paper addresses this gap by introducing a novel framework that enables carriers to make informed booking decisions(i.e., to accept only profitable non-contract-shipment requests) based on conceptual and quantitative evaluations that estimate each request's likely shipment itinerary and its contribution to overall profitability. In this framework, the objective is not to establish a detailed load plan at the time of booking, but rather to quickly assess origin–destination and time compatibility between shipments and scheduled services, along with an estimation of profitability. The proposed approach is designed for speed, responsiveness, and scalability, offering a practical alternative that is well aligned with the requirements of freight booking systems, where immediate decisions are often necessary.

We propose three models, formulated within a bin packing framework, each defined by a specific policy to address the dynamic nature of the booking environment. In such environments, booking decisions must be made immediately or within a short, controlled delay upon receiving non-contract-shipment requests. This inherent immediacy forces the carrier to decide with incomplete knowledge of potential future non-contract-shipment requests, creating a trade-off between accepting requests now and keeping acceptance decisions conservative in order to preserve opportunities for potentially more profitable requests later. When non-contract-shipment requests are highly uncertain and cannot be reliably forecast, our first myopic policy is the natural choice, making accept/reject decisions without anticipating future requests. When non-contract-shipment requests exhibit a degree of predictability, our second lookahead policy can exploit shipment request forecasts. In the context of intermodal barge transportation, the decision-making process can also be influenced by predictable environmental factors such as water-level variations, which affect navigability, vessel capacity, and feasible shipment routing (see, e.g., Payami et al. (2025a), Payami et al. (2025b)). Our third policy, a lookahead with environmental forecasting, extends the second one by also considering predictable changes in water levels. This enables more informed and robust planning decisions that account for both forecasted demand and anticipated environmental conditions.

The contributions of this paper are as follows: (i) We introduce a comprehensive booking-level decision problem for consolidation-based intermodal freight carriers operating over fixed service networks, where sequential accept/reject decisions must be made under incomplete information, accounting for both contractual shipment commitments and dynamic non-contract-shipment requests. (ii) We propose a bin packing–based modelling framework specifically tailored to the booking context. The framework is designed for rapid decision-making and scalability, making it particularly suitable for environments where immediate responses are required. (iii) We design booking policies reflecting different levels of predictability, including a myopic policy and a lookahead policy. (iv) While the proposed modelling framework is general, we apply it to inland barge transportation and evaluate the policies through comparative analyses that measure their impact on acceptance rates and overall profitability. (v) We conduct extensive computational experiments on diverse instances, analyzing the impact of the scheduled service network structure on the booking-level decision solutions. The analysis also examines how booking behavior differences (e.g., early vs. late booking tendencies), changes in the prediction confidence parameter, and environmental variations (such as smooth versus abrupt water-level fluctuation patterns) can affect the system's overall performance. The remainder of the paper is organized as follows. Section 2 presents a general overview of decision levels in freight transportation systems, with a focus on consolidation-based carriers, and reviews the literature related to acceptance and rejection decisions in such systems. Section 3 describes the problem setting in detail. Section 4 introduces the proposed methodology and provides the mathematical model formulation. Section 5 reports and discusses the experimental results. Finally, Section 6 concludes the paper by summarizing the key findings and suggesting directions for future research.

2 Decision-Making in Freight Transport: Supply Planning and Demand Control

The aim of this section is to position our study within the broader context of freight transportation planning and the different decision levels involved. Sub-section 2.1 provides an overview of these decision levels, with particular emphasis on their relevance to consolidation-based carriers. Sub-section 2.2 reviews the literature on acceptance and rejection, also referred to as demand control decisions in freight transportation.

2.1 Decision Levels in Freight Transportation

In a freight transportation system, particularly within the context of consolidation-based intermodal transport, one side consists of shippers who require the movement of freight between various origins and destinations. On the other side are carriers, particularly consolidation-based carriers for medium- and long-distance services, who operate scheduled services over a fixed network of terminals connected by predefined routes. Each service is operated using one or more resources (e.g., vessels, trucks, or trains) assigned to a fixed route with a defined origin and destination, and possibly intermediate stops, all with specified departure and arrival times. The physical and operational characteristics of these resources, such as load capacity and speed, determine the service's capacity and quality (e.g., express or standard). Effective planning is therefore required at multiple, interconnected decision levels, including tactical, booking, and operational, which together guide the carrier toward its ultimate objective of maximizing profitability. Within this multi-level decision structure, this sub-section clarifies the interplay between tactical and operational planning and explicitly positions the booking problem within this hierarchy. We frame booking-level decisions as an intermediate layer that connects long- and mid-term tactical planning with short-term operational planning. This conceptual positioning is central to our modeling approach and highlights our contribution in structuring the booking problem relative to existing tactical and operational formulations.

At the tactical level, carriers develop transportation plans that define the service network and schedule, resource utilization, manage transfer and consolidation activities at terminals, and determine preliminary demand routing. These plans are designed to ensure efficiency, profitability, and effective consolidation, while mitigating potential drawbacks of the consolidation strategy (such as increased delays or higher terminal handling costs) and maintaining the level of service quality essential for shipper satisfaction (see comprehensive reviews, e.g., SteadieSeifi et al. (2014); Elbert et al. (2020); Crainic and Rei (2024)). As illustrated in Figure 1, this planning process is based primarily on estimated regular demand, referring to shippers expected to generate business consistently throughout the upcoming season. Such expectations are grounded in established contracts and long-standing relationships with regular shippers, complemented by market forecasts that identify additional opportunities from spot shippers. The resulting plan, designed for a given time horizon (commonly referred to in the literature as the schedule length), is applied repeatedly throughout the season. It is important to note that tactical-level decisions are made before the start of the season, when demand can only be estimated at an aggregated level. In this context, acceptance and demand routing at the tactical level serve as an evaluation tool to guide the main tactical decisions. For example, to determine which potential services are most profitable and how limited resources and their capacities can be efficiently allocated to support those services while accommodating both regular and selected spot demand. These acceptance and routing decisions are not implemented at this level, as actual demand may differ from the forecasts used in tactical planning. Consequently, the implemented component of the tactical plan is comprised, in this case, of the cost-efficient services selected from the set of potential services, their schedules, and the associated resource allocations. This plan is fixed based on forecasted aggregated demand over the season. This fixed tactical configuration becomes the foundation for the next level of planning during which accept/reject decisions are made for individual requests as they arrive throughout the season.

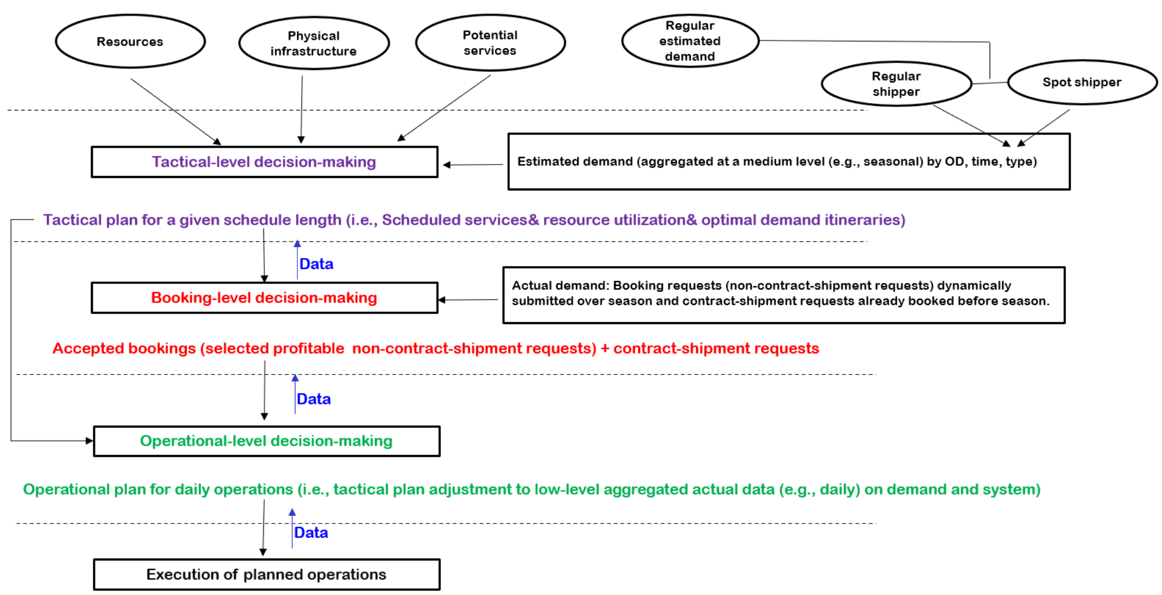

Figure 1: Multi-level decision-making framework

At the booking level, the carrier begins interacting with actual demand in the form of booking requests from non-contractual shippers arriving throughout the season. These must be considered alongside the obligation to fulfill the shipment requests of contractual shippers, whose shipments (specified in long-term agreements) will become available at predetermined times and origin–destination (O–D) terminals during the season. Even under the assumption that contractual shipments will be on time and on weight, the booking level does not rely on their preplanned itineraries from the tactical planning level. Instead, contractual shipments are re-routed at the booking level only to allow flow redistributions that can make room for newly arriving non-contractual requests and, when necessary, to permit outsourcing of part of the contractual demand. Any acceptance of a non-contract-shipment request is therefore evaluated under the constraint that all contractual commitments will be satisfied, either by using internal service capacity or through outsourcing, even though no actual allocation decisions are made at this stage. Each non-contract-shipment request is evaluated through a two-step framework introduced in this paper before an accept/reject decision is made. First, feasibility is assessed by verifying spatial and temporal compatibility within the tactical plan. To support this evaluation, we introduce the concept of *shipment–service ways*, defined as sequences of scheduled services with consistent spatial and temporal matching that enable the movement of a shipment from its origin to its destination within its availability and due dates. Second, the operational cost of routing the non-contract-shipment request along these feasible shipment–service ways, possibly including outsourcing alternatives, is estimated. Requests are accepted only if they remain profitable while ensuring that all seasonal obligations for contract-shipment requests are satisfied.

