

Service Network Design with Uncertainty on Water Levels for Intermodal River Transport

**Bitu Payami
Teodor Gabriel Crainic
Walter Rei
Iona Bilegan**

February 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 Palais-Prince, local 2415
Québec (Québec) G1V 0A6
Tél : 1-418-656-2073
Télécopie : 1-418-656-2624

Service Network Design with Uncertainty on Water Levels for Intermodal River Transport[†]

Bitá Payami¹, Teodor Gabriel Crainic^{1,*}, Walter Rei¹, Iona Bilegan²

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 8201Valenciennes

Abstract. Barge transport presents a sustainable alternative to road transport, offering the potential to reduce costs. However, increasingly frequent and severe drought seasons, attributed to climate change, challenge the resilience of this mode of transport. Decreased water levels restrict navigation, impact vessel size, and reduce vessel capacity. This study introduces a novel and integrated modeling framework for tactical planning in consolidation-based barge freight transportation. By examining the relationship between water levels and vessel load capacity, with a specific focus on vessel dimensions, the framework aims to evaluate possible impacts on the efficiency and profitability of inland waterway transport. We introduce two stochastic Water-Level- Constrained Scheduled Service Network Design with Resource and Revenue Management models to address uncertainties in resource capacities caused by water level variations, each addressing a distinct tactical plan adjustment strategy. These models aim to establish a tactical plan, given predicted water levels, that maximizes the expected carrier's revenue while accounting for future adjustments to the plan when information is revealed and predictions are reliably updated, to fulfill the demands of shippers and optimize the utilization of the carrier's resources. Through extensive experimentation using commercial software and a novel decision-based scenario clustering algorithm, we assess the quality of the solutions obtained with the different model variants and analyze the impact of the water level on the results.

Keywords: Scheduled service network design, tactical planning, two-stage stochastic programming, water level fluctuations, scenario reduction, clustering

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 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.

[†]Revised version of CIRRELT-2025-16

Results and views expressed in this publication are the sole responsibility of the authors and do not necessarily reflect those of CIRRELT.

Les résultats et opinions contenus dans cette publication ne reflètent pas nécessairement la position du CIRRELT et n'engagent pas sa responsabilité.

* Corresponding author: teodorgabriel.crainic@cirrelt.net

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

© Payami, Crainic, Rei, Bilegan and CIRRELT, 2026

1 Introduction

Barge transportation, a key component of intermodal freight transportation, is known for its cost-effectiveness and environmental sustainability and plays a crucial role in facilitating the exchange of freight between maritime ports and their hinterlands, as well as among river ports. Carriers in this sector often use consolidation strategies to combine small and large shipments, enhancing operational efficiency and cost-effectiveness. However, this requires precise coordination of shipping schedules, freight specifications, and service requirements. To deliver shipments at low cost and with high punctuality, carriers need effective planning of transportation activities—detailing the service network, resource utilization, shipment handling and transport operations—in a way that ensures efficiency, profitability, and effective consolidation, while maintaining the desired level of service quality that is crucial for shipper satisfaction. This problem, addressed at the tactical level, is supported by the Service Network Design (SND) methodology. This tactical planning problem is typically addressed using the Scheduled Service Network Design (SSND) model, which defines an optimized transportation plan that specifies which services and schedules to operate, and how shipments are routed across the network for a given schedule length (e.g., a week), executed repeatedly over a medium-term horizon (e.g., a season), with the objective of maximizing profit.

Most service network design cases, including SSND models, often overlook environmental factors that can directly impact the capacity of transportation resources, which are used for operating scheduled services efficiently across land-, air-, and water-based modes. Beyond disregarding infrastructure conditions, these models are predominantly deterministic, operating under the assumption that there are no significant variations in the system’s state—whether on the supply or demand sides—throughout the planning horizon (Crainic and Hewitt, 2021; Crainic and Rei, 2025). SSND models generally presume that the capacity offered by each resource supporting the service over the designed network remains known and constant throughout the planning horizon. This is often not the case, however. Hence, our research centers on the uncertainty in infrastructure and resource capacity, particularly relevant for inland waterway transportation, where water-level fluctuations can significantly impact vessel capacity. Lower water levels can decrease vessel capacity due to increased grounding risks, while higher water levels might allow for greater freight capacity but could introduce navigational challenges under bridges and through certain canal sections (Prandstetter et al., 2023). These fluctuations result from transient weather conditions, such as rainfall and short dry spells, and are characterized as randomness—foreseeable variations that can be described through random variables with known probability distributions (Christodoulou et al., 2020; Zheng and Kim, 2017). Expert-generated forecasts provide probability distributions for water levels at critical river segments and ports, specifying the expected variation in a vessel’s

loading capacity based on its type. Such predictions are made for the next season at a fairly high aggregation level and are later repeatedly updated during operations for short (e.g., the schedule length) horizons. When the updated predicted water levels are lower than initially anticipated, carriers face reduced vessel capacities, forcing them to either reject part of the demand to be transported and pay any associated penalties or adapt their transportation plans through costly re-optimization. These dynamic adjustments, which deterministic SSND models are unequipped to handle, underscore the need for a stochastic service network design framework that integrates infrastructure uncertainty and enables flexible, informed decision-making under environmental variability. While there is a substantial body of literature on SSND models for consolidation-based transportation, most studies focus on demand-side uncertainty, with only a few addressing supply-side uncertainty and even fewer considering uncertainties in infrastructure and resource capacity. To the best of our knowledge, no existing work has jointly tackled both the design of an efficient service network and the impact of water level uncertainty on vessel-supported service capacity. The practical consequences of this gap are increasingly evident. In 2022, European inland waterway freight volumes declined by 5.5%, and total transport performance dropped by 10.6% to 122 billion ton-kilometers. The Rhine alone recorded a 6.8% decrease in cargo traffic, a 12.2% reduction in container volumes, and a 42.5% surge in dry cargo freight rates driven directly by reduced vessel capacity under low-water conditions. On the Middle Danube, grain transport capacity collapsed by 80% due to prolonged drought, underscoring the system’s acute vulnerability to environmental fluctuations (Eurostat, 2023).

We aim to fill this gap by introducing the *Water-Level-Constrained Scheduled Service Network Design with Resource and Revenue Management* (WL-SSND-RRM) problem, which addresses the challenges of integrating water-level variability into tactical planning for barge transportation systems. We propose two models, formulated within a two-stage stochastic programming framework, each defined by a specific strategy to handle the impact of randomly changing water-level conditions. Both models aim to establish a tactical operations plan, given predicted water levels, that maximizes the expected carrier’s revenue while accounting for future possible adjustments to the plan as information is revealed and forecasts are updated—thereby fulfilling shipper demand and optimizing resource utilization.

In the proposed stochastic models, water level uncertainty is represented through a scenario set that captures random variations in service capacities. To reflect this variability accurately, the scenario set must be sufficiently comprehensive, which often requires a large number of scenarios and, consequently, increases the amount of contextual information that must be integrated into the optimization process. However, incorporating all scenarios into the stochastic WL-SSND-RRM models can significantly increase computational complexity, rendering them intractable in practice. To address this challenge, we adopt a decision-based scenario clustering approach inspired by Hewitt et al. (2022). However, our use of this approach differs from the original method in a meaningful way. Whereas Hewitt et al. (2022) cluster scenarios based on traditional sources of operational

uncertainty (e.g., demand), we cluster water-level scenarios according to the tactical decisions they induce through capacity changes. In Step 1, we form clusters of water-level scenarios that lead to similar tactical decisions, thereby reducing redundancy in the scenario set while minimizing information loss. At the same time, we also form clusters of scenarios that lead to dissimilar tactical decisions, ensuring that the reduced set captures diverse, decision-relevant information. In Step 2, the resulting scenario clusters are used to apply a series of bounding techniques. These techniques yield alternative feasible solutions whose objective values provide lower bounds for the stochastic models, while also producing an upper bound that serves as a certificate for evaluating the optimality gaps of the obtained solutions.

The contributions of this paper are the following: (i) We introduce the WL-SSND-RRM problem, which explicitly accounts for uncertainties in resource capacities caused by random water-level variations—a critical factor that directly impacts vessel capacity and service feasibility and has not been addressed in the literature; (ii) This work proposes two stochastic models, each defined by a specific approach to handling the impact of observed randomly changing water-level conditions on the execution of the tactical plan. The first model allows demand itinerary adjustments at a cost, enabling shipment flows to dynamically adapt to realized water-level conditions. The second model introduces penalties for unmet demand, maintaining a fixed tactical plan even when water levels deviate from expectations ; (iii) We complement the proposed stochastic models with a solution approach that makes them computationally tractable in practice. Specifically, we employ a decision-based scenario clustering analysis approach to reduce redundancy in the scenario set while preserving decision-relevant variability and limiting information loss. This reduction allows us to solve the models using a commercial solver such as Gurobi to generate bounds, without the need to develop problem-specific algorithms. Although the clustering method was originally proposed in the context of stochastic service network design, its effectiveness for the WL-SSND-RRM problem illustrates how large-scale stochastic models of this type can be solved and applied in realistic inland waterway planning settings; (iv) We conduct extensive computational experiments to evaluate the performance of our stochastic models compared to the deterministic model. Also, we analyze how different strategies to handling the impact of randomly changing water levels influence the structural characteristics of tactical plans. Additionally, we assess the impact of water-level fluctuations on the efficiency and profitability of the resulting tactical plans.

The paper is organized as follows: Section 2 presents the barge transportation system, including the physical network, supply side, and demand side. Section 3 provides a comprehensive review of the relevant literature, encompassing discussions on both deterministic and stochastic service network design, and examines various solution methods proposed in this field. Section 4 outlines the problem setting and introduces the proposed stochastic WL-SSND-RRM formulations for the barge transportation system. Section 5 explains the solution approach adopted in this research. Section 6 reports the experimental results. Finally, Section 7 summarizes the main findings and outlines future research

directions.

2 Barge transportation system description

Freight transportation via inland waterways involves the movement of shipments from their point of origin to their destination using vessels that operate along rivers, canals, and related infrastructure (Bilegan et al., 2022). A detailed description of the physical infrastructure involved in barge transportation, including waterways and terminals, will be provided in Section 2.1. The key stakeholders in this system are carriers and shippers, representing the supply and demand sides, respectively, which will be further explored in Sections 2.2 and 2.3.

2.1 Physical network

Similar to land- and air-based transportation modes, which require a certain set of infrastructure components, water-based transportation has its own physical network, which consists of waterways and ports. Waterways are corridors characterized by their length, width, and depth, connecting ports and allowing vessels to navigate between them to transport freight from one terminal to another. Depending on their geographic location, waterways are categorized as deep sea, short sea, coastal, or inland. In contrast to the deep sea, the coastal and inland waterways studied here are shallower, narrower, and shorter. Due to these physical characteristics, smaller vessels, such as barges, are used to navigate these waterways. These attributes limit not only the size of vessels but also the number of vessels that can travel simultaneously in both directions on a given section of the waterway within a specific time frame. While the width and length of rivers remain relatively constant over time, the depth—defined as the distance between the water surface and the highest point of the riverbed—fluctuates. Water depth is influenced by various factors over short, medium, and long-term periods, such as extreme weather events, seasonal changes, and annual warming and cooling cycles, making navigation more challenging at times.

Ports, the other key component of water-based transportation, are facilities located along coasts or rivers where vessels dock to load and unload freight. Ports are composed of one or more quays and terminals. A quay is the structure along the edge of the water where vessels are moored, and its length is often considered to match that of the terminal, especially in simplified terminal layouts. A terminal is a specialized facility within a port designed for handling freight. Terminals are equipped with cranes, warehouses, storage areas, and other infrastructure necessary for loading and unloading vessels and transferring freight to and from trucks, trains, or other modes of transportation. The main physical constraints of each terminal are the water depth at the berth, the number

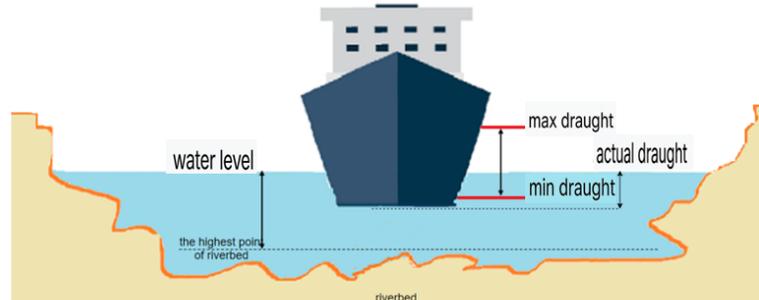


Figure 1: Schematic view of the most relevant measurements related to vessel's draught and water level

of load-handling machines, and the available storage space. The berthing capacity—the number of vessels that can dock simultaneously—is limited by the terminal's length. The water depth limits the maximum submerged vertical dimension of a loaded vessel, defining the safe docking depth at the terminal, while storage capacity limits the volume of freight that can be stored at the terminal for a given period. Therefore, each terminal's capacity is measured based on several dimensions, such as the volume of freight that can be loaded or unloaded per time unit (i.e., handling capacity), the number of vessels that can be docked at the terminal in a given length unit (i.e., berthing capacity), and the volume of freight that can be stored at the terminal per time unit (i.e., storage capacity). Carriers using the port have to pay the costs associated with each terminal. Load handling, load storage, and vessel holding costs are all included in the costs associated with each terminal. The first two costs—load handling cost and load storage cost—vary depending on the type of container used, while the vessel holding cost for each terminal depends on the type of the vessel.

2.2 The barge transportation system's supply side

Carriers in the barge transportation system offer scheduled services to transport freight between terminals. These services are designed to meet a variety of shipping needs by offering flexible routes and schedules to ensure shipments are delivered from origin to destination within a specified time window. Each potential scheduled service follows a specific waterway route within the network, which may be either a single-leg service (traveling directly from origin to destination without stops) or a multi-leg service (stopping at multiple terminals to load and unload freight). Each service is defined by its

origin, destination, intermediate terminals (if any), and scheduled arrival and departure times at each terminal. The total duration of any service, whether single-leg or multi-leg, is determined by the cumulative travel times for each leg and any waiting times at terminals (Crainic and Hewitt, 2021). Carriers offer a range of fare classes, which are price categories designed to reflect different levels of service quality that may vary in terms of delivery time and flexibility. These fare classes allow shippers to choose between different service quality options. For delivery time, shippers can opt for standard delivery, which involves longer transit times, or express delivery, which offers faster transit times. Overall, fare classes enable carriers to provide tailored quality of service levels that meet the varying needs of shippers.

Services are performed by vessels of various types in the barge transportation system. Each vessel type has specific technical specifications: length, typically measured in meters or feet; speed, measured in knots; and critical depth characteristics, such as maximum and minimum draught (See Figure 1). The maximum draught refers to the deepest point a fully loaded vessel can safely reach below the waterline, determining the maximum load it can carry without compromising safety. Conversely, the minimum draught represents the shallowest depth the vessel reaches when empty (Prandtstetter et al., 2023). These last two predefined specifications establish the safe loading range, or loading capacity, of each vessel, which can be measured in terms of weight (tons), volume (TEUs), or a combination of both. When a vessel is partially loaded (i.e., carrying less than its maximum capacity but not empty), its submerged depth—referred to as the actual draught—naturally falls between the minimum and maximum draughts, as illustrated in Figure 1. The actual draught can be described as an incremental function, where each added unit of load results in a proportional increase in draught (van Dorsser et al., 2020). In this context, the load-draught coefficient signifies how sensitive the draught is to changes in load, capturing this relationship whether it follows a simple linear progression or a more complex pattern.

2.3 The barge transportation system’s demand side

Shippers generate the demand for transportation services in the barge transportation system. These shippers include manufacturers, distributors, wholesalers, retailers, and other businesses that require the movement of shipments from one port to another. The shipments may consist of finished products, intermediate goods, raw materials, or other types of freight, which can be packaged in various forms such as boxes, bags, drums, bales, and rolls. In the intermodal barge transportation system, loads transported by vessels are typically placed into containers, offering several advantages, including improved load safety, reduced handling costs, and standardization (Bektaş and Crainic, 2008). Standard containers come in two lengths—20 feet and 40 feet. The Twenty-foot Equivalent Unit (TEU) is the standard measure of container volume, where a 20-foot container equals 1 TEU, and a 40-foot container equals 2 TEUs.

Shippers on the demand side of the barge transportation system can be categorized

based on their business relationship with carriers as regular or spot shippers; Regular shippers: have long-term contracts with carriers, ensuring consistent and predictable transportation demands and benefiting from guaranteed capacity. Spot shippers: on the other hand, request transportation services only when needed, without engaging in long-term contractual commitments (Bilegan et al., 2022). Regardless of whether a shipper is regular or spot, each shipper’s demand has several defining characteristics: the origin and destination of the shipment, specifying the pickup and delivery ports; the demand size, measured by volume (TEUs) and weight (total tonnage); and the type of container, such as open-top, open-side, or refrigerated. Additionally, shippers choose fare options based on their delivery needs, opting for either standard delivery or express delivery.

3 Literature Review

This section reviews the relevant literature on planning problems in intermodal barge transportation systems. The discussion emphasizes three interconnected aspects: planning studies related to service network design in barge transportation, the influence of infrastructure and environmental limitations on tactical decisions, and the development of optimization-based models and solution techniques proposed to support intermodal barge transportation at the tactical planning level.

Service network design problems have been extensively studied in the literature to address tactical planning issues for consolidation-based carriers (Crainic and Hewitt, 2021; Crainic and Rei, 2025). The reader may refer to reviews on this field for long-haul transportation by (Crainic, 2003), for rail by (Cordeau et al., 1998), for maritime transportation (Christiansen et al., 2007), for motor carriers (Bakir et al., 2021) and for intermodal transportation by (Crainic and Kim, 2007; Agamez-Arias and Moyano-Fuentes, 2017; SteadieSeifi et al., 2014). Despite the breadth of research on service network design, studies specifically targeting tactical planning in the context of barge transportation are relatively scarce. In intermodal barge transportation, tactical planning has been studied with a focus on service route selection, service frequency and capacity definitions, and in some cases, the repositioning of empty containers (Caris et al., 2012; Braekers et al., 2013; Riessen et al., 2015; Ypsilantis and Zuidwijk, 2019). Nevertheless, most of these studies concentrate on barge routing and repositioning without incorporating broader operational challenges, particularly those related to water depth variability and bridge height limitations—key factors that directly affect the feasibility of barge operations. Only a few studies explicitly consider these infrastructure-related constraints. A notable contribution is by Zhang et al. (2020), who propose the use of foldable containers to address height and depth limitations in river–sea intermodal transport. Their mixed-integer linear programming model optimizes empty container repositioning under these structural constraints in a deterministic setting. Müllerklein and Fontaine (2025) also work on integrating disruption risks into transportation planning by proposing a two-stage stochastic model that optimizes the selection and combination of resilience strate-

gies—such as rerouting, inventory positioning, and sourcing flexibility—under transport uncertainty. Although their work focuses on general supply chain networks, the approach highlights the importance of incorporating tactical and operational responses to infrastructure-related disruptions. This perspective is relevant for barge transportation systems, where fluctuating water levels can significantly impact route feasibility and vessel capacity, yet remain largely overlooked in the existing tactical planning literature.

A significant share of the literature on tactical planning for consolidation-based carriers leverages time-dependent service network design models, commonly referred to as Scheduled Service Network Design (SSND) models. These models incorporate the temporal dimensions of both demand and service through time-space networks, which are constructed by extending the network along the dimension of time for the duration of the schedule length. SSND models have been extensively applied in deterministic frameworks across transportation modes. Comprehensive overviews of SSND models are presented in (Crainic and Hewitt, 2021) and (Crainic and Rei, 2025), including model structures, applications, and solution strategies. SSND has recently been adapted to the context of intermodal barge transportation. Bilegan et al. (2022) developed a model tailored for inland waterway networks, enhancing the classical SSND framework by integrating resource and revenue management—resulting in the Scheduled Service Network Design with Resource and Revenue Management (SSND-RRM). This model jointly optimizes service selection, resource utilization, and demand routing to maximize carrier profitability while selecting high-value spot shipments and accommodating varied service quality requirements. Building on these contributions, our work develops a tactical planning framework that explicitly incorporates the stochastic nature of vessel-supported service capacity, which is affected by unavoidable water level fluctuations due to environmental and weather-related variability. We formulate two-stage stochastic programming models to capture this uncertainty—an aspect rarely addressed in SSND literature despite its operational significance. While demand uncertainty has been the primary focus of stochastic SSND models (Bai et al., 2014; Hewitt et al., 2019; Hoff et al., 2010; Lium et al., 2009; Wang et al., 2019; Wang and Qi, 2020), other sources of uncertainty remain underexplored. Most models are typically structured as the two-stage stochastic programming framework, with service selection in the first stage under probabilistic demand, and recourse actions such as rerouting or outsourcing in the second stage once actual demand is revealed (Crainic et al., 2011, 2014). Some contributions, such as (Hewitt et al., 2019), further extend this framework by including resource acquisition and allocation decisions in the first stage. In comparison, travel time uncertainty has received limited attention. Lanza et al. (2021) proposed using penalties for lateness as a recourse strategy, while Lanza et al. (2024) proposed a more adaptive approach involving departure rescheduling, flow rerouting, and outsourcing. In contrast, capacity uncertainty remains largely overlooked in the SSND literature. To date, (Sun et al., 2018) is the only study that explicitly models this aspect, using fuzzy chance constraints within a Mixed Integer Linear Programming (MILP) framework to capture rail capacity limitations in intermodal routing, ensuring feasibility with a certain probability. Notably, the model does not incorporate recourse decisions to address capacity shortfalls, leaving a significant

gap in adaptive planning under resource capacity uncertainty.

As mentioned, two-stage stochastic programming is the most commonly used formulation to model uncertainty in SSND problems. In these models, uncertainty is captured through a set of scenarios representing random variations, which must be large enough to provide an accurate and detailed representation of uncertainty. However, a larger scenario set significantly increases computational complexity in the optimization process. To address this challenge, various solution methods have been explored, including metaheuristics (Hoff et al., 2010), progressive-hedging-based meta-heuristic (Lanza et al., 2021), matheuristics that utilize column generation (Hewitt et al., 2019), and partial Benders decomposition methodology (Crainic et al., 2021). The complexity of these stochastic combinatorial optimization problems has led to the adoption of scenario clustering, where meta-heuristics first cluster scenarios and then apply progressive hedging to solve the subproblems within these clusters, demonstrating computational success in stochastic network design (Crainic et al., 2014; Jiang et al., 2021). Recently, (Hewitt et al., 2022) proposed a solution approach consisting of two steps: first, scenarios are grouped based on their similarity or dissimilarity in the solutions they induce; second, the scenario cluster information is leveraged to compute efficient bounds for stochastic two-stage service network design models.

4 Problem setting and modeling approach

This section presents two stochastic formulations of the Water-Level-Constrained Scheduled Service Network Design with Resource and Revenue Management (WL-SSND-RRM) problem developed for the tactical planning of barge transportation, where the plan is established for a schedule length and repeated over the planning horizon, referred to as a season. The WL-SSND-RRM problem involves designing a scheduled service network, managing resources to support services, optimizing profit through revenue management via capacity management, and routing demand under fluctuating water levels, ensuring demand fulfillment while maintaining profitability and operational efficiency. First, we introduce the physical system, demand, potential service network, and economic elements that apply to both stochastic variants of the WL-SSND-RRM problem in Section 4.1. This is followed by a detailed explanation of the specific aspects and mathematical formulations for each variant in Section 4.2.

4.1 Problem Setting and Notation

Physical System. The physical infrastructure system is represented by a physical network, $\mathcal{G}^{\text{ph}} = (\mathcal{N}^{\text{ph}}, \mathcal{A}^{\text{ph}})$, where \mathcal{N}^{ph} denotes the set of geographical ports and \mathcal{A}^{ph} represents the set of physical links. Each port $\eta \in \mathcal{N}^{\text{ph}}$ has a specific berthing capacity,

measured in length units and denoted by Q_η . Each physical link $a \in \mathcal{A}^{\text{ph}}$ is characterized by a specific water level at each time period, denoted by p_a , which can be represented either by a point estimate in the deterministic case or by a probability distribution in the stochastic case. The point estimate is typically derived statistically from historical data, often using the mode or average to represent the most likely water level on each physical link over the planning horizon. In contrast, water levels in the stochastic case are modeled using probability distributions to capture variability and uncertainty. These distributions reflect both the likelihood of water levels falling within a typical range and the less probable extreme values, providing a comprehensive representation of water level fluctuations over time.

Demand. We define the set of shipping demands $\mathcal{K} = \mathcal{K}^{\text{R}} \cup \mathcal{K}^{\text{P}} \cup \mathcal{K}^{\text{F}}$, where \mathcal{K}^{R} (regular shippers) includes demands that must be fully served, \mathcal{K}^{P} (partial-spot shippers) includes demands that may be partially met, and \mathcal{K}^{F} (full-spot shippers) includes demands that can either be fully met or not met at all. Each demand $k \in \mathcal{K}$ is characterized by its volume $d(k)$ in TEUs. The conversion factor from TEUs to tonnage, denoted by ω , relates the volume $d(k)$ to the corresponding weight. Each demand is further defined by its origin terminal $O(k) \in \mathcal{N}^{\text{ph}}$, destination terminal $D(k) \in \mathcal{N}^{\text{ph}}$, release time $\alpha(k)$ at the origin terminal, and due time $\beta(k)$ at the destination terminal. Shippers are assumed to request either standard or express delivery, indicated by $class(k)$, regardless of their category. Obviously, express deliveries are charged a higher fare than standard deliveries. The parameter $\phi(k)$ specifies the unit fare for each demand, corresponding to its fare class $class(k)$. Let Γ denote the set of container types, with $\gamma(k) \in \Gamma$ representing the container type for each demand. The type of container affects the handling, storage, and transportation costs, due to differing requirements and handling complexities.