The decision to accept or reject a non-contract shipment request follows predefined policies that govern how evaluations are conducted as requests are received during the season. These policies are characterized by two main components: the *timing of decision-making* and the *informational context used*. The timing of decision-making determines whether evaluation may occur immediately upon non-contract-shipment request submission (order-based) or after a controlled delay that allows multiple non-contract-shipment requests to be assessed together (batch-based). The informational context determines whether evaluation may rely solely on current and previously accepted non-contract-shipment requests (myopic approach) or may also incorporate forecasts of future non-contract-shipment requests (lookahead approach). In this case, such forecast requests are considered if their availability date occurs before the due date of the current request, even if their own due date extends beyond it.

As indicated in Figure 1, the set of accepted non-contract-shipment requests, combined with the tactical plan, forms the foundation for operational planning, which is typically conducted on the day of execution using the most up-to-date information available (e.g., changes in shipment availability times, destination updates, or sizes). The goal is to establish implementable decisions for the detailed shipment routing (see comprehensive reviews, e.g., Meng et al. (2014); Delbart et al. (2021); Ksciuk et al. (2023)).

One should note that the booking level is positioned between tactical and operational planning within this framework. Unlike the tactical and operational levels, which mainly emphasize supply-side activities such as scheduled service network design and shipment routing, the booking level focuses exclusively on demand-side control. It determines, ahead of operational planning and execution, which non-contract-shipment requests should be accepted to maximize system profitability. No actual capacity allocation or routing is performed at this level—neither for contract nor non-contract requests. Instead, carriers conceptually estimate the net cost of potential shipment–service way assignments to shipment requests or outsourcing alternatives to inform early acceptance decisions. The need to explicitly consider contract-shipment requests arises from the fact that such requests, which require a commitment of the available service supply in the transportation system, are important to factor in when assessing the expected profitability of dynamically arriving non-contract requests.

2.2 Literature Review on Acceptance and Rejection Decisions

As discussed earlier, decision-making in freight transportation systems spans three interconnected levels, with the role of shipment routing differing at each stage. At the tactical level, shipment routing supports main tactical planning decisions such as selecting cost-efficient services and defining their schedules—referred to as the scheduled service network design. The booking level relies on this tactical plan but does not yet make routing implementation decisions; instead, it uses shipment routing as a decision-support tool to evaluate the time-space compatibility of each non-contracted shipment request with the selected service network and to estimate its expected profitability. At this level, accepted non-contract-shipment requests, together with all contract-shipment requests, are assumed to be accommodated later, but no specific resource-supported service capacity is yet assigned. The operational level then makes routing implementation decisions by allocating the capacity of available resource-supported services to shipments that are available for transportation operations.

The existing literature has primarily addressed shipment routing in the context of tactical or operational planning, with less emphasis on its role in booking-level decision-making and resource allocation mechanisms within a revenue management framework. In the latter context—referred to here as an acceptance/rejection-based resource allocation problem—shipment routing functions as a decision-support mechanism for allocating limited resource-supported services, such as transportation services on single-leg or network-based routes, with the objective of maximizing carrier revenue through optimal use of available capacity (Meng et al., 2019). The remainder of this section focuses on acceptance/rejection-based resource allocation problems, as they closely resemble the booking-level decision-making process. Both involve evaluating each incoming request and deciding whether to accept or reject it based on feasibility and profitability considerations. The key distinction lies in the role of shipment routing: in booking problems, routing is used to assess both the feasibility and profitability of a request, whereas in acceptance/rejection-based resource allocation problems, routing determines the specific shipment itinerary through the scheduled service network, and the residual capacities of those services are updated once a shipment is accepted.

The literature on the acceptance/rejection-based resource allocation problem highlights two primary strategic approaches: booking limits and bid-price controls. Booking limits define the maximum number of shipments that may be accepted and typically reflect a first-come, first-served strategy, where requests are accepted sequentially until the limit is reached. In contrast, the bid-price strategy captures the opportunity cost of allocating capacity to a current request versus reserving it for a potentially more profitable future request (Littlewood, 1972). Several studies have compared these strategies. (Kapetanović et al., 2018) analyze a bid-price-based dynamic acceptance/rejection strategy applied to a network train service. A dynamic programming formulation and deterministic approximations are used to demonstrate the potential revenue benefits of bid-price strategies compared to the traditional first-come, first-served (FCFS) strategy, particularly under high-demand situations. Similarly, (Bilegan et al., 2015) and (Wang et al., 2016) use a bid-price-based dynamic acceptance/rejection strategy for rail and intermodal freight transportation, and show advantages in comparison with FCFS. Both papers propose a probabilistic mixed-integer program on a time–space network, incorporating the probability distributions of future requests (considering volume uncertainty) in a rolling horizon framework to handle dynamic request arrivals. (Wang et al., 2017) formulate a dynamic and stochastic resource allocation problem in single-leg intermodal freight transportation as a Markov decision process, where multiple shipper arrivals with random shipment volumes are considered at each decision epoch, thereby improving capacity management and planning under uncertainty. In addition to the aforementioned strategies that highlight the importance of informational context—such as the contrast between FCFS, which bases decisions solely on current requests, and bid-price approaches, which incorporate expectations of future demand—some studies have emphasized the value of postponing decisions as a means to improve resource allocation outcomes. For instance, (Lee et al., 2009) study a dynamic and stochastic resource allocation problem in single-leg maritime transportation and propose a stochastic dynamic programming model. The paper considers both contract-shipment requests and non-contract-shipment requests and introduces flexibility by allowing the carrier to postpone contractual shipment deliveries to better manage available capacity.

Service disruptions have not been extensively addressed in the existing literature on resource allocation models. Early studies such as (Wang, 2016a,b) were among the first to introduce resource allocation problems under service disruptions and random resource capacities. However, both works focus on static settings, dealing only with single-period resource allocation problems. The dynamic extension is proposed by (Wang, 2017), who formulate the multi-period version as a dynamic programming model. Their study shows that accounting for disruption risk enhances the robustness of booking and allocation decisions. The authors consider two distinct cases: under uniform resource consumption rates, they characterize the monotonicity of the optimal solution and propose a tailored backward induction algorithm. In contrast, for general consumption rates, they demonstrate through a counterexample that the optimal solution may not be monotone, and provide an alternative backward induction algorithm to solve the model. In the context of global synchromodal transportation, (Guo et al., 2021) investigate a dynamic and stochastic shipment matching problem, in which a platform makes real-time decisions to accept or reject shipment requests and assign them to multimodal services (ship, barge, rail, truck), while facing uncertainties in both demand and travel times. The problem is formulated using a hybrid stochastic model, combining sample average approximation to represent the stochastic information of future requests and chance-constrained programming to handle uncertain travel times. The model is embedded within a rolling horizon framework, where decisions are made over time as new information becomes available. Their model explicitly incorporates service disruptions and infeasible transshipments into the planning process, with the goal of maximizing total profit while minimizing infeasibility and delivery delays. In intermodal barge transportation systems, (Cui et al., 2024) adopt a bid-price-based strategy and introduce a mixed-integer programming model on a time–space network that incorporates probability distributions of future shipment requests. The model is also embedded within a rolling horizon framework to evaluate whether to accept or reject a current non-contract-shipment request in order to maximize expected revenue. In addition, the model allows rerouting of accepted but not yet delivered shipments to enhance capacity utilization. Importantly, the study addresses the impact of environmental disruptions—such as weather events and fluctuating water levels—by including both rerouting options and penalty costs for outsourced demand resulting from capacity shortfalls.

Although acceptance/rejection-based resource allocation models have been widely studied, they are primarily situated at the operational planning level. By contrast, this paper explicitly addresses booking-level decision-making in intermodal transportation—a topic that remains largely unexplored. To the best of our knowledge, this is among the first studies to address this problem with a methodology that is neither dynamic programming–based nor network flow–based. Focusing on intermodal barge transportation, the proposed framework further accounts for environmental disruptions, particularly water-level variations, which directly affect vessel capacities, navigability, and feasibility of shipment routing solutions.

3 Problem Setting

The problem setting is inspired by the booking challenges in barge transportation systems, which serve as a critical mid- and long-haul component of the freight transportation network. These systems operate between major regions connected by an inland waterway network—a system of navigable rivers, canals, and channels enabling vessel and barge movement between ports equipped with infrastructure for barge operations, including loading and unloading facilities, as well as warehouses and storage yards for temporary freight storage. While the setting is general and applicable to various freight networks, it is illustrated here in the context of barge transportation.

Within this context, the booking system involves multiple shippers—including producers, wholesalers, and retailers—who request the transportation of their shipments across the network. These shipment requests are categorized as either contractual or non-contractual. Contractual shipments originate from long-term agreements and are assumed to be on time and on weight, in accordance with the agreed contractual terms. In contrast, non-contractual requests have no pre-established agreement and arrive dynamically. Each shipment request—regardless of its type—is characterized by an origin–destination terminal pair, an availability time, a due time, a weight, and an economic value for the carrier. The economic value depends on the requested delivery time (generally express or standard) as well as the timing of the booking submission (early or late), with express delivery and late bookings typically associated with higher values. On the other hand, the booking system involves the availability of a transportation network composed of scheduled services operated by barges. These services are used to construct shipment–service ways, defined as sequences of temporally and spatially connected barge movements between inland terminals. Each shipment–service way is described by its origin (departure terminal of the first service), destination (arrival terminal of the last service), availability time (departure time of the first service), travel time (including transit and terminal processing durations), capacity (the minimum capacity among the path of services along the way), and total cost (covering both transportation and handling costs). A shipment–service way is considered item-feasible if it matches the shipment request's origin and destination terminals, starts no earlier than its availability time, and finishes no later than its due time.

Building on this system description, the booking problem is defined over a finite booking length (e.g., several days), within which booking decisions are made at specific decision points determined by the adopted policy. Under a request-based policy, each arriving request triggers an immediate decision, so that every request constitutes a decision point. Under a delay-based policy, the booking length is partitioned into fixed time intervals (e.g., half-day periods), during which multiple non-contractual requests accumulate and are jointly evaluated at the end of each interval. Although request-based decision-making may be applied in high-frequency transportation contexts (e.g., hourly services), where shipment consolidation is not a primary operational requirement, it is generally less suitable for the barge transportation systems considered in this study, which rely on consolidation-based operations and operate at lower service frequency.

The composition of the shipment request set and the associated service ways considered at each decision point depends on the adopted booking policy, which is defined by the informational context available. Under a myopic policy, only current and past information is used: the shipment request set includes (i) all contractual shipment requests, which must be satisfied due to contractual obligations, and previously accepted non-contractual shipment requests, and (ii) the current batch of available non-contractual shipment requests. Under a lookahead policy, the informational scope is extended to also include (iii) predicted non-contractual shipment requests that are expected to become available within the time window defined by the current batch—namely, the interval between the earliest availability time and the latest due time among current items.

Given a shipment request and its feasible service ways at each decision point, the booking system uses an optimization model to evaluate the profitability of non-contractual shipment requests. This evaluation accounts for transportation costs (i.e., assignment to feasible service ways) and potential outsourcing costs, under the constraint that all obligated shipments—including contractual shipment requests and previously accepted non-contractual shipment requests—must be fulfilled. To provide flexibility and improve utilization of service-way capacities, the system permits late pickups (assigning items to services departing after their availability time) and early deliveries (arrivals before their due time). The objective at each decision point is to maximize profit by fully serving obligated shipments while selectively accepting only profitable non-contractual requests.

This setting is modeled as a bin-packing-based framework, where shipment requests are represented as items to be packed into item-feasible service ways that act as bins with limited capacity. Outsourcing alternatives are also modeled as auxiliary bins, incurring higher costs. Our methodology builds upon the classical bin-packing framework, extending it to address the specific requirements of the booking problem in consolidation-based freight transportation. Unlike most bin-packing formulations that operate over a single decision period, our models consider a multi-period planning horizon to capture the temporal dynamics inherent in real-world freight operations. To the best of our knowledge, time-window constraints—critical for aligning with shipper delivery requirements—are rarely incorporated in the bin-packing literature. Our work extends the study of Crainic et al. (2021), which applied a multi-period bin-packing approach in a single-corridor setting, by generalizing it to a network context. Furthermore, the framework explicitly incorporates item profits and distinguishes between items, providing a richer and more realistic representation of demand heterogeneity. The primary objective of the assignment model is to evaluate profitability rather than determine the physical arrangement of items—a focus more relevant to operational-level planning. In our case, the item "size" is approximated solely by weight, which is consistent with the operational reality of inland waterway transport systems where service capacity is strongly influenced by water levels. Specifically, a service-way is considered feasible only if the total weight of the assigned items remains within the service's adjusted capacity, which depends on the water level due to constraints such as vessel draught limitations and grounding risk. Therefore, we do not adopt a full three-dimensional bin-packing formulation. This abstraction allows us to preserve computational tractability while still capturing the critical interaction between weight-based capacity and environmental limitations in barge transportation networks, directly influencing the profitability of accept/reject decisions.

4 Methodology

Based on the problem definition introduced in Section 3, in this section we first define the notation and the parameters used and then present the mathematical formulation developed for the problem.

4.1 Notation

Shipment request. Shipment requests originate from two categories of shippers: contractual shippers, who hold long-term agreements guaranteeing transportation services, and non-contractual shippers, who do not have such agreements and whose requests are subject to the carrier's acceptance decision. In this study, *compulsory items* ($\mathcal{I}^C$) comprise all not-yet-moved requests, including those from contractual shippers and previously accepted non-contractual requests, all of which must be fully satisfied. The term *non-compulsory items* ($\mathcal{I}^{NC}$) refers to current requests from non-contractual shippers awaiting an acceptance or rejection decision, while *predicted non-compulsory items* ($\mathcal{I}^{PNC}$) denote forecasted future requests from non-contractual shippers. The complete set of items is thus $\mathcal{I} = \mathcal{I}^C \cup \mathcal{I}^{NC} \cup \mathcal{I}^{PNC}$.

Each item $i \in \mathcal{I}$, whether compulsory ($i \in \mathcal{I}^C$) or non-compulsory ($i \in \mathcal{I}^{NC} \cup \mathcal{I}^{PNC}$), is described by three categories of attributes: *physical*, *temporal–spatial*, and *economic*. The *physical* attribute includes the weight w_i of the item. The *temporal–spatial* attributes capture the availability time $\underline{t}_i$, representing when the item becomes available for transport at the origin terminal $o(i)$; the delivery due time $\bar{t}_i$ at the destination terminal $d(i)$, specifying the preferred arrival time at the destination; and the reservation time r_i, indicating when the transportation request is submitted. The *economic* attributes are expressed through the revenue ϕ_i associated with transporting item i, which depends on the requested delivery time: either standard or express, with express requests being charged higher fares. Items are further classified as early or late based on the anticipation window $\omega_i = \underline{t}_i - r_i$, relative to a predefined threshold τ: if $\omega_i \geq \tau$, the item is categorized as early; otherwise, it is considered late. Early bookings benefit from base fares, while late bookings may incur surcharges due to increased operational complexity. Thus, each item generates revenue based on its requested delivery time and anticipation window. All fares are determined in advance and remain fixed, with no negotiation process required.

Shipment–Service ways. Let Σ denote the set of all scheduled services. A shipment–

service way is an ordered sequence of one or more services from Σ that connects an origin-destination (OD) terminal pair, under a given schedule. Let $\mathcal{J}$ be the set of all such sequences in the network. The services within any sequence must be temporally and spatially matched (the arrival of each service precedes the departure of the next one, and consecutive terminals are consistent). For each $j \in \mathcal{J}$, let $\mathrm{Serv}(j) = \big(\sigma_1(j), \sigma_2(j), \ldots, \sigma_{|\mathrm{Serv}(j)|}(j)\big)$, with $\sigma_k(j) \in \Sigma\ \forall k$, denote the ordered sequence of scheduled services composing j. The origin and destination terminals of j are given by the first and last services in the sequence, respectively, denoted $o(j)$ and $d(j)$. Each scheduled service $\sigma \in \mathrm{Serv}(j)$ has a nominal capacity, an availability time, a travel time, and cost components (transportation and handling such as loading/unloading). The aggregated attributes of j are defined as follows: its *physical* capacity W_j equals the minimum capacity across all $\sigma \in \mathrm{Serv}(j)$. Since water level directly affects the draught and thus the effective load capacity of vessels, each W_σ is a function of the water level $l \in \mathcal{L}$, where $\mathcal{L}$ denotes the set of considered water levels (e.g., discretized levels over the planning horizon). Consequently, the capacity of j is also water-level-dependent and is expressed as: $W_j(l) = \min_{\sigma \in \mathrm{Serv}(j)} W_\sigma(l),$; the *temporal–spatial* attributes are given by t_j, the availability time of the first service, and α_j, the sum of the travel and terminal times of all scheduled services in the sequence; and its *economic* attribute C_j is the sum of transportation and handling costs of the scheduled services used by j. An additional cost component β_i is considered to capture the cost of using an outsourced service to transport the item $i \in \mathcal{I}$. Given an item $i \in \mathcal{I}$, we denote by $\mathcal{J}_i \subseteq \mathcal{J}$ the set of *item-feasible service ways* for i (i.e., service ways that match the origin and destination of item i and whose departure and arrival times fall within the item's availability and due time window): $\mathcal{J}_i = \big\{j \in \mathcal{J} : \mathrm{o}(j) = o(i),\ \mathrm{d}(j) = d(i),\ t_j \geq \underline{t}_i,\ t_j + \alpha_j \leq \bar{t}_i\big\}$.