Service Network. Let \mathcal{V} denote the set of vessel types, and let F_v represent the maximum number of vessels of type $v \in \mathcal{V}$ available. Each vessel $v \in \mathcal{V}$, pre-assigned to a service, follows a circular sequence of services, starting and ending at the same port. Vessels are characterized by their nominal speed and nominal capacity, measured as $cap_w(v)$ in tonnage and $cap_{vol}(v)$ in TEUs. Additional characteristics include length $len(v)$, maximum draught $dh^+(v)$, minimum draught $dh^-(v)$ and the load-draught coefficient $\theta(v)$, measured in meters per tonnage, which quantifies how changes in loaded weight affect the vessel's draught. The set of potential services the carrier may operate to meet transportation demand is denoted by Σ . Each potential service $\sigma \in \Sigma$ is defined by its route in the physical network, by its schedule, and by its vessel type. The route of a potential service σ is specified as an ordered set of consecutive stops, including the origin, destination, and intermediate stops, denoted as $\mathcal{N}^{\text{ph}}(\sigma) = \{\eta_i(\sigma) \mid i = 0, \dots, n(\sigma)\}$. Here, $n(\sigma) = |\mathcal{N}^{\text{ph}}(\sigma)| - 1$ and i indicates the i^{th} stop of the service, with $\eta_0(\sigma) = O(\sigma)$ and $\eta_n(\sigma) = D(\sigma)$ as the origin and destination of the service, respectively. These services are characterized by schedules that specify the departure and arrival times, $\alpha(\eta_i(\sigma))$ and $\beta(\eta_i(\sigma))$, respectively, at each terminal $\eta_i(\sigma)$ in $\mathcal{N}^{\text{ph}}(\sigma)$. Each service has a total duration $\delta(\sigma)$, which includes the time spent at each stop as well as the travel time for each leg. A leg $l_i(\sigma)$ is defined as the segment between each pair of consecutive stops and

is expressed as $(\eta_{i-1}(\sigma), \eta_i(\sigma))$ for $i = 1, \dots, n(\sigma)$.

The WL-SSND-RRM problem is formulated on a time-space network $\mathcal{G} = (\mathcal{N}, \mathcal{A})$. This network is based on time discretization over the schedule length T , divided into equal-length time periods, $t \in \{0, 1, \dots, T-1\}$. The node set \mathcal{N} is defined as $\{(\eta, t) \mid \eta \in \mathcal{N}^{\text{ph}}, t = 0, \dots, T-1\}$, representing all terminals in the physical network at each time instant. The arc set \mathcal{A} is the union of moving arcs and holding arcs, $\mathcal{A} = \mathcal{A}^M \cup \mathcal{A}^H$. The set \mathcal{A}^M represents movements between nodes and is defined as: $\mathcal{A}^M = \{((\eta, t), (\eta', t')) \mid \eta, \eta' \in \mathcal{N}^{\text{ph}}, t, t' \in \{0, \dots, T-1\}, t < t'\}$. This indicates movements between nodes η and η' , departing at time t and arriving at time t' . The set \mathcal{A}^H is defined as: $\mathcal{A}^H = \{((\eta, t), (\eta, t+1)) \mid \eta \in \mathcal{N}^{\text{ph}}, t \in \{0, \dots, T-1\}\}$. These represent a one-time period waiting at terminal η at time t for vessels, demand, and services. According to the definition of moving arcs in the time-space network, each service leg corresponds to a moving arc. Specifically, a moving arc standing for service leg $l_i(\sigma) = \{(\eta_{i-1}(\sigma), \eta_i(\sigma)) \mid i = 1, \dots, n(\sigma), \sigma \in \Sigma\}$, is defined as $a_{l_i(\sigma)} = ((\eta_{i-1}(\sigma), \alpha(\eta_{i-1}(\sigma))), (\eta_i(\sigma), \beta(\eta_i(\sigma))))$. This arc indicates the departure of the service leg from terminal $\eta_{i-1}(\sigma)$ at time $\alpha(\eta_{i-1}(\sigma))$ and its arrival at terminal $\eta_i(\sigma)$ at time $\beta(\eta_i(\sigma))$. The second type of arc referred to as a holding arc, is thus defined as $a_{\eta t} = ((\eta, t), (\eta, t+1))$, where $(\eta, t) \in \mathcal{N}$. As we formulated the problem on a time-space network, the characteristics of nodes and arcs in the physical network are represented in the time-space network. The water levels, which are defined for each physical link over different periods, can be directly mapped onto moving arcs in the time-space network. Specifically, the water level for a physical arc, denoted as p_a , is represented as $p_{a_{l_i(\sigma)}}$ for moving arcs $a_{l_i(\sigma)} \in \mathcal{A}^M$. Similarly, the berthing capacity Q_η for nodes is represented as $Q_{a_{\eta t}}$ for holding arcs $a_{\eta t} \in \mathcal{A}^H$.

In this setting, we consider the following information revelation and decision-making process. At the tactical planning stage, a transportation plan is developed that encompasses the establishment of a network of scheduled services operated by assigned vessels and the determination of itineraries for regular and profitable spot shippers' demand. Since each service is directly linked to a vessel, the planning process implicitly manages vessel scheduling as part of resource management. However, it also explicitly includes additional resource management decisions, such as determining the number of temporarily idle vessels awaiting their next service departure and the total number of vessels utilized in the service plan. These decisions—encompassing the establishment of a network of scheduled services, the determination of demand itineraries, and explicit resource management—are made during the planning stage, before observing actual water level realizations, and are based solely on statistical distributions of water levels. Once these decisions are made and just before service operations begin (i.e., at the start of the schedule length), we assume that more precise water level forecasts become available. This assumption is consistent with the scope of the paper: incorporating multiple updates within the schedule length would correspond to operational planning, which involves short-term adjustments during execution and lies outside the scope of this study. Such forecasts may be obtained through specialists' interpretation of updated short-term weather and hydrological data, and they are expected to be highly accurate because the schedule length is

sufficiently short to maintain predictive precision and reflect observed weather and hydrological patterns. Thus, improving the accuracy of water level estimations shortly before transportation services begin is both feasible and highly advantageous, as it enables carriers to incorporate the latest data into their decision-making processes, refine tactical plans, and adjust them based on updated water level information. This ensures better alignment with operational realities and helps address limitations caused by fluctuating water levels, which can lead to partial or total reductions in vessel capacity, directly affecting the feasibility of demand itineraries by restricting the amount of demand that can be transported over planned services. Specifically, carriers might re-plan demand itineraries by modifying itineraries or selectively mitigating demand over planned demand itineraries by paying penalties. These adjustments ensure that the transportation plan remains feasible, efficient, and adaptable to dynamic conditions.

Costs. The following unit costs are defined: Let $h(\eta, v)$ denote the holding cost associated with each vessel type at the terminal for one time period, and $\mu(v)$ the cost associated with each vessel type that is not used in the optimal plan. Let $f(\sigma)$ denote the fixed cost of setting up and operating the service σ , and let $c_i(\gamma(k), v(\sigma))$ represent the transportation cost of a container of type $\gamma(k)$ by a vessel of type $v(\sigma)$ on the i^{th} leg of service σ . Here, $\gamma(k) \in \Gamma$ denotes the container type for demand d , and $v(\sigma) \in \mathcal{V}$ specifies the type of the vessel assigned to service σ . Let $c(\eta, \gamma(k))$ denote the cost of holding a container of type $\gamma(k)$ at terminal η for one period and $\kappa(\eta_i(\sigma), \gamma(k))$ represent the loading/unloading cost of a container of type $\gamma(k)$ on the i^{th} leg of service σ . Finally, $b(k)$ represents the penalty cost incurred when failing to meet demand $k \in \mathcal{K}$.

4.2 Two-stage Stochastic WL-SSND-RRM Model

Building on the previously described information revelation and decision-making process, we adopt a two-stage stochastic programming framework to establish a tactical plan that maximizes the carrier's overall profit. This plan is enriched by the expected adjustment (recourse) costs incurred at the operational level in response to realized water-level conditions, which impact vessel-supported service capacities. These recourse costs depend not only on the specific water-level realizations and their resulting impact on capacity availability but also on the type of adjustment strategy implemented. They form the foundation for two distinct two-stage stochastic models: *Two-Stage Stochastic Programming with Demand Flow Adjustment Recourse* (2-SPDFA) and *Two-Stage Stochastic Programming with Selective Demand Mitigation Recourse* (2-SPSDM).

In the 2-SPDFA model, the first-stage tactical plan includes selecting services and managing resources, while the second-stage decisions determine itineraries for both regular and selected spot shippers in response to the updated capacity conditions. This recourse strategy (*Demand Flow Adjustment Strategy*) emphasizes flexibility by adjusting demand itineraries after water levels are realized, aiming to minimize routing costs while ensuring demand fulfillment across the service network under all water-level realiza-

tions. To guarantee feasibility, the model introduces artificial origin–destination arcs that represent outsourcing options, ensuring that all regular and selected spot demand can be fulfilled even when vessel capacity is insufficient because of extremely low water-level conditions.

In contrast, some carriers may prefer stability over flexibility when reacting to capacity reductions. For these carriers, the 2-SPSDM model may represent a more suitable option. The 2-SPSDM model establishes a complete tactical plan in the first stage—including service selection, resource management, and itineraries for both regular and selected spot shippers—and addresses capacity shortfalls through the *Selective Demand Mitigation Strategy*. In this case, lower-priority or less profitable demands are selectively refused, while high-priority shipments are prioritized even if penalties are incurred for unfulfilled demands. The objective of this recourse strategy is to mitigate the impact of capacity constraints on critical shipments while balancing the penalties associated with unmet demands.

These two models therefore represent fundamentally different ways in which carriers may respond to water-level uncertainty—one emphasizing flexibility through flow adjustments, and the other emphasizing stability through selective demand control. The objective of each model is to maximize the carrier’s overall profit across the two stages by considering the costs associated with tactical-level decisions in the first stage and the expected adjustment (recourse) costs in the second stage. The recourse costs in the 2-SPSDM model consist of penalty costs for unmet demand, whereas in the 2-SPDFA model they encompass adjustment costs for demand itineraries, including demand routing costs (e.g., moving, loading, and unloading) and demand holding costs. In these stochastic models, a finite scenario set S is used to approximate the underlying probability space, where each scenario $s \in S$ represents a possible realization of water levels on each service leg. Each scenario is assigned a probability ρ^s , with $\sum_{s \in S} \rho^s = 1$. The detailed formulations are presented in the following sections.

4.2.1 The 2-SPDFA model

To formulate the 2-SPDFA model, we define four sets of decision variables. The first three belong to the first stage, while the fourth is associated with the second stage. The decision variables are defined as follows:

- **Service Selection**

- $y(\sigma) \in \{0, 1\}$: Equals 1 if transportation service $\sigma \in \Sigma$ is selected in the tactical plan, and 0 otherwise.

- **Shipper Selection**

- $\xi(k) \in [0, 1]$: Proportion of the volume of partial-spot shipper demand $k \in \mathcal{K}^P$ that is selected and serviced.
- $\zeta(k) \in \{0, 1\}$: Equals 1 if full-spot shipper demand $k \in \mathcal{K}^F$ is serviced, and 0 otherwise.

- **Resource Management**

- $z(v, a_{\eta_t}) \in \mathbb{Z}_{\geq 0}$: Number of temporarily idle vessels of type v waiting at holding arc a_{η_t} for the departure of the next service they support.
- $B(v) \in \mathbb{Z}_{\geq 0}$: Total number of vessels of type v used in the service plan. Due to the circular nature of the schedule, $B(v)$ remains constant across all time periods, although vessels may be moving or idle in ports at any given time.

- **Demand Flow Distribution (Recourse Decisions)**

- $x(k, a_{l_i}(\sigma))^s \geq 0$: Volume of demand $k \in \mathcal{K}$ transported on the i^{th} leg of service σ under scenario s .
- $x^{\text{out}}(k, a_{l_i}(\sigma))^s \geq 0$: Volume of demand $k \in \mathcal{K}$ unloaded from the i^{th} leg of service σ under scenario s .
- $x^{\text{in}}(k, a_{l_i}(\sigma))^s \geq 0$: Volume of demand $k \in \mathcal{K}$ loaded onto the i^{th} leg of service σ under scenario s .
- $x^{\text{hold}}(k, a_{\eta_t})^s \geq 0$: Volume of demand $k \in \mathcal{K}$ held at terminal η during the time interval $(t, t + 1)$ under scenario s .