4.2 Mathematical formulation

We develop three models for the booking problem. The first is the *Myopic Bin-Packing Booking Problem*(MBBP), in which acceptance and rejection decisions are based solely on past and current booking information. This model serves as a baseline for comparison with the *Lookahead Bin-Packing Booking Problem* (LBBP), which incorporates predictions of future booking requests to support decision-making. The third model, LBBP-WL, extends the LBBP by explicitly incorporating predictions of water-level conditions and their impact on the capacity of item-feasible service ways. To formulate the three models, we first introduce the following binary variables:

- $y_{ij} = \begin{cases} 1, & \text{if item-feasible service way } j \in \mathcal{J}_i \text{ is selected for item } i \in \mathcal{I}, \\ 0, & \text{otherwise;} \end{cases}$

- $x_{ij} = \begin{cases} 1, & \text{if item } i \in \mathcal{I} \text{ is assigned to item-feasible service way } j \in \mathcal{J}_i, \\ 0, & \text{otherwise;} \end{cases}$

- $u_i = \begin{cases} 1, & \text{if item } i \in \mathcal{I} \text{ is assigned to outsourcing service,} \\ 0, & \text{otherwise.} \end{cases}$

The MBBP model. The MBBP can be formulated as follows:

$$\max \sum_{i \in \mathcal{I}} \phi_i \left[\sum_{j \in \mathcal{J}_i} x_{ij} + u_i \right] - \sum_{i \in \mathcal{I}} \beta_i \, u_i - \sum_{j \in \mathcal{J}_i} C_j \sum_{i \in \mathcal{I}} x_{ij} \tag{1}$$

subject to:

$$\sum_{j \in \mathcal{J}_i} x_{ij} + u_i = 1, \quad \forall i \in \mathcal{I}^C \tag{2}$$

$$\sum_{j \in \mathcal{J}_i} x_{ij} + u_i \leq 1, \quad \forall i \in \mathcal{I}^{NC} \tag{3}$$

$$\sum_{i \in \mathcal{I}} w_i \, x_{ij} \leq W_j(l) \, y_{ij}, \quad \forall j \in \mathcal{J}_i, \quad \forall i \in I \tag{4}$$

$$\sum_{i \in I} \sum_{\substack{j \in J_i: \\ \sigma \in \text{Serv}(j)}} w_i x_{ij} \leq W_\sigma(l), \quad \forall \sigma \in \Sigma \tag{5}$$

$$y_{ij}, x_{ij}, u_i \in \{0, 1\}, \quad \forall i \in \mathcal{I}, j \in \mathcal{J}_i \tag{6}$$

The objective function (1) maximizes the profit by adding the revenues from items assigned to item-feasible service ways and outsourcing paths, and subtracting the corresponding costs for using those service ways or outsourcing paths. Constraint (2) ensures that each compulsory item is either assigned to a item-feasible service way or sent to an alternative path, while constraint (3) states that non-compulsory items may or may not be assigned. Constraint (4) enforces the capacity limits of each item-feasible service way, explicitly considering the corresponding water level. Constraint (5) ensures that the total load assigned to all service ways sharing the same scheduled service does not exceed that service's physical capacity, thereby linking capacity usage across different item-feasible service ways that include the same scheduled services. Finally, constraint (6) defines all decision variables as binary.

The LBBP model. Each item $i \in \mathcal{I}^{PNC}$ is characterized by a point-estimated weight, denoted by $\hat{w}_i$. Additionally, a confidence parameter $\theta_i \in [0,1]$ is associated with each predicted item i, representing the reliability of the prediction for that item. This parameter scales the revenue term of predicted items in the objective function, allowing the model to prioritize or de-prioritize these items based on their prediction accuracy.

The LBBP can be formulated as follows:

$$\begin{aligned} \max \quad & \sum_{i \in \mathcal{I}^C \cup \mathcal{I}^{NC}} \phi_i \left[\sum_{j \in \mathcal{J}_i} x_{ij} + u_i \right] + \sum_{i \in \mathcal{I}^{PNC}} \theta_i \, \phi_i \left[\sum_{j \in \mathcal{J}_i} x_{ij} + u_i \right] \\ & - \sum_{i \in \mathcal{I}^C \cup \mathcal{I}^{NC} \cup \mathcal{I}^{PNC}} \beta_i \, u_i - \sum_{j \in \mathcal{J}_i} C_j \sum_{i \in \mathcal{I}^C \cup \mathcal{I}^{NC} \cup \mathcal{I}^{PNC}} x_{ij} \end{aligned} \tag{7}$$

subject to:

$$\sum_{j \in \mathcal{J}_i} x_{ij} + u_i = 1, \quad \forall i \in \mathcal{I}^C \tag{8}$$

$$\sum_{j \in \mathcal{J}_i} x_{ij} + u_i \leq 1, \quad \forall i \in \mathcal{I}^{NC} \cup \mathcal{I}^{PNC} \tag{9}$$

$$\sum_{i \in \mathcal{I}^C \cup \mathcal{I}^{NC}} w_i \, x_{ij} \; + \sum_{i \in \mathcal{I}^{PNC}} \hat{w}_i \, x_{ij} \; \leq \; W_j(l) \, y_{ij}, \quad \forall j \in \mathcal{J}_i, \quad \forall i \in I \tag{10}$$

$$\sum_{i \in \mathcal{I}^C \cup \mathcal{I}^{NC}} \sum_{\substack{j \in \mathcal{J}_i: \\ \sigma \in \mathrm{Serv}(j)}} w_i \, x_{ij} \; + \sum_{i \in \mathcal{I}^{PNC}} \sum_{\substack{j \in \mathcal{J}_i: \\ \sigma \in \mathrm{Serv}(j)}} \hat{w}_i x_{ij} \; \leq \; W_\sigma(l), \quad \forall \sigma \in \Sigma. \tag{11}$$

$$y_{ij}, x_{ij}, u_i \in \{0,1\}, \quad \forall i \in \mathcal{I}, j \in \mathcal{J}_i \tag{12}$$

The objective function (7) maximizes profit by adding the revenues from items assigned to item-feasible service ways and outsourcing paths, with predicted items scaled by their confidence parameter, and subtracting the corresponding outsourcing path and item-feasible service way assignment costs. Constraint (8) ensures that each compulsory item is either assigned to an item-feasible service way or sent to theoutsourcing path. Constraint (9) allows non-compulsory and predicted items to be optionally assigned. Constraints (10) and (11) enforce individual item-feasible service way capacities and shared scheduled-service capacities, respectively. Finally, constraint (12) defines all decision variables as binary.

The LBBP-WL model. The LBBP-WL extends the LBBP by replacing the fixed service capacities with point-estimated capacities. Specifically, the parameters W_j and W_σ in constraints (10)–(11) are replaced by their point-estimated counterparts $\hat{W}_j(l)$ and $\hat{W}_\sigma(l)$, which reflect the variations in capacity induced by the water level l.

5 Experimental Results and Analysis

In this section, we present a series of computational experiments designed to assess the performance of the proposed decision-making models and to evaluate the influence of diverse operational and environmental conditions on booking outcomes. The analyses aim to provide both a quantitative comparison of alternative policies and qualitative insights into their managerial implications. The organization of this section is as follows: in Section 5.1, we present the characteristics of the instances generated for the computational experiments—based on realistic cases—and the rolling-horizon booking simulation framework. We examine the behavior of the model in terms of computational time in Section 5.2. In Section 5.3, we define the performance indicators required for analysing the computational results. Five experimental settings are then examined in Subsections 5.3. The first compares myopic, lookahead, and lookahead-with-water-level-variation policies to assess the value of anticipating future item and environmental changes. The second explores how booking behavior heterogeneity, including early and late booking tendencies, affects acceptance rates and profit. The third examines the impact of tactical planning choices—such as service selection and fleet composition—on downstream booking efficiency and profitability, highlighting the interdependence between the planning levels. The fourth evaluates the system's sensitivity to the prediction confidence parameter, offering insights into its robustness against forecast accuracy. Finally, the fifth investigates the impact of smooth versus abrupt water level variations on booking decisions.

All implementations are conducted using the Pyomo software package and the Gurobi solver on a machine equipped with an Intel(R) Xeon(R) CPU E5-2630 v4 @ 2.20GHz and 256 GB of memory.

5.1 Test instances

Shipment-service way generation. We define a set of scheduled services inspired by the structure introduced in Payami et al. (2025b), following a one-week horizon divided into 14 half-day periods, consistent with observed practices of morning and afternoon departures. We apply the procedure detailed in Payami et al. (2025b), of which we just recall the main lines. Each scheduled service corresponds to a tactical planning decision and is characterized by a fixed route connecting a predefined origin–destination terminal pair, an availability time selected from the discrete set of 14 half-day periods $\{0, 1, \ldots, 13\}$, and a travel duration selected from $\{2, 3, 4\}$, corresponding to a minimum of one day and a maximum of two days. Each scheduled service incurs a fixed operational cost of 5 units, while an outsourcing option is available at a fixed cost of 8 units. We consider two vessel classes: small vessels with a nominal capacity of 35 units and large vessels with a nominal capacity of 50 units. We consider three distinct scheduled service networks, each representing a tactical plan obtained in advance through a prior tactical planning process and reflecting different assumptions about navigability at the time of planning. In the VC-U (Uniform) setting, services are evenly distributed across vessel classes, representing a balanced fleet composition with no structural bias toward small or large vessels. The VC-SF (Split Fleet) setting introduces a heterogeneous composition in which half of the services are operated by small vessels and the other half by large vessels, allowing the model to benefit from service flexibility across vessel classes. Finally, the VC-R (Restricted) setting represents a scenario in which reduced navigability is anticipated at the tactical planning stage, reflected in the allocation of 75% of services to small vessels and only 25% to large vessels, resulting in an asymmetric and capacity-limited network. Based on each of these service sets, we construct shipment–service ways by combining temporally and spatially feasible sequences of scheduled services. As a result, both VC-U and VC-R yield 36 distinct service ways, while VC-SF produces 72, due to the larger number of feasible combinations arising from the mixed fleet composition.

Item generation. We generate a heterogeneous shipment request set composed of compulsory and non-compulsory items. Compulsory items ($\mathcal{I}^C$) represent contractual shipments that must be served, while non-compulsory items ($\mathcal{I}^{NC}$) may be accepted or rejected depending on their contribution to profitability. Each item i is defined by its availability time, due time, weight, and fare. Availability times are selected from the discrete set of 14 half-day periods. Due times are defined by specifying a flexible delivery window relative to the item's availability time, corresponding to approximately one to five days after availability, thereby introducing variability in delivery time tightness. Item weights are set according to a uniform distribution $w_i \sim U(5, 15)$, and fares are calculated as $\phi_i = m \cdot b$, where $b \in \{1, \ldots, 9\}$ is a base fare obtained by discretizing a linear mapping of item weight onto nine tariff levels, ensuring that heavier shipments are systematically associated with higher tariff levels, and $m \in \{1.0, 1.5, 2.0, 2.5\}$ is a multiplier capturing the service/booking combination (standard-early, express-early, standard-late, express-late). Both demand sets include the same 60 compulsory items. The difference lies in the composition of non-compulsory items. The low-fare *LF-480* set contains 480 non-compulsory requests, of which approximately 34% belong to the high-fare category (express-late), with the remainder distributed among the other service/booking combinations. The high-fare *HF-528* set contains 528 non-compulsory requests, with a substantially higher proportion-about 65%-in the high-fare category, the remainder again covering the other combinations. Predicted items follow the same fare generation rules as observed items, while their weights are drawn from a uniform distribution $w_i \sim U(0, 15)$. This structure ensures that both compulsory and non-compulsory item patterns are consistent across experiments, while allowing variation in total volume, fare composition, and the share of high-value shipments.

Rolling-horizon booking simulation. We evaluate booking policies using a rolling-horizon approach, where at each decision step h, the model observes the fixed set of compulsory items $\mathcal{I}^C$ and a batch of non-compulsory arrivals $\mathcal{I}_h^{NC}$ of randomized size $|\mathcal{I}_h^{NC}| \in [120, 180]$. Accepted non-compulsory items are carried forward as compulsory in the next step, while rejected items are removed from the system and are no longer available for future consideration. This observe–optimize–roll procedure is repeated iteratively, with full re-optimization at each step, enabling greater flexibility in evaluating the profitability of dynamically arriving shipment requests. In the lookahead variant, we simulate forecasted demand by introducing a global set of predicted non-compulsory items, denoted $\mathcal{I}^{PNC}$, which is generated a priori at the start of the simulation. At each decision step h, a relevant subset is extracted based on the time window defined by the currently observed non-compulsory batch—specifically, the interval between the minimum availability time and the maximum due time in $\mathcal{I}_h^{NC}$. These predicted (not-yet-submitted) items contribute in expectation, weighted by a confidence coefficient $\theta \in [0, 1]$; we set $\theta = 0.5$ by default to remain neutral (higher θ trusts forecasts more, lower θ is conservative). Water level fluctuations are modeled using a discrete probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where each realization $\omega \in \Omega$ corresponds to a specific water level scenario. For each service σ, a random capacity reduction factor $\gamma_\sigma(\omega) \in (0, 1]$ is defined, representing the impact of water levels on that specific service under realization ω. The realized capacity of service σ under scenario ω is given by: $W_\sigma(\omega) = \gamma_\sigma(\omega) \cdot W_\sigma^{\text{base}}$, where W_σ^{base} denotes the nominal capacity. The expected capacity of service σ is thus: $\mathbb{E}[W_\sigma(l)] = \sum_{\omega \in \Omega} \mathbb{P}(\omega) \cdot \gamma_\sigma(\omega) \cdot W_\sigma^{\text{base}}$. Since a shipment–service way $j \in \mathcal{J}_i$ consists of a sequence of services $\sigma \in \text{Serv}(j)$, its expected capacity is determined by the most constraining component: $\mathbb{E}[W_j(l)] = \min_{\sigma \in \text{Serv}(j)} \mathbb{E}[W_\sigma(l)]$.We consider two distinct water-level fluctuation scenarios, each inducing different magnitudes of capacity reduction. In the smooth scenario, water-level variations remain moderate, leading to an average capacity reduction of approximately 9%, meaning that the expected capacity remains close to 91% of its nominal value (with the lower bound reaching 75%). In contrast, the abrupt scenario reflects more extreme fluctuations in water levels, causing a much sharper average reduction of approximately 42%, with the expected service capacity dropping to about 58% of its nominal value (and potentially falling as low as 35%). These contrasting scenarios allow us to assess the model's sensitivity to both mild and severe environmental disruptions.

5.2 Model performance

Tables 1 to 3 summarize the computational characteristics of the three models—MBBP, LBBP, and LBBP-WL—across different network topologies (VC-U, VC-SF, and VC-R) and two shipment-request configurations (LF-480 and HF-528). For each model and instance, the tables report the average and standard deviation of (i) the number of decision variables, (ii) the number of constraints, and (iii) the total solution time in seconds. These metrics assess the scalability and computational efficiency of each model under varying problem sizes and network complexities. Tables 1–3 reveal consistent patterns across the three models—MBBP (myopic), LBBP (lookahead), and LBBP-WL (lookahead with water-level variation)—reflecting both the structural characteristics of the formulations and the computational effects of the rolling-horizon approach. Increasing

Table 1: MBBP performance with shipment-request/Item structure

Path	Item	# of DV		# of constraints		Time (s)	
		Avg	Std	Avg	Std	Avg	Std
VC-U	LF-480	9315.6	1976.19	358.8	53.41	22.27	5.82
	HF-528	10447.8	2756.26	389.4	74.49	72.68	11.61
VC-SF	LF-480	20176.2	5177.49	491.4	70.92	95.12	33.16
	HF-528	21811.4	6160.04	513.8	84.38	120.56	21.03
VC-R	LF-480	9382.2	2008.18	360.6	54.27	50.35	12.26
	HF-528	10433	2766.84	389	74.77	105.77	14.29

the shipment-request set size from LF-480 to HF-528 leads to a clear expansion in problem size, with higher numbers of decision variables and constraints, and, in most cases, longer average solution times. This is expected because larger request sets create more item–service-way combinations, thereby increasing the number of binary selection decisions. The effect is particularly pronounced in the VC-SF topology, which consistently yields the largest problem sizes and often the longest solution times; this topology offers more shipment–service-way options, resulting in a denser assignment structure. Comparing models, LBBP consistently produces the largest formulations because it accounts for both revealed and predicted requests at each decision stage. Incorporating predictions increases the number of possible alternatives for assigning requests to services and adds more capacity-related constraints. MBBP, by contrast, works only with revealed requests and thus has fewer variables and constraints, resulting in smaller problem sizes and generally shorter solve times. LBBP-WL lies between these two in size, as the introduction of water-level-dependent capacity constraints removes some assignment possibilities when low-water scenarios occur, effectively reducing available service capacity. Interestingly,

Table 2: LBBP performance with shipment-request/Item structure

Path	Item	# of DV		# of constraints		Time (s)	
		Avg	Std	Avg	Std	Avg	Std
VC-U	LF-480	13777.8	3074.96	479.4	83.10	572.47	217.24
	HF-528	15183.8	3506.28	517.4	94.76	1073.69	262.404
VC-SF	LF-480	28834	6583.57	610	90.18	1526.39	269.95
	HF-528	31213.8	7514.14	642.6	102.93	1027.23	183.06
VC-R	LF-480	13763	3067.32	479	82.90	526.97	123.38
	HF-528	15176.4	3504.4	517.2	94.71	346.98	64.79

Table 3: LBBP-WL performance with shipment-request/Item structure

Path	Item	# of DV		# of constraints		Time (s)	
		Avg	Std	Avg	Std	Avg	Std
VC-U	LF-480	13037.8	2835.94	459.4	76.64	121.163	19.49
	HF-528	14887.8	335.60	509.4	90.15	155.40	39.92
VC-SF	LF-480	26556.4	5880.78	578.8	80.55	1549.70	203.38
	HF-528	29534.8	6650.33	619.6	91.10	1343.25	237.31
VC-R	LF-480	13074.8	2844.50	460.4	76.87	133.62	26.25
	HF-528	14880.4	338.56	509.2	90.23	114.32	23.04

larger problem size does not always translate into proportionally higher solution times. A key reason lies in the nature of the bin packing problems that are solved iteratively within the rolling-horizon procedure. The computational difficulty of bin packing models often stems not from the number of items or bins alone, but from the degree of symmetry in the solution space—namely, the number of equivalent solutions (in terms of objective value) that differ only in their decision variables (i.e., the specific item-to-bin assignments). When there is a large number of small items relative to bin capacity (as in the lookahead models, which incorporate predicted items), and when bins are highly homogeneous—such as in VC-SF or VC-R configurations where bins share identical or nearly identical capacities—the number of equivalent item-to-bin assignments increases substantially. This symmetry can make some instances harder to solve, even when the nominal problem size is smaller. This behaviour aligns with what we observe in Tables 2 and 3, where LF-480 requires longer solution times than HF-528 despite involving fewer items. Overall, the computational results align with expectations: LBBP is the largest and often the slowest due to its predictive nature, MBBP is the smallest and generally the fastest, and LBBP-WL occupies a middle ground—except in cases where its additional constraints significantly accelerate convergence by tightening the feasible region.