Mathematical formulation. The problem is formulated as follows:

$$\begin{aligned}
 \max \quad & \sum_{k \in \mathcal{K}^R} \phi(k)d(k) + \sum_{k \in \mathcal{K}^P} \phi(k)\xi(k)d(k) + \sum_{k \in \mathcal{K}^F} \phi(k)\zeta(k)d(k) \\
 & - \sum_{v \in \mathcal{V}} \mu(v)(F_v - B(v)) - \sum_{\sigma \in \Sigma} f(\sigma)y(\sigma) - \sum_{a_{\eta_t} \in \mathcal{A}^H} \sum_{v \in \mathcal{V}} h(\eta, v)z(v, a_{\eta_t}) \\
 & - \sum_{s \in \mathcal{S}} \rho^s G^s(\chi)
 \end{aligned} \tag{1}$$

Subject to

$$B(v) = \sum_{\eta \in \mathcal{N}^{\text{ph}}} z(v, a_{\eta_0}) + \sum_{\sigma \in \Lambda_{0l}} y(\sigma), \quad \forall v \in \mathcal{V} \tag{2}$$

$$B(v) \leq F_v, \quad \forall v \in \mathcal{V} \tag{3}$$

$$\sum_{\sigma \in \Sigma_{\eta t v}^-} y(\sigma) + z(v, a_{\eta_{t-1}}) = \sum_{\sigma \in \Sigma_{\eta t v}^+} y(\sigma) + z(v, a_{\eta_t}), \quad \forall v \in \mathcal{V}, a_{\eta_{t-1}}, a_{\eta_t} \in \mathcal{A}^H \tag{4}$$

$$\sum_{v \in \mathcal{V}} \text{len}(v) \left(\sum_{\sigma \in \Sigma_{\eta t v}^-} y(\sigma) + z(v, a_{\eta_{t-1}}) \right) \leq Q_{a_{\eta_t}}, \quad \forall a_{\eta_t} \in \mathcal{A}^H \tag{5}$$

$$y(\sigma) \in \{0, 1\} \quad \forall \sigma \in \Sigma \quad (6)$$

$$\xi(k) \in [0, 1] \quad \forall k \in \mathcal{K}^P \quad (7)$$

$$\zeta(k) \in \{0, 1\} \quad \forall k \in \mathcal{K}^F \quad (8)$$

$$z(v, a_{\eta_t}) \geq 0 \quad \forall v \in \mathcal{V}, \quad a_{\eta_t} \in \mathcal{A}^H \quad (9)$$

$$B(v) \geq 0, \text{ integer} \quad \forall v \in \mathcal{V} \quad (10)$$

where $G^s(\boldsymbol{\chi})$ is the optimal value of the second-stage problem, and $\boldsymbol{\chi}$ is the vector of the first-stage decisions.

$$\begin{aligned} G^s(\boldsymbol{\chi}) = \min & \sum_{a_{l_i}(\sigma) \in \mathcal{A}^M} \sum_{k \in \mathcal{K}} c_i(\gamma(k), v(\sigma)) x(k, a_{l_i}(\sigma))^s + \sum_{a_{\eta_t} \in \mathcal{A}^H} \sum_{k \in \mathcal{K}} c(\eta, \gamma(k)) x^{\text{hold}}(k, a_{\eta_t})^s \\ & + \sum_{a_{l_i}(\sigma) \in \mathcal{A}^M} \sum_{k \in \mathcal{K}} \kappa(\eta_i(\sigma), \gamma(k)) (x^{\text{in}}(k, a_{l_i}(\sigma))^s + x^{\text{out}}(k, a_{l_i}(\sigma))^s) \end{aligned} \quad (11)$$

Subject to

$$\begin{aligned} & x^{\text{hold}}(k, a_{\eta_t})^s + \sum_{a_{l_i}(\sigma) \in \mathcal{A}^M, (\eta_{i-1}=\eta, \alpha(\eta_{i-1})=t)} x^{\text{in}}(k, a_{l_i}(\sigma))^s \\ & = \begin{cases} d(k), & \forall k \in \mathcal{K}^R, \eta = O(k), t = \alpha(k) \\ \xi(k)d(k), & \forall k \in \mathcal{K}^P, \eta = O(k), t = \alpha(k) \\ \zeta(k)d(k), & \forall k \in \mathcal{K}^F, \eta = O(k), t = \alpha(k) \end{cases} \end{aligned} \quad (12)$$

$$\sum_{\alpha(k) < t \leq \beta(k)} \sum_{a_{l_i}(\sigma) \in \mathcal{A}^M, (\eta_i=\eta, \beta(\eta_i)=t)} x^{\text{out}}(k, a_{l_i}(\sigma))^s = \begin{cases} d(k), & \forall k \in \mathcal{K}^R, \eta = D(k), \\ \xi(k)d(k), & \forall k \in \mathcal{K}^P, \eta = D(k) \\ \zeta(k)d(k), & \forall k \in \mathcal{K}^F, \eta = D(k) \end{cases} \quad (13)$$

$$\begin{aligned} & x^{\text{hold}}(k, a_{\eta_{t-1}})^s + \sum_{a_{l_i}(\sigma) \in \mathcal{A}^M, (\eta_i=\eta, \beta(\eta_i)=t)} x^{\text{out}}(k, a_{l_i}(\sigma))^s - x^{\text{hold}}(k, a_{\eta_t})^s \\ & - \sum_{a_{l_i}(\sigma) \in \mathcal{A}^M, ((\eta_{i-1}=\eta, \alpha(\eta_{i-1})=t)} x^{\text{in}}(k, a_{l_i}(\sigma))^s = 0 \\ & \forall (\eta, t) \neq (O(k), \alpha(k)) \forall \eta \neq D(k), \forall k \in \mathcal{K} \end{aligned} \quad (14)$$

$$x^{\text{in}}(k, a_{l_i}(\sigma))^s - x(k, a_{l_i}(\sigma))^s = 0, \forall a_{l_i}(\sigma) \in \mathcal{A}^M, \eta_{i-1}(\sigma) = O(\sigma), k \in \mathcal{K} \quad (15)$$

$$x(k, a_{l_i}(\sigma))^s - x^{\text{out}}(k, a_{l_i}(\sigma))^s = 0, \quad \forall a_{l_i}(\sigma) \in \mathcal{A}^M, \eta_i(\sigma) = D(\sigma), k \in \mathcal{K} \quad (16)$$

$$\begin{aligned} & x(k, a_{l_{i-1}}(\sigma))^s - x^{\text{out}}(k, a_{l_{i-1}}(\sigma))^s + x^{\text{in}}(k, a_{l_i}(\sigma))^s - x(k, a_{l_i}(\sigma))^s = 0, \\ & \forall \sigma \in \Sigma, \eta_{i-1} \neq O(\sigma), \eta_i \neq D(\sigma), k \in \mathcal{K} \end{aligned} \quad (17)$$

$$\omega \sum_{k \in \mathcal{K}} x(k, a_{l_i}(\sigma))^s \leq \text{cap}_w(v(\sigma))y(\sigma), \quad \forall \sigma \in \Sigma, a_{l_i}(\sigma) \in \mathcal{A}^M \quad (18)$$

$$\sum_{k \in \mathcal{K}} x(k, a_{l_i}(\sigma))^s \leq \text{cap}_{vol}(v(\sigma))y(\sigma), \quad \forall \sigma \in \Sigma, a_{l_i}(\sigma) \in \mathcal{A}^M \quad (19)$$

$$\theta(v(\sigma))\omega \sum_{k \in \mathcal{K}} x(k, a_{l_i}(\sigma))^s + dh^-(v(\sigma)) \leq p_{a_{l_i}(\sigma)}^s \forall \sigma \in \Sigma, a_{l_i}(\sigma) \in \mathcal{A}^M \quad (20)$$

$$x(k, a_{l_i}(\sigma))^s \geq 0 \quad \forall k \in \mathcal{K}, a_{l_i}(\sigma) \in \mathcal{A}^M \quad (21)$$

$$x^{\text{out}}(k, a_{l_i}(\sigma))^s \geq 0 \quad \forall k \in \mathcal{K}, a_{l_i}(\sigma) \in \mathcal{A}^M \quad (22)$$

$$x^{\text{in}}(k, a_{l_i}(\sigma))^s \geq 0 \quad \forall k \in \mathcal{K}, a_{l_i}(\sigma) \in \mathcal{A}^M \quad (23)$$

$$x^{\text{hold}}(k, a_{\eta_t})^s \geq 0 \quad \forall k \in \mathcal{K}, a_{\eta_t} \in \mathcal{A}^H \quad (24)$$

The mathematical formulation is presented in Eqs. (1)-(24). The objective function (Eq. (1)) maximizes profit by considering the revenue from satisfying demand across three categories, represented by the first three terms. From this, we subtract the costs associated with first-stage decisions, such as resource management and service selection, as well as the expected second-stage costs, which are captured in the final term—specifically, the cost of demand flow over the scheduled network design across all scenarios Eq. (11). Constraints (2)-(5) ensure the feasibility of the first-stage decisions, while constraints (6)-(10) define the boundaries for these decisions. Equation (2) computes the number of vessels used in the plan as the sum of vessels idling in ports or moving between them performing services. Due to the resource management concerns and the resulting circular vessel routes, $B(v)$ is the same at all periods, only the relative proportion of idle versus active vessels being different at different time periods. We therefore compute this number for the first period, i.e., $t = 0$, the set $\Lambda_{0l} = \{\sigma \in \Sigma, v(\sigma) = v | (\alpha_n(\sigma) \bmod T) < \beta_0(\sigma) \text{ and } \beta_0(\sigma) \geq 0\} \subseteq \Sigma$ containing all services, of the appropriate vessel type, that operate one of its legs during the first period. Constraints (3) enforce the fleet size for each vessel type, while Equations (4) are the so-called design-balance constraints, enforcing the vehicle-flow conservation at terminals (the number of services and vessels entering a node equals the number exiting the node), where sets $\Sigma_{\eta tv}^-$ and $\Sigma_{\eta tv}^+$

$$\Sigma_{\eta tv}^- = \{\sigma \in \Sigma \mid D(\sigma) = \eta, \beta_n(\sigma) = t, v(\sigma) = v\} \quad (25)$$

$$\Sigma_{\eta tv}^+ = \{\sigma \in \Sigma \mid O(\sigma) = \eta, \alpha_0(\sigma) = t, v(\sigma) = v\} \quad (26)$$

group the services with a vessel type v that arrive at their destination or depart from their origin, respectively. Finally, Constraints (5) enforce the terminal berthing capacity at each time period. Constraints (12)–(20) ensure the feasibility of the second-stage decisions. In particular, flow conservation is enforced for containers of all shipper types at their origins, destinations, and intermediate terminals through (12), (13), and (14), respectively. Similarly, container flow is conserved along each service route at its origin,

destination, and intermediate legs via (15), (16), and (17). Eqs. (18) and (19) ensure that the weight and volume of demand $k \in \mathcal{K}$ transported by service σ on leg i do not exceed the nominal capacity of the vessel v operating service σ . Constraint (20) ensures that the vessel's draught does not exceed the available water level under each scenario $s \in S$. Finally, constraints (21)–(24) define the non-negativity domains of the associated recourse variables.

4.2.2 The 2-SPSDM model

To formulate the 2-SPSDM model, we define five sets of decision variables. The first-stage variables are identical to those in the 2-SPDFA model and include service selection, shipper selection, resource management, and the demand flow variables. The second-stage variable $\lambda(k, a_{l_i}(\sigma))^s \geq 0$ represents the volume of demand $k \in \mathcal{K}$ that cannot be transported on leg i of service σ under water-level scenario $s \in S$.

Mathematical formulation. The problem is defined as follows:

$$\begin{aligned}
\max \quad & \sum_{k \in \mathcal{K}^R} \phi(k)d(k) + \sum_{k \in \mathcal{K}^P} \phi(k)\xi(k)d(k) + \sum_{k \in \mathcal{K}^F} \phi(k)\zeta(k)d(k) \\
& - \sum_{v \in \mathcal{V}} \mu(v)(F_v - B(v)) - \sum_{\sigma \in \Sigma} f(\sigma)y(\sigma) - \sum_{a_{\eta_t} \in \mathcal{A}^H} \sum_{v \in \mathcal{V}} h(\eta, v)z(v, a_{\eta_t}) \\
& - \sum_{a_{l_i}(\sigma) \in \mathcal{A}^M} \sum_{k \in \mathcal{K}} c_i(\gamma(k), v(\sigma))x(k, a_{l_i}(\sigma)) - \sum_{a_{\eta_t} \in \mathcal{A}^H} \sum_{k \in \mathcal{K}} c(\eta, \gamma(k))x^{\text{hold}}(k, a_{\eta_t}) \\
& - \sum_{a_{l_i}(\sigma) \in \mathcal{A}^M} \sum_{k \in \mathcal{K}} \kappa(\eta_i(\sigma), \gamma(k)) (x^{\text{in}}(k, a_{l_i}(\sigma)) + x^{\text{out}}(k, a_{l_i}(\sigma))) - \sum_{s \in S} \rho^s G^s(\boldsymbol{\chi})
\end{aligned} \tag{27}$$

Subject to

constraints defined in Eqs. (2)–(10), (12)–(19), and (21)–(24).

where $G^s(\boldsymbol{\chi})$ is the optimal value of the second-stage problem, and $\boldsymbol{\chi}$ is the vector of the first-stage decisions.

$$G^s(\boldsymbol{\chi}) = \min b(k) \sum_{a_{l_i}(\sigma) \in \mathcal{A}^M} \sum_{k \in \mathcal{K}} \lambda(k, a_{l_i}(\sigma))^s \tag{28}$$

Subject to:

$$\begin{aligned}
\theta(v(\sigma))\omega \sum_{k \in \mathcal{K}} (x(k, a_{l_i}(\sigma)) - \lambda(k, a_{l_i}(\sigma))^s) + dh^-(v(\sigma)) &\leq p_{a_{l_i}(\sigma)}^{I^s}, \quad \forall \sigma \in \Sigma, a_{l_i}(\sigma) \in \mathcal{A}^M \\
\lambda(k, a_{l_i}(\sigma))^s &\leq x(k, a_{l_i}(\sigma)), \quad \forall k \in \mathcal{K}, a_{l_i}(\sigma) \in \mathcal{A}^M
\end{aligned} \tag{29}$$

$$\lambda(k, a_{l_i}(\sigma))^s \geq 0, \quad \forall k \in \mathcal{K}, a_{l_i}(\sigma) \in \mathcal{A}^M \quad (31)$$

The 2-SPSDM model, formulated in Eqs. (27)–(31), aims to maximize net profit. The first three terms of the objective function represent revenues generated from regular shippers, partial-spot shippers, and full-spot shippers, respectively. The first-stage costs include both fixed and variable components. Fixed costs comprise penalties for unused vessels, fixed costs for establishing and operating services, and the cost of vessels idling at ports while awaiting their next departure. Variable costs are associated with the distribution of demand flows, including moving, holding, loading, and unloading containers. These costs are driven by tactical decisions made prior to the realization of water-level scenarios. The final term in the objective function represents the expected second-stage cost, specifically penalties incurred for unmet demand across all scenarios, each corresponding to a potential water-level realization (Eq. (28)). The model constraints ensure the feasibility of the first-stage decisions. These include flow conservation constraints (Eqs. (12), (13), (14)), flow balance across services (Eqs. (15), (16), (17)), service capacity limitations (Eqs. (18), (19)), and constraints on fleet size, fleet balance, and terminal berthing capacity (Eqs. (2), (3), (4), (5)). The remaining constraints define the boundaries of first-stage decisions. It is important to note that the structure of these constraints remains consistent with the previous model; however, in the present formulation, they are defined independently of any scenario. Second-stage feasibility is ensured through constraints (29), (30), and (31). Specifically, constraint (29) ensures that the vessel’s draught does not exceed the available water level for each scenario $s \in S$. Constraints (30) and (31) define the feasible region for second-stage decisions.