5.3 Rolling Horizon Experiment

5.3.1 Performance metrics

Before presenting the performance metrics, we first recall that the booking process is dynamic: demand arrives over a finite booking length, and booking decisions must be made at specific points in time. To handle this sequential structure, the proposed bin packing models are solved iteratively within a rolling-horizon procedure. In our setting, the booking length is divided into five decision stages to reflect the sequential nature of booking arrivals and to capture the operational reality that barge services typically experience a small number of well-defined booking waves over a weekly cycle, ranging from early to late requests. Using five stages therefore provides a realistic level of temporal granularity—sufficiently fine to capture booking dynamics while remaining consistent with operational practice—and offers a consistent basis for comparing the performance of alternative booking policies. The following metrics are therefore used to evaluate system performance across the five decision stages: **Total profit:** The overall profit obtained by the carrier calculated as the total revenue minus the total costs. This metric reflects the system's ability to maximize economic efficiency while limiting the use of costly alternatives (i.e., outsourcing paths). **Acceptance rate:** The ratio of all accepted non-compulsory items to the total number of non-compulsory items requested throughout the horizon. **Total cost:** The cumulative cost incurred over the horizon, comprising *way cost* (the total cost associated with assigning items to shipment-service ways) and *path cost* (the total cost associated with routing items via outsourcing paths).

5.3.2 Impact of Anticipative Information on Booking System Performance

In this subsection, we assess the use of the three proposed models within the context of the rolling-horizon procedure: a myopic model (MBBP), a lookahead model (LBBP), and a lookahead model with water-level variations (LBBP-WL). The first model relies solely on realized information available at each decision point, whereas the second incorporates anticipative contextual information, such as expected future requests. The third model extends this anticipative framework by additionally accounting for random water-level variations that reduce available capacities, thereby capturing a more constrained and operationally realistic environment. To facilitate interpretation, the experiment is organized into two distinct settings. The first setting corresponds to cases in which water-level variability is not considered; in this environment, MBBP and LBBP operate under identical capacity assumptions and can thus be directly compared in terms of decision quality and efficiency. The second setting explicitly incorporates water-level–induced capacity reductions. In this more restrictive environment, LBBP-WL is not evaluated as a competing policy against MBBP or LBBP, but is instead used to quantify how water-level variability alters booking decisions and overall system performance relative to a setting in which such variations are ignored.

Table 4 summarizes the key performance indicators (KPIs)—including total profit, way cost, path cost, and acceptance rate—highlighting how different levels of anticipative information influence booking performance across these distinct informational and environmental contexts. Table 4 shows that the average profit in LBBP is notably higher than in MBBP (1580.66 vs. 1369.83). This improvement results from the anticipative nature of LBBP: by incorporating predictions of future requests, the model strategically accepts fewer current non-compulsory shipments (as reflected in its lower acceptance rate, 0.49 vs. 0.51 in MBBP) to preserve capacity for potentially more profitable future demand. Comparing MBBP with LBBP-WL reveals an even larger drop in acceptance rate (0.51 vs. 0.44), which reflects a combination of anticipative demand forecasting and more restrictive operational conditions. Similar to LBBP, the LBBP-WL model incorporates forecasts of future shipments, encouraging conservative acceptance decisions to maintain future flexibility. In addition, the updated water-level information available to LBBP-WL allows the model to anticipate reduced future capacity under low-water conditions. Consequently, it rejects more current requests to avoid future infeasibility, which directly contributes to its lower average profit relative to MBBP (1301.55 vs. 1369.83).

Table 4: KPI comparison for MBBP, LBBP, and LBBP-WL

KPI	MBBP		LBBP		LBBP-WL	
	Avg	Std	Avg	Std	Avg	Std
Profit	1369.83	556.45	1580.66	600.25	1301.55	545.49
Way.Cost	831.16	260.99	968.66	277.82	458.66	141.06
Path.Cost	479.20	424.80	797.06	550.177	1390.93	579.96
AC	0.51	0.07	0.49	0.09	0.44	0.09

The comparison between LBBP and LBBP-WL highlights the impact of ignoring water-level variability and the operational consequences that emerge when capacity reductions are taken into account. In LBBP, service-way capacities remain fixed and optimistic, allowing the model to consolidate more shipments and rely more heavily on regular services. From a cost perspective, this results in the highest way cost (968.66), as the model can select more regular services to handle both current and predicted demand. In contrast, LBBP-WL updates service-way capacities based on anticipated low-water conditions, leading the model to expect tighter capacity in upcoming stages. As regular-service capacity becomes limited, the system must rely more frequently on outsourced services to serve already-accepted and contractual shipments. This shift produces the lowest way cost (458.66), due to reduced feasible capacity for regular services, but simultaneously results in a higher path cost because of the increased dependence on outsourcing paths Overall, all three models remain profitable, but their outcomes reflect the informational and environmental conditions under which they operate. LBBP yields the highest profit because anticipative demand information enables it to make more selective acceptance decisions and reserve capacity for potentially more profitable future requests. MBBP, which relies solely on realized information, accepts a larger share of current requests but at the expense of lower overall profit. In contrast, the reduced profit observed in LBBP-WL is not due to weaker decision-making, but results from the more restrictive capacity conditions imposed by anticipated low-water levels, which limit feasible acceptance opportunities and increase reliance on expensive outsourced services.

5.3.3 Item impact

To isolate the effect of demand composition on model performance, we analyze how the fare structure of the shipment-request set influences booking outcomes. Specifically, we distinguish between two input categories: LF-480 (low-fare–dominant) and HF-528 (high-fare–dominant). Figure 2a and 2b provide a visual comparison of the total number of accepted items and the resulting profits across the three models for each input type.

This comparison helps assess the sensitivity of each booking policy to variations in revenue potential, and highlights how predictive or capacity-constrained models adjust their decisions based on expected profitability. Across all models, the fare composition of the input set has a strong and systematic influence on booking outcomes. When the request set is dominated by low-fare items (LF-480), accepted volumes are markedly lower and profits are more limited, with median acceptance ranging from approximately 30 items in LBBP-WL to about 40 in MBBP and LBBP, and median profits between roughly 900 and 1300. Under high-fare–dominant inputs (HF-528), both accepted items and profits increase substantially: median acceptance rises to around 55–60 in MBBP, 60–65 in LBBP, and remains near 60 in LBBP-WL, while median profits improve to roughly 1800 in MBBP, exceed 2000 in LBBP (with maxima approaching 2800), and remain close to 1800 in LBBP-WL. The improvement from LF to HF corresponds to a gain of about 15–25 accepted items and 700–900 in profit across models. Prediction of future demand (LBBP) consistently increases profit relative to the myopic case, with the largest gains observed in HF-528. Incorporating water-level-induced capacity reduction (LBBP-WL) reduces acceptance and profit in LF-480 more sharply, while having a more moderate effect under HF-528, where the high unit revenue enables the model to sustain profitability despite fewer accepted bookings.