5 Solution Approach

In our proposed models, the effects of water level uncertainty are captured using a set of scenarios that represent how service capacities vary randomly. This scenario set must be representative, providing a sufficiently high level of confidence in capturing the variability of observed service capacity values. Generating such a comprehensive set may require a large number of scenarios and, consequently, the incorporation of substantial contextual information into the optimization process to achieve a detailed and accurate representation of the uncertainty. Therefore, processing and utilizing this information effectively in decision-making can place significant demands on computational resources. Incorporating all scenarios into stochastic WL-SSND-RRM models often renders them numerically intractable, requiring impractical amounts of time and computational power to solve them. This dual challenge—managing the large volume of information while ensuring computational feasibility—represents significant obstacles to the practical application of these models. To address this general challenge, we adopt the two-step approach introduced by Hewitt et al. (2022) and apply it to our problem setting. Details of both steps are presented in the following subsections.

5.1 Step 1: Scenario Clustering Analysis

Given scenario set $S = \{s_1, s_2, \dots, s_n\}$, which represents possible realizations of water levels, the primary objective of this step is to interpret these scenarios in terms of the decisions or plans they induce. This analysis aims to better understand how each scenario influences the decisions to be made. We begin by considering the vector of first-stage decisions \mathcal{X} , as introduced in Section 4.1 for the 2-SPSDM model and in Section 4.2 for the 2-SPDFA model. These decisions are constrained by a feasible region $\mathcal{F}_{\mathcal{X}}$, which includes all possible solutions that satisfy the model's constraints, as specified in their respective sections. To assess the performance of these decisions, we solve a restricted single-scenario version of the 2-SPSDM and 2-SPDFA models, using the objective function, g , where $g : \mathcal{F}_{\mathcal{X}} \times S \rightarrow \mathbb{R}^+$. In this context, $g(\chi, s_j)$ represents the objective value (total profit) for a given combination of first-stage decisions $\chi \in \mathcal{F}_{\mathcal{X}}$ and scenario $s_j \in S$. Among all feasible first-stage decisions in $\mathcal{F}_{\mathcal{X}}$, the optimal first-stage decision for a specific scenario s_j is denoted by $\chi_{s_j}^*$ and is defined as: $\chi_{s_j}^* = \arg \max_{\chi \in \mathcal{F}_{\mathcal{X}}} g(\chi, s_j)$. The opportunity cost is then defined as $\theta(s_i | s_j) = g(\chi_{s_j}^*; s_j) - g(\chi_{s_i}^*; s_j)$, where $\theta(s_i | s_j)$ measures the dissimilarity between decisions optimized for scenario s_i when evaluated under scenario s_j . Since each scenario s_i is compared against every other scenario s_j ($j \neq i$), the *opportunity cost matrix* Θ is formed, where each entry $\theta(s_i | s_j)$ reflects the relative performance of the decisions across different scenarios. However, the opportunity cost can be *asymmetric* ($\theta(s_i | s_j) \neq \theta(s_j | s_i)$), and this asymmetry can introduce biases in the clustering process. To address this, we construct a *symmetric dissimilarity matrix* D by summing the opportunity costs: $D(s_i, s_j) = \theta(s_i | s_j) + \theta(s_j | s_i)$. This symmetric matrix provides a balanced and mutual measure of dissimilarity between any two scenarios, ensuring its suitability for clustering. We use the symmetric matrix D and apply the K-means Spectral Clustering method, as described in (Hewitt et al., 2022), to form clusters of scenarios based on their decision impact. This process can be approached in two ways: one that seeks to minimize dissimilarity among grouped scenarios and another that aims to maximize dissimilarity. To clearly distinguish between these two approaches, we introduce the following notations.

- Let C_k^{\min} , $k = 1, \dots, K^{\min}$, represent the set of clusters obtained by minimizing dissimilarity. This approach allows us to form clusters of scenarios that lead to similar decisions, thereby reducing redundancy in the scenario set and simplifying it while minimizing information loss.
- Let C_k^{\max} , $k = 1, \dots, K^{\max}$, represent the set of clusters obtained by maximizing dissimilarity. This approach enables us to cluster scenarios that result in different decisions, ensuring that each cluster captures diverse, decision-relevant information, thus preserving the richness of uncertainty.

Notably, the number of clusters is determined using the Silhouette Score (Rousseeuw, 1987), a widely used metric for assessing clustering quality.

5.2 Step 2: Bounds Computation

In this section, we present the second step of the solution approach, which focuses on leveraging the scenario clusters obtained in the previous step to apply a series of bounding techniques. These techniques, based on the use of scenario clusters, provide both:

1. A set of alternative feasible solutions, whose objective function values define alternative lower bounds for the stochastic models under consideration.
2. An upper bound value, which serves as a bounding certificate for evaluating the optimality gaps of the obtained feasible solutions.

The *Medoid Lower Bound* is derived using C_k^{\min} , $k = 1, \dots, K^{\min}$. This approach is based on the idea of using clusters to reduce the size of the scenario set utilized in the optimization process, thereby retaining a representative subset of scenarios while keeping computational complexity manageable. This representative subset, denoted as M , consists of the *medoid* of each cluster, denoted as M_k . The medoid M_k is defined as the scenario within each cluster that has the smallest average dissimilarity to all other scenarios in that cluster. The probability assigned to each medoid, denoted as p_{M_k} , is calculated as the sum of the probabilities of all scenarios within the same cluster. Formally, if ρ^{s_i} represents the probability of scenario s_i , then the probability of the medoid for cluster C_k^{\min} is given by $p_{M_k} = \sum_{s_i \in C_k^{\min}} \rho^{s_i}$. Given the representative scenarios (medoids) and their corresponding probabilities, the optimal decision $\bar{\chi}$ for the 2-SPSDM and 2-SPDFA models is then determined by maximizing the expected profit over the representative scenario set. Formally, this is given by $\bar{\chi} = \arg \max_{\chi \in \mathcal{F}_X} \sum_{M_k \in M} g(\chi, M_k) \times p_{M_k}$. The lower bound is then computed by evaluating this solution over all scenarios, expressed as $\sum_{s_i \in S} g(\bar{\chi}; s_i) \rho^{s_i}$, where $g(\bar{\chi}; s_i)$ represents the objective value for scenario $s_i \in S$ given $\bar{\chi}$.

In contrast, the other two bounds, namely the *Cluster Lower Bound* and the *Cluster Upper Bound*, are derived using clusters C_k^{\max} , $k = 1, \dots, K^{\max}$. For these bounds, the clustering process partitions the original scenario set in such a way as to produce a representative subset within each cluster that preserves the overall structure of the scenario set. For each cluster C_k^{\max} , we define the total probability $P_{C_k^{\max}}$ as: $P_{C_k^{\max}} = \sum_{s_i \in C_k^{\max}} \rho^{s_i}$. The optimal decision $\bar{\chi}_{C_k^{\max}}$ for each cluster is then determined by maximizing the expected profit over the scenarios within that cluster, with scenario probabilities normalized within the cluster. Formally, this is given by: $\bar{\chi}_{C_k^{\max}} = \arg \max_{\chi \in \mathcal{F}_X} \sum_{s_i \in C_k^{\max}} g(\chi, s_i) \times \frac{\rho^{s_i}}{P_{C_k^{\max}}}$. The upper bound is calculated by aggregating the results from all clusters and is represented as $\sum_{k=1}^{K^{\max}} \sum_{s_i \in C_k^{\max}} g(\bar{\chi}_{C_k^{\max}}; s_i) \times \rho^{s_i}$. The lower bound is derived by comparing results across clusters and is represented as $\min_{k=1, \dots, K^{\max}} \sum_{s_i \in C_k^{\max}} g(\bar{\chi}_{C_k^{\max}}; s_i) \times \rho^{s_i}$.

6 Numerical experiments

In this section, we present a set of computational experiments and analyses. These serve two main objectives. The first is to assess the performance of the proposed solution approach and validate its ability to balance computational efficiency and solution quality. The second objective is to explore the impact of explicitly considering water level uncertainty on the tactical planning process. Section 6.1 presents the characteristics of the instances used in the computational experiments, which are based on realistic case scenarios. We then begin by evaluating the performance of the solution approach in Section 6.2. We also investigate several critical managerial questions in the remainder of the section. Section 6.3 examines how incorporating the effects of stochastic water levels within the tactical planning process can improve upon the traditional approach of relying on predictions and applying a deterministic optimization model based on those predictions. In this study, we further propose two recourse strategies to account for the impact of reacting to randomly changing water levels. These strategies both aim to generate more profitable tactical plans (i.e., plans that can lead to higher profits for carriers by better adapting to uncertainty) yet they may result in plans with different structural characteristics. Section 6.4 assesses the effectiveness of these alternative recourse strategies. Finally, Section 6.5 investigates the impact of water level fluctuations on metrics such as the number of vessels used, consolidation, cost, and profit.

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.

6.1 Test instances

The test instances are designed to reflect operational conditions commonly observed in Canadian and European inland waterway systems, where major ports are aligned along linear corridors. Representative examples include the Saint Lawrence River corridor (Montréal–Québec City–Sept-Îles–Port-Cartier) and the Rhine River system (Upper, Middle, and Lower Rhine). Accordingly, our network consists of four connected major terminals (Figure 2a) with berthing capacities between 700 and 1000 meters, enabling multiple vessels to operate simultaneously. Costs associated with terminal operations are considered as follows: the loading and unloading of containers incur a cost of 2 units per operation, terminal activities such as storage are charged to carriers at a rate of 3 units per container per period, and vessel holding costs at terminals are set at 5 units per period. Water levels are monitored at critical river points, ranging between 150 and 350 cm, consistent with historical observations reported in Christodoulou et al. (2020). To capture realistic variability, we adopt the copula-based method of Kaut (2014), which preserves observed marginal distributions and correlations across terminals. A Beta(2,5) distribution is used to reproduce the right-skewed behavior typically seen in drought con-

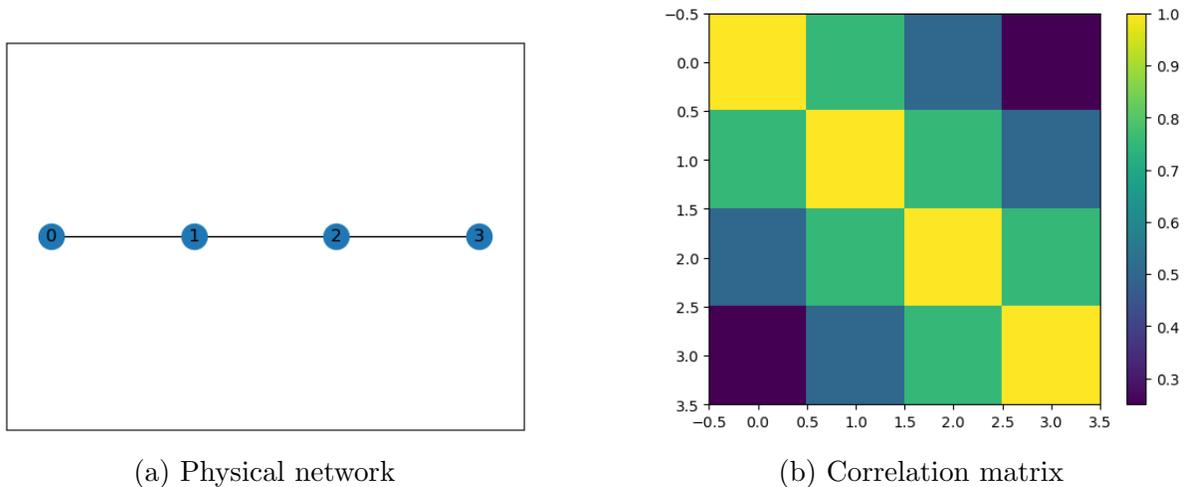


Figure 2: Physical network and correlation matrix

ditions, with distance-based correlations (0.75 for adjacent nodes, 0.25 for the farthest). Scenarios are then sampled to represent distinct realizations of water-level conditions, each assigned equal probability.