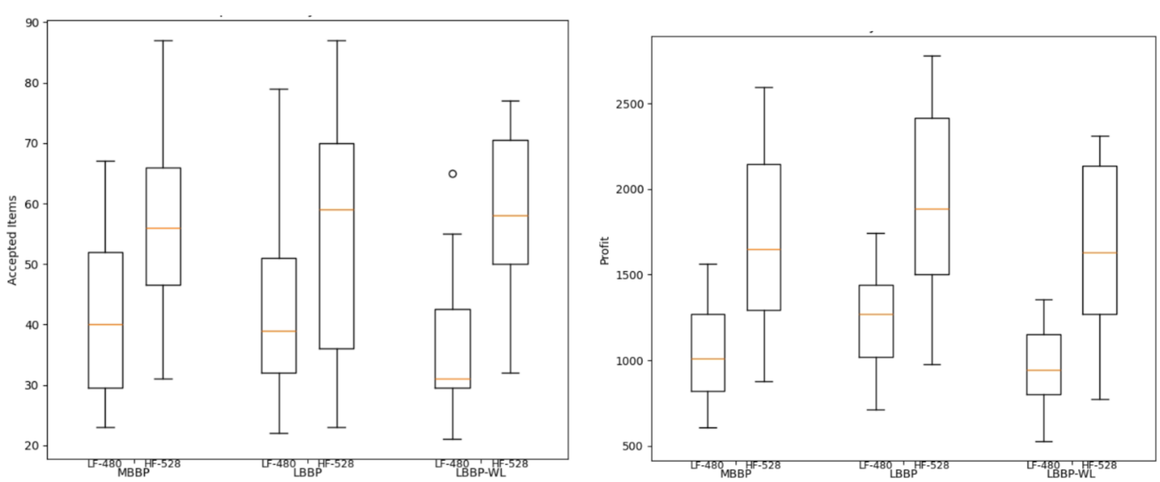

(a) Accepted item by model and item

(b) Profit by model and item

Figure 2: Item impact

Overall, the results reveal a clear and systematic pattern driven by the revenue composition of the request set. When the input is dominated by low-fare items (LF-480), all models accept fewer shipments and generate lower profits, as each accepted item contributes little revenue. Under these conditions, the system becomes far more sensitive to operational restrictions: once capacity is reduced due to low water levels (LBBP-WL), acceptance drops sharply. In contrast, when the request set consists primarily of high-fare items (HF-528), acceptance levels and total profits increase substantially. The high revenue per item makes each accepted booking worthwhile, enabling the lookahead model (LBBP) to better exploit profitable opportunities. Moreover, the negative impact of water-level–induced capacity reductions is noticeably weaker in this setting, because even a smaller number of accepted high-fare shipments is sufficient to maintain strong overall profitability. In other words, how strongly capacity reductions affect the system depends entirely on the revenue potential of the incoming items: they amplify losses when fares are low and dampen them when fares are high.

5.3.4 Shipment-request ways impact

This section examines how the configuration of shipment-service ways influences the performance of the booking system across different models. Two representative configurations are considered: VC-R, which features a higher share of services supported by small vessels and thus lower total network capacity, and VC-SF, which provides greater capacity due to the inclusion of large-vessel services. The goal is to assess how these structural differences affect acceptance decisions and profitability under the MBBP, LBBP, and LBBP-WL models.

Figure 3 illustrates the effect of two shipment-service way configurations on booking system performance across the MBBP, LBBP, and LBBP-WL models. The VC-R configuration represents a network with lower aggregate capacity due to a higher share of small vessel supported-services, whereas VC-SF provides substantially greater total network capacity. In terms of accepted shipments (Figure 3a), increasing total capacity from VC-R to VC-SF leads to a rise in median acceptance across all models, but this increase is modest relative to the scale of capacity growth. For MBBP, the median rises from roughly 45 to 55, with the overall range expanding from about 20–70 in VC-R to 25–85 in VC-SF. In LBBP, the median increases from about 50 to 65, with the upper range exceeding 80, representing the largest change among the models. In LBBP-WL, the median grows from around 40 to 55 but remains below LBBP due to the additional consideration of water-level–induced capacity reductions. These patterns confirm that acceptance decisions are primarily driven by the economic value of shipments rather than the nominal capacity of the network, as even in the higher-capacity case, the system accepts only as many shipments as remain profitable. For profit (Figure 3b), the increase from VC-R to VC-SF is more pronounced. In MBBP, the median profit rises from approximately 1200 to 1500–1800. In LBBP, the growth is more substantial—from around 1600 to over 2000, with a maximum near 2600. In LBBP-WL, the median improves from about 1100 to 1400–1500, with the maximum reaching around 2000. The primary driver of this improvement is the reduced reliance on costly outsourced paths in the VC-SF configuration. Higher in-network capacity allows a greater proportion of shipments to be routed internally at lower cost, thereby increasing total profit even if accepted items do not grow substantially. Overall, higher capacity (VC-SF) yields limited gains in accepted volumes but significant improvements in profit across all models. This outcome indicates that the booking system's acceptance policy is profitability-oriented rather than purely capacity-driven, with capacity expansion primarily enhancing economic efficiency by reducing dependence on expensive outsourced services.

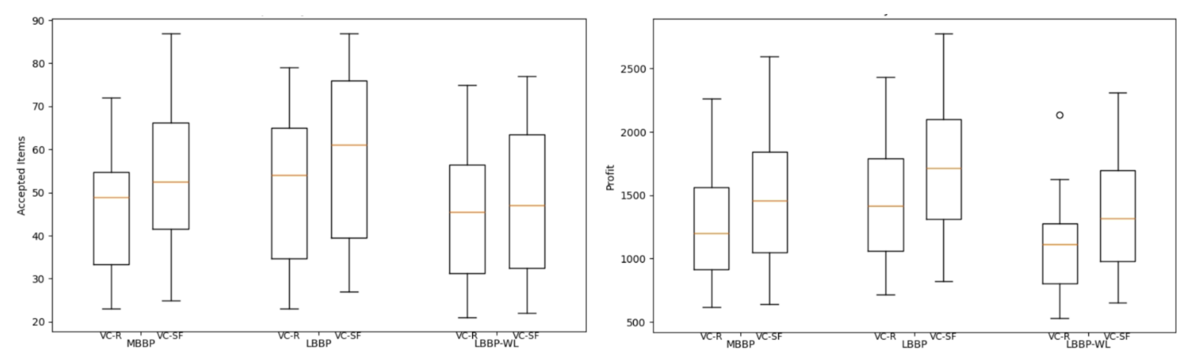

(a) Accepted item by model and ways (b) Profit by model and ways

Figure 3: Shipment-request ways impact

5.3.5 Prediction confidence impact

This section explores how the level of confidence associated with predicted demand (θ) influences booking decisions and overall system performance. Table 5 presents the results across three confidence levels—$\theta = 0.25$, 0.5, and 0.85—for each combination of way configuration and item mix. The parameter θ reflects the model's trust in the accuracy of predicted item attributes and plays a crucial role in shaping acceptance behavior, particularly under uncertainty. Table 5 shows that increasing prediction confidence (θ) from 0.25 (low confidence)to 0.85 (high confidence) systematically increases profit while

Table 5: Prediction confidence impact

Path	Item	$\theta = 0.25$		$\theta = 0.5$		$\theta = 0.85$	
		profit	AC	profit	AC	profit	AC
VC-U	LF-480	5450.50	0.40	5703.00	0.39	6048.10	0.38
	HF-528	8931.75	0.57	9170.50	0.57	9510.50	0.56
VC-SF	LF-480	6624.50	0.48	6964.00	0.48	7387.10	0.46
	HF-528	10345.50	0.63	10634.0	0.62	11024.0	0.62
VC-R	LF-480	5483.50	0.41	5735.00	0.39	6085.10	0.38
	HF-528	8967.00	0.58	9210.50	0.57	9546.50	0.56

slightly reducing acceptance rates across all path configurations and item mixes. Average profit rises from 7134 at $\theta = 0.25$ to 7758 at $\theta = 0.85$ ($\approx$8.7% gain), with the lowest value observed for LF-480/VC-U (5450.50) and the highest for HF-528/VC-SF (11024.00). Conversely, the average acceptance rate declines from 0.512 to 0.493, indicating that higher confidence leads the model to reject more current requests in favor of keeping flexibility to accommodate predicted high-value arrivals. This trade-off is most pronounced in high-fare–dominant cases, where profits grow disproportionately, and in the high-capacity VC-SF configuration, where reduced reliance on costly outsourced paths amplifies gains. Overall, the results confirm that profitability in the booking system is driven more by selective acceptance of high-value requests than by maximizing the number of accepted items.

5.3.6 Water level impact

We consider two contrasting capacity-reduction scenarios as representative worst cases: smooth fluctuations and abrupt fluctuations. These cases are chosen to stress-test the model under challenging water-level conditions: the smooth case reflects moderate reductions where capacity falls at most to 75% of nominal, while the abrupt case represents severe reductions where capacity can drop as low as 35%. Evaluating such extreme scenarios is advantageous because it highlights the robustness of booking decisions and reveals how system performance changes when service availability is most constrained.

Table 6 reports the evolution of profit, path cost, and way cost over five discrete decision steps under smooth and abrupt water-level fluctuations. In both scenarios, profit increases from step 1 to step 4 and then exhibits a slight decline at step 5. However, profit levels differ substantially across scenarios: under smooth fluctuations, profit remains consistently higher and reaches a maximum of approximately 1,350 at step 4, whereas under abrupt fluctuations the peak profit is lower, around 1,200. This gap directly

Table 6: Impact of smooth ($\gamma = 0.75$) and abrupt ($\gamma = 0.35$) water-level fluctuations on system performance

Decision step	$\gamma = 0.75$ (Smooth)			$\gamma = 0.35$ (Abrupt)		
	Profit	Path cost	Way cost	Profit	Path cost	Way cost
1	650	380	580	530	650	350
2	940	760	600	790	1050	350
3	1080	980	640	920	1280	380
4	1350	1400	730	1200	1720	430
5	1280	1080	590	1100	1450	280

reflects the impact of tighter capacity restrictions under low-water conditions, which constrain revenue generation even when demand remains available. Path costs also rise with successive decision steps in both scenarios, but their magnitude is markedly higher under abrupt fluctuations. In the smooth case, path costs increase broadly in line with profit and reach approximately 1,400 at step 4. Under abrupt fluctuations, however, path costs escalate much more sharply, exceeding 1,700 at step 4 and clearly surpassing profit levels. This behavior indicates that severe capacity reductions force the system to rely on more expensive routing and allocation options, thereby eroding margins and overall profitability. Way costs remain comparatively stable across decision steps, varying within a limited range in both scenarios. They are systematically lower under abrupt fluctuations, as stricter capacity constraints reduce the number of feasible service-way allocations, leading to lower overall way usage while shifting a larger share of demand toward costlier path-level decisions. Overall, the results indicate that smooth water-level fluctuations allow the system to preserve a closer balance between revenues and costs, whereas abrupt fluctuations disrupt this balance by disproportionately increasing path costs and limiting achievable profits. The slight decline in profit observed at the final decision step in both scenarios further suggests diminishing returns once accumulated commitments and tight capacity limits reduce the flexibility of feasible allocations.

6 Conclusions

This paper studied the booking problem for consolidation-based intermodal freight carriers, incorporating accept/reject decisions for non-contractual requests while ensuring fulfillment of contractual commitments. To our knowledge, this is the first work to address booking-level control in this context using a bin-packing framework, thereby linking tactical and booking planning. The approach first identifies time–space feasible service ways and then applies bin-packing models to evaluate profitability and make acceptance decisions. Three models were proposed: a myopic version (MBBP), a lookahead version (LBBP), and an extended lookahead model with water-level constraints (LBBP-WL). Together, they provide a unified framework to analyze how different informational and operational conditions influence booking decisions and system performance. Computational experiments on realistic-size instances revealed clear trade-offs between profit, acceptance rates, and computational effort. While MBBP was computationally efficient, LBBP achieved higher profits by exploiting demand predictions, and LBBP-WL captured the impact of capacity reductions due to water-level variability. Results highlighted the importance of anticipative decision-making and accounting for capacity variability to improve profitability. Additional analyses showed that fare heterogeneity, service and fleet configuration, prediction confidence, and water-level scenarios significantly affect solution structure, service utilization, and overall profit.

Finally, several directions remain open for future research. Extending the framework to consider order-based booking decisions, rather than batch-based decisions, would enable a more detailed representation of how shipments arrive and are processed, and allow comparisons between the two approaches to evaluate how the timing of decision-making influences system profitability and booking structures. Another promising extension is to consider flexible delivery times, where late deliveries are permitted with explicit penalty costs. The modelling structure developed in this paper naturally supports such an extension: since each shipment is assigned to a service way with well-defined timing attributes, the model can be augmented to represent cases where the assigned service departs or arrives after the shipment's due time. This would allow late deliveries to be explicitly captured and penalty costs to be computed as a function of the degree of lateness (e.g., linearly scaled with delay duration). Implementing this feature would provide additional flexibility in planning and enable a systematic evaluation of how different penalty structures influence the profitability and feasibility of booking decisions. Furthermore, while contractual shipment requests were treated as deterministic commitments in this study, relaxing this assumption and considering stochastic patterns even for contractual flows would provide a more realistic decision environment. Integrating such uncertainty into the accept/reject process for non-contractual requests, combined with efficient heuristic or decomposition methods, would improve scalability and extend the applicability of the models to larger and more complex problem instances.

Acknowledgements

While working on the project, the first author was Ph.D. candidate at the School of Management, Université du Québec à Montréal (UQAM), Canada, and member of CIRRELT. The third author held the UQAM Chair in Intelligent Logistics and Transportation Systems Planning, and was an Adjunct Professor in the Department of Computer Science and Operations Research at the Université de Montréal, while the fourth author held the Canada Research Chair in Stochastic Optimization of Transport and Logistics Systems. We gratefully acknowledge the financial support provided by the Natural Sciences and Engineering Research Council of Canada (NSERC), through its Collaborative Research and Development, and Discovery grant programs. We also gratefully acknowledge the support of Fonds de recherche du Québec through their infrastructure grants. Finally, during the preparation of this work the authors used ChatGPT-4, under its most strict privacy enforcing settings, for text edit purposes exclusively. After using this tool, the authors reviewed and further edited the text as needed and take full responsibility for the content of the publication.

References

T. Bektaş and T.G. Crainic. A Brief Overview of Intermodal Transportation. In Taylor, G.D., editor, *Logistics Engineering Handbook*, chapter 28, pages 1–16. Taylor and Francis Group, Boca Raton, FL, USA, 2008.

I. C. Bilegan, L. Brotcorne, D. Feillet, and Y. Hayel. Revenue management for rail container transportation. *EURO Journal on Transportation and Logistics*, 4(2):261–283, 2015.

T. Crainic and M. Hewitt. Service network design. In T. Crainic, M. Gendreau, and B. Gendron, editors, *Network Design with Applications in Transportation and Logistics*, chapter 12, pages 347–382. Springer, Boston, MA, 2021.

T. Crainic and W. Rei. 50 years of operations research for planning consolidation-based freight transportation. Research Report CIRRELT-2024-11, Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique et le transport, Université de Montréal, Université de Montréal, Montréal, 2024. Forthcoming in EURO Journal on Transportation and Logistics.

T. G. Crainic, F. D. Fomeni, and W. Rei. Multi-period bin packing model and effective constructive heuristics for corridor-based logistics capacity planning. *Computers & Operations Research*, 132:105308, 2021.

Y. Cui, I. C. Bilegan, E. Duchenne, and D. Duvivier. Demand rerouting mechanisms with revenue management for intermodal barge transportation networks. *Transportmetrica B: Transport Dynamics*, 12(1):2416182, 2024.

T. Delbart, Y. Molenbruch, K. Braekers, and A. Caris. Uncertainty in intermodal and synchromodal transport: Review and future research directions. *Sustainability*, 13(7): 3980, 2021.

R. Elbert, J. P. Müller, and J. Rentschler. Tactical network planning and design in multimodal transportation–a systematic literature review. *Research in Transportation Business & Management*, 35:100462, 2020.

W. Guo, B. Atasoy, W. B. van Blokland, and R. R. Negenborn. Global synchromodal transport with dynamic and stochastic shipment matching. *Transportation Research Part E: Logistics and Transportation Review*, 152:102404, 2021.

M. Kapetanović, N. Bojović, and M. Milenković. Booking limits and bid price based revenue management policies in rail freight transportation. *European Journal of Transport and Infrastructure Research*, 18(1), 2018.

J. Ksciuk, S. Kuhlemann, K. Tierney, and A. Koberstein. Uncertainty in maritime ship routing and scheduling: A literature review. *European Journal of Operational Research*, 308(2):499–524, 2023.

L. Lee, E. Chew, and M. Sim. A revenue management model for sea cargo. *International Journal of Operational Research*, 6(2):195–222, 2009.

K. Littlewood. Forecasting and control of passenger bookings. *The Airline Group of the International Federation of Operational Research Societies*, 12:95–117, 1972.

Q. Meng, S. Wang, H. Andersson, and K. Thun. Containership routing and scheduling in liner shipping: overview and future research directions. *Transportation Science*, 48 (2):265–280, 2014.

Q. Meng, H. Zhao, and Y. Wang. Revenue management for container liner shipping services: Critical review and future research directions. *Transportation Research Part E: Logistics and Transportation Review*, 128:280–292, 2019.

B. Payami, I. C. Bilegan, T. G. Crainic, and W. Rei. Tactical network planning for intermodal barge transportation considering varying water levels. Research Report CIRRELT-2025-17, Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique et le transport, Université de Montréal, Université de Montréal, Montréal, 2025a.

B. Payami, I. C. Bilegan, T. G. Crainic, and W. Rei. Service network design with uncertainty on water levels for intermodal river transport. Research Report CIRRELT-2025-16, Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique et le transport, Université de Montréal, Montréal, 2025b.

M. SteadieSeifi, N. P. Dellaert, W. Nuijten, T. Van Woensel, and R. Raoufi. Multimodal freight transportation planning: A literature review. *European journal of operational research*, 233(1):1–15, 2014.

H. Wang, X. Wang, and X. Zhang. Dynamic resource allocation for intermodal freight transportation with network effects: Approximations and algorithms. *Transportation Research Part B: Methodological*, 99:83–112, 2017.

X. Wang. Optimal allocation of limited and random network resources to discrete stochastic demands for standardized cargo transportation networks. *Transportation Research Part B: Methodological*, 91:310–331, 2016a.

X. Wang. Stochastic resource allocation for containerized cargo transportation networks when capacities are uncertain. *Transportation Research Part E: Logistics and Transportation Review*, 93:334–357, 2016b.

X. Wang. Static and dynamic resource allocation models for single-leg transportation markets with service disruptions. *Transportation Research Part E: Logistics and Transportation Review*, 103:87–108, 2017.

Y. Wang, I. C. Bilegan, T. G. Crainic, and A. Artiba. A revenue management approach for network capacity allocation of an intermodal barge transportation system. In *Computational Logistics: 7th International Conference, ICCL 2016, Lisbon, Portugal, September 7-9, 2016, Proceedings 7*, pages 243–257. Springer, 2016.