The schedule is defined over a one-week horizon divided into 14 half-day periods, consistent with operational practice in inland waterway transportation in which departures and arrivals typically occur in morning and afternoon slots. Potential services are designed over the linear network and include all possible non-stop and one-stop routes between the four terminals. In total, 6 non-stop and 4 one-stop routes are considered, each available in both directions, yielding 20 routes. These routes are assigned selected departure times within the 14-period weekly cycle, resulting in 180 potential services—sufficient to ensure diversity in the service network. Each service can be operated by either a large or a small vessel, producing 360 service alternatives overall. Service durations depend on the travel time between terminals (ranging from a minimum of one day to a maximum of two days) and the time spent at each stop (one half-day period for every origin, destination, or intermediate terminal). Large vessels (200 m, 500 t, 50 TEUs, draught 135 cm) and small vessels (100 m, 200 t, 20 TEUs, draught 100 cm) are both assumed to be able to navigate the entire network when empty, even under minimum water-level conditions. We generate five distinct demand sets, totaling 298 demands (55–66 per set), to ensure diversity and reduce instance-specific bias. Demands are defined on the linear network using random origin–destination pairs, with volumes between 5 and 25 TEU units. Each demand is characterized by shipper type (regular, partial spot, or full spot) and delivery-time requirements (express or standard), where express shipments are assigned unit fares twice those of standard shipments to reflect real-world tariff structures. Pickup time windows are uniformly distributed across the 14 half-day periods in the one-week horizon. Express deliveries require 1–3 periods, whereas standard deliveries require twice the express time for the same route.

Finally, we introduce the instance categories labeled *Instance-1* through *Instance-7*. Each category consists of three demand instances selected from the five demand sets described above and includes the corresponding recourse parameters: the penalty cost $b(k)$ in the 2-SPSDM model (ranging from 5 to 60 units), and the per-leg transportation cost $c_i(\gamma(k), v(\sigma))$ in the 2-SPDFA model (ranging from 3 to 7 units). Using a single instance could lead to biased or unrepresentative results due to specific parameter realizations. By averaging performance across multiple instances, we obtain a more robust comparison that mitigates the effects of randomness in the data.

6.2 Performance of the solution approach

Preliminary experiments revealed a clear distinction in tractability between the two models. All instances of the 2-SPSDM model could be solved within a reasonable time using Gurobi, whereas none of the 2-SPDFA instances reached optimality within acceptable computational limits. These observations motivated the adoption of a scenario clustering-based analysis to evaluate model performance under limited computational resources. Importantly, the objective of this analysis is not to establish superiority over general-purpose solvers such as Gurobi, but rather to assess the ability of the proposed approach to provide high-quality bounds and meaningful performance insights in settings where solving the proposed stochastic models to optimality is computationally impractical or not attainable.

This section describes the experimental setup used to assess the performance of the proposed solution approach. The experimental analysis follows the procedure outlined in Section 5 and consists of two main steps. In Step 1, the dissimilarity matrix is constructed and the clustering analysis is performed. In Step 2, several bounds are computed using both the medoid-based and the cluster-based approaches. The total computational time required for the solution approach is recorded, including the time spent on computing the dissimilarity matrix, performing the clustering, and calculating the different bounds (namely, the medoid lower bound as well as the cluster-based lower and upper bounds). The recorded runtimes indicate that, while the cluster-based approach is computationally efficient, it remains slower than the medoid-based alternative. Focusing first on the 2-SPSDM model, the experiments are designed to evaluate the behavior of the proposed solution approach for large-scale tactical planning problems under limited computational resources. Although optimal solutions can be obtained for this model by directly solving it with Gurobi when sufficient time is available, we intentionally restrict Gurobi's maximum runtime to match the average runtime of the cluster-based approach. This setup allows us to assess how the proposed scenario-based approach performs when the model remains solvable in principle, but computational resources are constrained. For the more complex 2-SPDFA model, optimal solutions cannot be obtained using Gurobi within reasonable computational limits. As a result, the proposed scenario-based approach serves as the primary means of analyzing model performance. Gurobi is nevertheless allocated

a substantially larger time limit of 7200 seconds to serve as a reference and to confirm that the problem remains computationally intractable even when additional computational effort is allowed. This experimental design highlights the ability of the proposed approach to produce informative bounds and performance insights in such challenging settings.

Table 1 and Table 2 present the performance evaluation of the proposed solution approach on the 2-SPSDM and 2-SPDFA models, respectively, across instance categories labeled Instance-1 through Instance-7. The "Gap (%)" columns show the average optimality gaps (in percentage) achieved by Gurobi, the Medoid method, and the Cluster method over the three sub-instances associated with each main instance, indicating how closely each method approximates the optimal solution. The optimality gap for the cluster-based approach is calculated as the relative difference between the cluster upper bound and the cluster lower bound obtained through this method. For the medoid-based approach, the optimality gap is determined by calculating the relative difference between the medoid lower bound and the cluster upper bound. The "Time (s)" column under the Cluster method represents the average computational time (in seconds) required to apply the cluster method to solve each problem instance. Rows labeled "Mean," "Median," "Std Dev," "Max," and "Min" summarize the performance metrics across all instances, providing insights into the variability and efficiency of each method.

Table 1: Performance Evaluation of the Proposed Solution Approach on the 2-SPSDM Model

Instance	Gurobi	Medoid	Cluster	
	Gap (%)	Gap (%)	Gap (%)	Time (s)
Instance-1	4.99	2.67	1.71	19.97
Instance-2	17.3	0.51	1.56	22.05
Instance-3	2.74	5.29	2.74	25.41
Instance-4	3.77	7.34	2.44	23.81
Instance-5	4.72	5.15	0.97	21.15
Instance-6	10.1	9.50	2.73	20.23
Instance-7	5.74	1.22	2.02	20.74
Mean	7.05	4.38	2.03	21.34
Median	4.99	5.15	2.02	21.15
Std Dev	5.18	3.14	0.60	2.07
Max	17.30	9.50	2.74	25.41
Min	2.74	0.51	0.97	19.97

The performance evaluation for the 2-SPSDM model (Table 1) shows that the Cluster-based approach is able to generate tight and stable bounds within short computation times. Across all instances, the observed optimality gaps remain consistently small, with an average value of 2.03% and a maximum of 2.74%, indicating that the proposed

approach provides reliable and informative performance indicators under limited computational resources. Gurobi is used solely as a reference under the same restricted computational conditions. Under this setting, Gurobi exhibits larger and more variable gaps, with an average gap of 7.05% and a maximum gap of 17.30%. These results reflect the difficulty of extracting high-quality information from the stochastic model when computation time is limited. When the time restriction is removed, Gurobi is able to solve the instances to optimality, confirming that the model is solvable in principle and reinforcing the relevance of evaluating alternative approaches in resource-constrained computational settings. The Medoid-based approach also produces useful bounds in several instances; however, its performance shows greater variability across cases, with an average gap of 4.38%, a maximum gap of 9.50% (Instance-6), and a minimum gap of 0.51% (Instance-2). This variability suggests that while the Medoid method can be effective, its performance is less consistent compared to the Cluster method. In particular, the medoid bound is based on applying scenario reduction, which effectively eliminates information when solving the optimization problem. This reduction simplifies the problem but also introduces a trade-off, as removing scenarios inherently limits the model’s ability to fully capture the variability in the uncertainty representation. In contrast, the cluster bounds leverage the full original scenario set, preserving all information contained within it.

Table 2: Performance Evaluation of the Proposed Solution Approach on the 2-SPDFA Model

Instance	Gurobi	Medoid	Cluster	
	Gap (%)	Gap (%)	Gap (%)	Time (s)
Instance-1	5.88	7.53	0.19	140.05
Instance-2	11.1	6.70	0.30	104.55
Instance-3	5.43	2.83	1.19	300
Instance-4	42.9	18.53	11.54	255.88
Instance-5	3.03	6.40	1.24	399.58
Instance-6	9.38	10.50	0.53	203.65
Instance-7	6.25	4.22	0.01	268.70
Mean	11.28	8.24	2.29	238.21
Median	6.25	6.70	0.53	255.88
Std Dev	13.53	5.32	4.11	109.91
Max	42.90	18.53	11.54	399.58
Min	3.03	2.83	0.01	104.55

Table 2 reports the performance evaluation of the proposed solution approach for the 2-SPDFA model across multiple instances. The results summarize the quality of the bounds and computation times obtained using the Medoid-based and Cluster-based approaches, with Gurobi used solely as a reference under a fixed runtime of 7200 seconds. For this more complex model, the Cluster-based approach is able to generate informative and stable bounds across all instances. The observed gaps remain relatively small, with

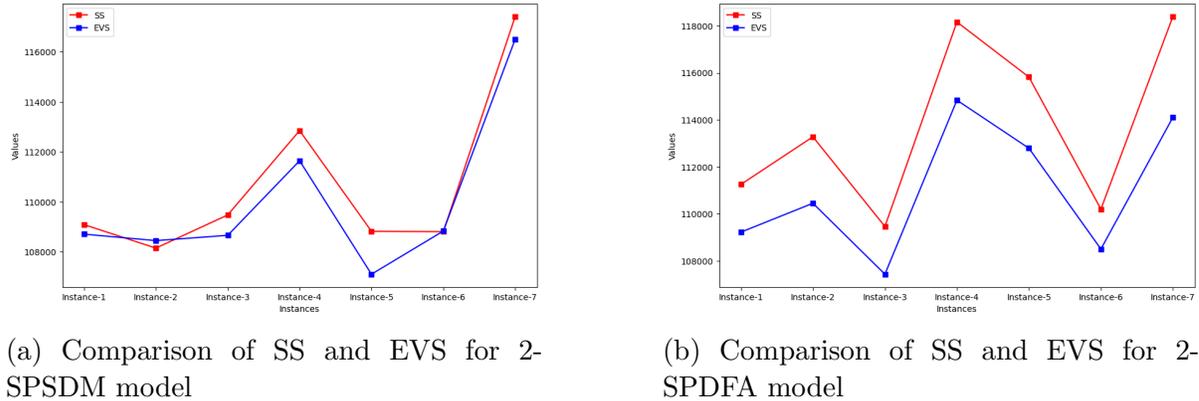


Figure 3: Value of the Stochastic Solution

an average value of 2.29%, a minimum gap of 0.01% (Instance-7), and a maximum gap of 11.54% (Instance-4). These results indicate that the proposed approach can effectively extract meaningful performance information from the model in settings where direct solution methods face substantial computational challenges. When used as a reference with a substantially longer runtime, Gurobi still exhibits wide variability in the resulting gaps, ranging from 3.03% (Instance-5) to 42.9% (Instance-4), with an average gap of 11.28%. This behavior confirms the intrinsic computational difficulty of the 2-SPDFA model and motivates the use of alternative approaches capable of providing reliable bounds within practical time frames. The Medoid-based approach also produces useful bounds for the 2-SPDFA model, achieving an average gap of 8.24% and substantially reducing the variability of the results, as reflected by a decrease in standard deviation from 13.53 to 5.32. This reduction in variability is consistent with the scenario reduction mechanism underlying the medoid approach, which simplifies the problem and improves computational stability, albeit at the cost of discarding part of the uncertainty representation.

6.3 Importance of modeling uncertainty

The Value of Stochastic Solution (VSS) is a measure used in stochastic programming to quantify the overall benefit of using a stochastic model over a deterministic model when making decisions under uncertainty. VSS measures the expected improvement in performance—here profit maximization—achieved by basing decisions on a stochastic model rather than a deterministic approximation that instead uses a prediction for the stochastic parameters.

To calculate the VSS, two key values need to be determined. *Expected Value Solution (EVS)*: This is the outcome of applying the solution from a deterministic model that does not account for uncertainty and instead uses expected values of uncertain water levels. *Stochastic Solution (SS)*: This represents the best upper bound obtained using the proposed solution approach for the stochastic model, which incorporates a set of possible

scenarios and their associated probabilities for water levels. Figure 3a and Figure 3b present the VSS results for the 2-SPSDM and 2-SPDFA models, respectively. In these figures, the blue line with square markers represents the EVS, derived from a deterministic model that uses average water levels without accounting for uncertainty. The red line with square markers represents the SS, which explicitly accounts for random fluctuations in water levels.

Figure 3a, we observe that, in most instances (except for Instance 2), the SS performs better or similarly to the EVS. In Instance 2, the SS performs marginally worse because the stochastic model is not solved to optimality for one of the three instances included in Instance 2. In Instance 6, both models exhibit nearly identical performance, while in Instances 4, 5, and 7, the SS significantly outperforms the EVS, with the greatest improvement seen in Instance 5. This pattern highlights the importance of solving the 2-SPSDM stochastic model when the infeasibilities in the tactical plan due to varying water levels are penalized by applying a cost to the amount of demand that cannot be transported. A similar pattern is observed in Figure 3b, which compares the SS and EVS values for the 2-SPDFA model. As shown, the SS consistently outperforms the EVS across all instances, underscoring the effectiveness of the 2-SPDFA model in providing a tactical plan that is better suited to deal with uncertainty in water levels.

6.4 Model Evaluation

This section includes a set of comparative tests designed to evaluate the performance of the proposed models for the WL-SSND-RRM problem in a stochastic environment. The objective is to assess the models' ability to generate profitable tactical plans and to analyze how different recourse strategies lead to plans with distinct structural characteristics. To provide a more accurate basis for evaluating the obtained solutions with a higher level of confidence as to their actual value (the so-called ground truth), we define a Ground Truth set consisting of 500 scenarios sampled from the same distribution described in Section 6.1. Each test begins by independently solving each model to obtain its tactical plan. Subsequently, the performance of each solution is evaluated by applying it either to a different model or within its own framework using the Ground Truth scenarios. To facilitate this evaluation, we define two types of benchmarks: Cross-Model Benchmark and Internal Model Benchmark:

- **Cross-Model Benchmark:** This benchmark includes the Deterministic-in-2-SPSDM, Deterministic-in-2-SPDFA, and 2-SPSDM-in-2-SPDFA frameworks. These frameworks involve fixing certain decisions derived from one model (e.g., the deterministic solution or 2-SPSDM) as input to another stochastic model framework (e.g., 2-SPSDM or 2-SPDFA). The resulting restricted stochastic models are then solved by exploring the remaining decision subspace to complete the full solution using the ground truth scenarios.

- **Internal Model Benchmark:** This benchmark includes the 2-SPSDM Baseline and 2-SPDFA Baseline frameworks. In these frameworks, each model’s optimized solution is evaluated within its original framework using the ground truth scenarios, establishing a baseline for that model’s optimal performance under its specific recourse strategy.

Based on these benchmarks, four comparative tests are performed. The first two tests, $\Delta 1$: Deterministic-in-2-SPSDM vs. 2-SPSDM Baseline and $\Delta 2$: Deterministic-in-2-SPDFA vs. 2-SPDFA Baseline, are designed to validate the VSS analyses previously conducted by performing out-of-sample assessments for the two stochastic models. The third test, $\Delta 3$: 2-SPSDM-in-2-SPDFA vs. 2-SPDFA Baseline, assesses whether a tactical plan obtained using a model that applies itinerary infeasibility penalty costs can serve as a good proxy for a plan obtained using a model that applies the more involved and complex flow optimization recourse costs. Specifically, this test seeks to determine if there is additional value, from a tactical planning perspective, in solving the more complex 2-SPDFA stochastic model. The final test, $\Delta 4$: Deterministic-in-2-SPSDM vs. Deterministic-in-2-SPDFA, compares the effectiveness of two stochastic models, each defined by a distinct approach to handling the impact of randomly changing water level conditions.

Table 3: Comparison tests

Test	Benchmark	RD 1 to 2 (%)	RP 1 to 2 (%)
$\Delta 1$	1: Deterministic-in-2-SPSDM	47.21	-3.04
	2: 2-SPSDM Baseline		
$\Delta 2$	1: Deterministic-in-2-SPDFA	-25.23	-2.03
	2: 2-SPDFA Baseline		
$\Delta 3$	1: 2-SPSDM-in-2-SPDFA	-19.86	1.90
	2: 2-SPDFA Baseline		
$\Delta 4$	1: Deterministic-in-2-SPSDM	-39.73	-8.81
	2: Deterministic-in-2-SPDFA		

In Table 3, we compare the performance of two stochastic programming models across four tests ($\Delta 1$, $\Delta 2$, $\Delta 3$, and $\Delta 4$), focusing on two key metrics: Relative Difference in Recourse Decision Value (RD) and Relative Difference in Profit Value (RP). These metrics evaluate each model’s adaptability and profitability when subjected to uncertainty. The RD column represents the percentage difference in recourse decision values between the two benchmarks, illustrating each model’s effectiveness in meeting demand under variable conditions. RD in $\Delta 1$ is computed as the relative difference in the total amount of unmet demand, highlighting how each model manages demand satisfaction under uncertainty. RD in $\Delta 2$ and $\Delta 3$ capture the relative difference in the total amount of met demand, providing a direct comparison of the models’ performance in fulfilling demand.

RD in Δ_4 is calculated by first determining the total met demand in the first benchmark as the difference between total demand and unmet demand, and then comparing this value to the total met demand in the second benchmark. The RP column, on the other hand, measures the relative difference in profit between the two benchmarks.

The RD values for Δ_1 and Δ_2 highlight that when service selections and demand itineraries are planned without accounting for variability, these plans become less reliable under fluctuating water conditions. In contrast, the tactical plans obtained using the 2-SPSDM model and the 2-SPDFA model demonstrate better performance in fulfilling demand, with 47.21% less unmet demand and 25.23% more transported demand, respectively. Similarly, the RP values for Δ_1 and Δ_2 highlight the overall benefit of using the two stochastic models over a deterministic model when making decisions under uncertainty, with an improvement in profit of 3.04% and 2.03%, respectively.

Δ_3 highlights the tactical plans obtained using the 2-SPSDM model can nonetheless lead to higher profitability when itineraries are adjusted based on the revealed water levels. This can be explained by the fact that the plans obtained tend to select fewer services. Here, an average profit increase of 1.90% is observed across all results obtained. On the other hand, the plans obtained by solving the 2-SPDFA model include a higher number of services, allowing for more transportation paths when defining itineraries. However, these additional options come with higher fixed costs, which are adverse to profit maximization. Nevertheless, having access to a greater number of potential itineraries significantly decreases the amount of unmet demand, with an observed average reduction of 19.86% in serviced demands when using the 2-SPSDM model instead of the 2-SPDFA model. Second, the difference in profits can be attributed to the fact that the instances are not guaranteed to be solved to optimality. Thus, additional profit-maximizing benefits could still be gained from using the 2-SPDFA model by further improving the quality of the solutions through additional computational efforts. Lastly, this comparison illustrates a general trade-off between prioritizing profit maximization and the amount of serviced demand. Solving the 2-SPSDM model can produce tactical plans conducive to profit maximization, even when such plans are implemented in a different setting where transportation itineraries are allowed to be adjusted based on revealed water levels. However, such plans are not necessarily set up to easily adjust the flows according to varying water level conditions, leading to a reduction in the amount of serviced demand. Considering that there may be significant opportunity losses in not servicing demands, the 2-SPDFA model can provide additional value by leading to plans that accommodate more alternative demand itineraries.

In Δ_4 , the goal is to compare the efficiency of different recourse strategies in stochastic models in terms of profit maximization and the amount of serviced demand when dealing with changing water levels. As shown in the table, when the carrier adopts a more flexible recourse strategy, such as demand flow adjustment in the 2-SPDFA model, it achieves a higher profit (8.81%) and can transport significantly more demand (39.73%). These results highlight the substantial benefits of adjusting itineraries in response to water level

variations.

6.5 Water Level Impact

In this section, we analyze the performance, efficiency, and structural characteristics of solutions obtained by solving the deterministic WL-SSND-RRM model under three fixed and a priori known water-level settings—maximum, mean, and first-quarter—ranging from 150 to 350 cm, as defined by Christodoulou et al. (2020). Rather than modeling stochastic water-level uncertainty through probability distributions, these experiments rely on point forecasts of water levels, while keeping the demand set and the service network identical to those used in the previous experiments. These water levels serve as the basis for evaluating how fluctuations impact key metrics such as overall system performance, customer satisfaction, and whether any systematic structural changes emerge as the water levels vary. Our goal is to measure the relationship between water level fluctuations and the level of multi-service usage, a critical factor in barge transportation network design for minimizing operational costs and optimizing resource use. A decrease in the number of demands using multiple paths (services) to reach their destinations typically indicates reduced flexibility and efficiency in the network. Additionally, we investigate the relationship between water levels and the distribution of small and large vessels, as vessel deployment is directly influenced by water depth. Changes in water levels affect vessel draught, which in turn impacts the types of vessels that can be effectively deployed under varying conditions. We also assess the impact of water levels on operational costs, driven by changes in the network structure—specifically, the shifts in multi-service usage and vessel size. Furthermore, we investigate the proportion of spot shippers that are satisfied at each water level, as customer satisfaction directly influences the network’s profitability. We also examine capacity usage as a measure of how efficiently the deployed transport capacity is utilized under different water levels. This comprehensive analysis aims to provide a clearer understanding of how water level fluctuations affect the network’s performance, cost-efficiency, and financial outcomes.

6.5.1 Performance Indicators (PIs)

To evaluate the experimental results, we use the following performance indicators for each water level (maximum, mean, and first quarter):

Structural Impact PIs

- **Multi-path Usage:** For each solution, we record the number of services used by each demand across all instances. We then produce a histogram to show the distribution of path usage frequencies across all demands. For example, if a demand uses two paths, it is counted as such.

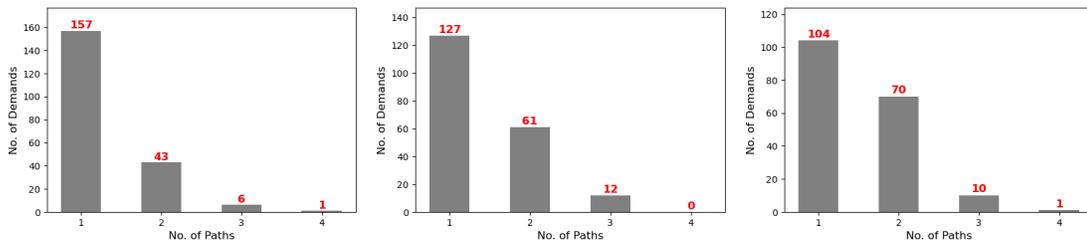
- **Vessel Size:** We record the different sizes of vessels used to transport demands for each solution and produce a histogram showing the distribution of vessel sizes. This allows us to assess how water levels influence the deployment of small and large vessels.

Performance PIs

- **Profit:** A reflection of the overall system’s performance, considering revenue obtained and total costs.
- **Demand Flow Costs:** The costs incurred for moving, loading, and unloading demands through the network.
- **Service Costs:** The expenses associated with the transport services used.
- **Shipper Satisfaction (TEUs):** Measured as the percentage of the volume of spot shippers who are fully or partially met under each water level condition.
- **Capacity Usage:** Defined as the ratio of total volume-kilometers moved to total capacity-kilometers operated.

6.5.2 Results and analysis

In the figures presented, we examine multi-path usage and vessel usage at different water levels—350 cm, 250 cm, and 200 cm—highlighting how the transportation network responds to fluctuating water conditions. These figures provide important insights into the operational challenges and adaptations the system may make as water levels change.



(a) Water level = 350 cm (b) Water level = 250 cm (c) Water level = 200 cm

Figure 4: Multi-path usage at different water levels: 350 cm, 250 cm, and 200 cm.

Figure 4 illustrates the multi-path usage at the three water levels. At 350 cm, the network satisfies 207 total demands, with 157 of those demands satisfied by a single service and 50 demands requiring multiple services. This suggests that at higher water

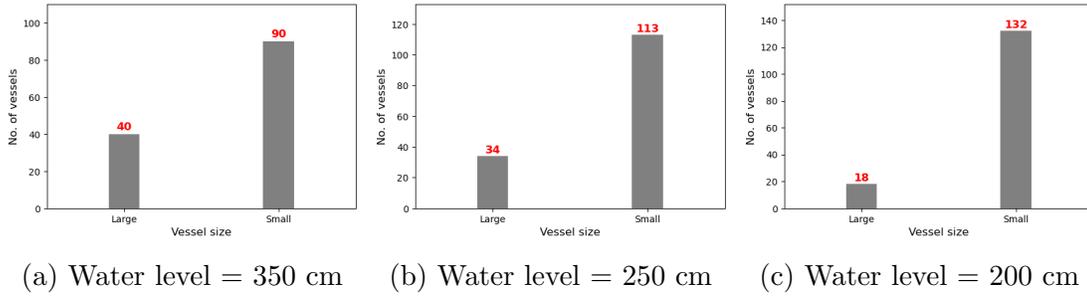


Figure 5: Vessel usage at different water levels: 350 cm, 250 cm, and 200 cm.

levels, most demands are transported via one service, showing less need for alternative services. As we move to 250 cm, the total number of satisfied demands decreases slightly to 200, and the distribution of service usage shifts. 127 demands now use a single service, while 73 demands (61 using two services, 12 using three) require more than one service to reach their destination. This increase in multi-path usage reflects the growing operational challenge as the water level drops, forcing the network to rely on more services to maintain service levels. At 200 cm, the system is under more strain, satisfying only 185 demands. The number of single-service demands decreases to 104, while 70 demands require two services, 10 demands need three, and 1 demand uses four services. In total, 81 demands rely on multiple paths, which is a 62% increase compared to the 350 cm level. This shows that as water levels decrease, the network faces greater operational pressure, leading to more fragmented routing and less efficient transportation.

Figure 5 highlights the vessel usage at the same water levels. At 350 cm, the network operates with 40 large vessels and 90 small vessels, indicating that 130 services are activated to meet demand. As the water level drops to 250 cm, the number of large vessels decreases to 34, while the number of small vessels increases to 113, resulting in 147 total services being activated. This shift shows that as the water becomes shallower, the system compensates by deploying more small vessels, which are capable of operating under these conditions but require more services due to their limited capacity. At 200 cm, this trend becomes even more pronounced, with only 18 large vessels in operation and a significant increase in the number of small vessels to 132, resulting in a total of 150 activated services—20 more than at the 350 cm water level. This reflects the severe limitations imposed by low water levels on the network’s ability to use larger vessels. Small vessels become the primary means of transportation, but their lower cargo capacity necessitates higher numbers of activated services. This leads to increased operational costs (as highlighted in Table 4).

Overall, the figures show that multi-path usage increases as water levels drop, with more demands requiring multiple routes. This indicates a more fragmented and less efficient system at lower water levels. Additionally, vessel usage indicates a clear shift towards smaller vessels as water levels decrease, leading to an increased number of activated services. The total number of services rises from 130 at 350 cm to 150 at 200

cm, reflecting the system’s need to compensate for reduced capacity by deploying more smaller vessels. This combination of increased multi-path usage and reliance on smaller vessels at lower water levels results in higher operational costs, reduced profitability, and lower shipper satisfaction, as shown in the table 4.

The table reveals noticeable variations in profit, costs, and shipper satisfaction as water levels decrease. At 350 cm, the profit is 568,140. However, as water levels drop to 250 cm, the profit decreases by 0.42% to 565,725, and at 200 cm, it drops further by 1.46% to 559,936. While the absolute change might appear small, the decline is consistent and shows that lower water levels negatively impact profitability. This reduction in profit reflects the system’s increasing inefficiency, where higher operational costs are driven by the need to compensate for reduced capacity by using more small vessels and increasing multi-path usage. This is further evidenced by the decline in capacity usage, which drops from 61.28% at 350 cm to 56.31% at 250 cm, and then to 51.47% at 200 cm. Despite deploying a higher number of vessels and services, the system transports less freight per unit of available capacity, indicating rising inefficiency under lower water levels. The demand flow cost, although it decreases slightly from 134,506 at 350 cm to 132,518 at 200 cm (a 1.48% reduction), masks the real inefficiencies in the system. The marginal reduction in flow cost likely results from fewer demands being met, as indicated by the drop in shipper satisfaction, rather than an improvement in operational efficiency. In contrast, service costs exhibit more noticeable fluctuations. From 11,900 at 350 cm, service costs increase by 6.48% to 12,670 at 250 cm, reflecting the increased use of smaller vessels that are less efficient and require more frequent services. By 200 cm, service costs drop slightly to 11,760 (a 7.18% drop compared to 250 cm). This reduction at 200 cm may indicate reduced activity in the network due to fewer overall demands being satisfied, but it does not imply improved operational efficiency. The drop in shipper satisfaction is stark, particularly in the full spot shippers category, which declines from 57.61% at 350 cm to 51.09% at 200 cm, representing an 11.32% decrease. Similarly, partial spot shipper satisfaction falls by 8.65%, from 62.39% to 57.00%. These declines reflect the system’s increasing inability to fully meet demand as water levels decrease. The reduced capacity, due to the shift towards smaller vessels, directly limits the number of shipments that can be satisfied, causing a significant drop in customer satisfaction.

Table 4: Performance Evaluation of Solutions at Different Water Levels

Water level	Profit	Costs		Shipper Satisfaction (%)		Capacity Usage (%)
		Demand Flow Cost	Service Cost	Partial Spot	Full Spot	
350 cm	568140	134506	11900	62.39	57.61	61.28
250 cm	565725	133927	12670	59.43	53.18	56.31
200 cm	559936	132518	11760	57.00	51.09	51.47

In summary, the drop in water levels from 350 cm to 200 cm leads to a cumulative reduction in profit by 1.46%, driven by increased operational inefficiencies. Although de-

mand flow costs appear stable, the 6.48% spike in service costs at 250 cm highlights the significant inefficiencies that arise as the network compensates by using smaller vessels more frequently. Additionally, the drop in shipper satisfaction—particularly the 11.32% decline in full spot shippers—illustrates the growing challenges in meeting demand. Overall, these results underscore the direct impact of lower water levels on profitability and customer satisfaction, driven by higher service costs and reduced network capacity.

7 Conclusions

In this study, we introduced the Water-Level-Constrained Scheduled Service Network Design with Resource and Revenue Management (WL-SSND-RRM) problem, which addresses the challenges of integrating water-level uncertainty into tactical planning for barge transportation systems. To explicitly account for the effects of water level uncertainty in tactical planning for barge transportation, we propose two alternative stochastic models, each defined by a specific approach to handling the impact of randomly changing water level conditions. The first stochastic model assumes that demand itineraries can be adjusted based on observed water levels, allowing for flows to dynamically adapt to fluctuating conditions. The second one, instead, relies on the application of penalties directly tied to the amount of demand that cannot be routed under the established tactical plan. To address the stochastic models, we apply a general decision-based scenario clustering analysis approach, which enables the computation of a series of alternative bounding techniques. Through numerical experiments, these techniques are shown to provide high-quality bounds and meaningful performance insights within practical computation times, particularly in settings where solving the full stochastic models is computationally challenging. This study also presented a set of comparative tests designed to evaluate the performance of the proposed models. Compared to the deterministic model, these tests highlight the importance of explicitly accounting for uncertainty in tactical planning. Moreover, the results compare the two stochastic models to assess the structural differences introduced in the tactical plan and their varying effectiveness in handling water level uncertainty. Finally, the study highlighted the significant influence of water levels on key tactical decisions (specifically, service selection and demand routing), as well as profitability and operational efficiency.

Several avenues for future research remain open. One direction is the development of novel exact solution methods to solve the proposed stochastic WL-SSND-RRM model. Additionally, future studies could extend the problem by incorporating other sources of uncertainty, such as variability in travel times, to enhance the robustness and applicability of the proposed methodologies. Another promising direction involves exploring alternative recourse strategies. For instance, explicitly managing the vessel fleet could enable dynamic reassignment of vessels to services, allowing the recourse strategy to effectively address infeasibilities caused by random fluctuations in water levels.

Acknowledgments

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 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.

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

- A.-d.-M. Agamez-Arias and J. Moyano-Fuentes. Intermodal transport in freight distribution: a literature review. *Transport Reviews*, 37(6):782–807, 2017.
- R. Bai, S. W. Wallace, J. Li, and A. Y.-L. Chong. Stochastic service network design with rerouting. *Transportation Research Part B: Methodological*, 60:50–65, 2014.
- I. Bakir, A. Erera, and M. Savelsbergh. Motor carrier service network design. In T. Crainic, M. Gendreau, and B. Gendron, editors, *Network Design with Applications in Transportation and Logistics*, chapter 16, pages 427–467. Springer, Boston, MA, 2021.
- 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, T. G. Crainic, and Y. Wang. Scheduled service network design with revenue management considerations and an intermodal barge transportation illustration. *European Journal of Operational Research*, 300(1):164–177, 2022.
- K. Braekers, A. Caris, and G. K. Janssens. Optimal shipping routes and vessel size for intermodal barge transport with empty container repositioning. *Computers in Industry*, 64(2):155–164, 2013.
- A. Caris, C. Macharis, and G. K. Janssens. Corridor network design in hinterland transportation systems. *Flexible Services and Manufacturing Journal*, 24:294–319, 2012.
- M. Christiansen, K. Fagerholt, B. Nygreen, and D. Ronen. Maritime transportation. In C. Barnhart and G. Laporte, editors, *Transportation*, volume 14 of *Handbooks in Operations Research and Management Science*, pages 189–284. North-Holland, Amsterdam, 2007.
- A. Christodoulou, P. Christidis, and B. Bisselink. Forecasting the impacts of climate change on inland waterways. *Transportation Research Part D: Transport and Environment*, 82:102159, 2020.
- J.-F. Cordeau, P. Toth, and D. Vigo. A survey of optimization models for train routing and scheduling. *Transportation Science*, 32(4):380–404, 1998.
- T. Crainic. Long-haul freight transportation. In R. Hall, editor, *Handbook of Transportation Science*, pages 451–516. Kluwer Academic Publishers, Norwell, MA, second edition, 2003.
- 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 K. Kim. Intermodal transportation. In C. Barnhart and G. Laporte, editors, *Transportation*, volume 14 of *Handbooks in Operations Research and Management Science*, chapter 8, pages 467–537. North-Holland, Amsterdam, 2007.
- 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, 2025. Forthcoming in *EURO Journal on Transportation and Logistics*.
- T. G. Crainic, X. Fu, M. Gendreau, W. Rei, and S. W. Wallace. Progressive hedging-based metaheuristics for stochastic network design. *Networks*, 58(2):114–124, 2011.
- T. G. Crainic, M. Hewitt, and W. Rei. Scenario grouping in a progressive hedging-based meta-heuristic for stochastic network design. *Computers & Operations Research*, 43: 90–99, 2014.
- T. G. Crainic, M. Hewitt, F. Maggioni, and W. Rei. Partial benders decomposition: general methodology and application to stochastic network design. *Transportation Science*, 55(2):414–435, 2021.
- Eurostat. Inland waterways freight transport statistics–2022. Rapport technique, European Commission, Luxembourg, 2023.
- M. Hewitt, T. G. Crainic, M. Nowak, and W. Rei. Scheduled service network design with resource acquisition and management under uncertainty. *Transportation Research Part B: Methodological*, 128:324–343, 2019.
- M. Hewitt, J. Ortmann, and W. Rei. Decision-based scenario clustering for decision-making under uncertainty. *Annals of Operations Research*, 315(2):747–771, 2022.
- A. Hoff, A.-G. Lium, A. Løkketangen, and T. G. Crainic. A metaheuristic for stochastic service network design. *Journal of Heuristics*, 16:653–679, 2010.
- X. Jiang, R. Bai, S. W. Wallace, G. Kendall, and D. Landa-Silva. Soft clustering-based scenario bundling for a progressive hedging heuristic in stochastic service network design. *Computers & Operations Research*, 128:105182, 2021.
- M. Kaut. A copula-based heuristic for scenario generation. *Computational Management Science*, 11:503–516, 2014.
- M. Kaut, H. Vladimirov, S. W. Wallace, and S. A. Zenios. Stability analysis of portfolio management with conditional value-at-risk. *Quantitative Finance*, 7(4):397–409, 2007.
- G. Lanza, T. G. Crainic, W. Rei, and N. Ricciardi. Scheduled service network design with quality targets and stochastic travel times. *European Journal of Operational Research*, 288(1):30–46, 2021.

- G. Lanza, T. Crainic, M. Passacantando, and M. Scutellà. Continuous-time service network design with stochastic travel times. Research Report CIRRELT-2024-29, Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique et le transport, Université de Montréal, Université de Montréal, Montréal, QC, Canada, 2024.
- A.-G. Lium, T. G. Crainic, and S. W. Wallace. A study of demand stochasticity in service network design. *Transportation Science*, 43(2):144–157, 2009.
- D. Müllerklein and P. Fontaine. Resilient transportation network design with disruption uncertainty and lead times. *European Journal of Operational Research*, 322(3):827–840, 2025.
- M. Prandtstetter, P. Widhalm, H. Leitner, and I. Czege. A novel visualisation tool for reliable inland navigation. *Transportation Research Procedia*, 72:3537–3544, 2023.
- B. V. Riessen, R. R. Negenborn, R. Dekker, and G. Lodewijks. Service network design for an intermodal container network with flexible transit times and the possibility of using subcontracted transport. *International Journal of Shipping and Transport Logistics*, 7(4):457–478, 2015.
- P. J. Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987.
- 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.
- Y. Sun, G. Zhang, Z. Hong, and K. Dong. How uncertain information on service capacity influences the intermodal routing decision: A fuzzy programming perspective. *Information*, 9(1):24, 2018.
- C. van Dorsser, F. Vinke, R. Hekkenberg, and M. van Koningsveld. The effect of low water on loading capacity of inland ships. *European Journal of Transport and Infrastructure Research*, 20(3):47–70, 2020.
- X. Wang, T. G. Crainic, and S. W. Wallace. Stochastic network design for planning scheduled transportation services: The value of deterministic solutions. *INFORMS Journal on Computing*, 31(1):153–170, 2019.
- Z. Wang and M. Qi. Robust service network design under demand uncertainty. *Transportation Science*, 54(3):676–689, 2020.
- P. Ypsilantis and R. Zuidwijk. Collaborative fleet deployment and routing for sustainable transport. *Sustainability*, 11(20):5666, 2019.
- R. Zhang, C. Huang, and X. Feng. Empty container repositioning with foldable containers in a river transport network considering the limitations of bridge heights. *Transportation Research Part A: Policy and Practice*, 133:197–213, 2020.

Y. Zheng and A. M. Kim. Rethinking business-as-usual: Mackenzie river freight transport in the context of climate change impacts in northern canada. *Transportation Research Part D: Transport and Environment*, 53:276–289, 2017.

Appendix. Stability test

While increasing the number of scenarios improves the representation of the water level distribution, it also complicates the process of obtaining an optimal solution. Our scenario generation method is inherently random, meaning that if the procedure is rerun with the same inputs, the resulting scenario tree varies. To ensure robustness, we use the approach proposed by (Kaut et al., 2007). We solve our two-stage stochastic models with a specific number of scenarios across 10 different randomly generated scenario trees with Gurobi. We then calculate the average and standard deviation of the objective function values. By repeating this process for different scenario counts, we study the effect of the number of scenarios on the stability of the results.

Table 5: Stability Results for 2-SPSDM model

Scenario Size	Overall Average	Overall SD
10	178943.86	4788.49
20	174989.42	5576.04
30	172455.44	3950.13
40	170642.14	3138.39

Table 6: Stability Results for 2-SPDFA model

Scenario Size	Overall Average	Overall SD
10	113088.30	2315.56
20	110212.41	1651.58
30	109274.73	1311.68
40	108167.22	1308.50

Table 5 presents the optimal results obtained from the 2-SPSDM model, indicating that as the number of scenarios increases, the Standard Deviation (SD) decreases, which translates to an increase in in-sample stability. As the number of scenarios rises, the Coefficient of Variation (CV)—the ratio of the standard deviation to the average—drops to 2.29%. However, when the number of scenarios increases from 30 to 40, the CV further decreases to 1.83%. Considering the computational cost of using 40 scenarios compared to 30, alongside the slight reduction in the CV value, we determine that 30 is the appropriate choice for the following experiments.

In Table 6, the best-known solutions for the 2-SPDFA model are presented. We set a runtime limit of 6 hours and obtain these results. The results indicate a similar trend to the first stochastic model. Based on these findings, we identify 20 scenarios as the best option for this model, achieving a CV of 1.49%.