UNIVERSITÉ DU QUÉBEC À MONTRÉAL

PLANNING INTERMODAL INLAND WATERWAY TRANSPORTATION UNDER ENVIRONMENTAL UNCERTAINTY AND NAVIGATIONAL CONSTRAINTS

THESIS

PRESENTED

AS PARTIAL REQUIREMENT

TO THE PH.D IN BUSINESS ADMINISTRATION

BY

BITA PAYAMI SHABESTARI

JANUARY 2026

UNIVERSITÉ DU QUÉBEC À MONTRÉAL

PLANIFICATION DU TRANSPORT INTERMODAL PAR VOIES NAVIGABLES INTÉRIEURES SOUS INCERTITUDES ENVIRONNEMENTALES ET CONTRAINTES DE NAVIGATION

THÈSE

PRÉSENTÉE

COMME EXIGENCE PARTIELLE

DU DOCTORAT EN ADMINISTRATION DES AFFAIRES

PAR

BITA PAYAMI SHABESTARI

JANVIER 2026

REMERCIEMENTS

This thesis was carried out under the supervision of Professor Teodor Gabriel Crainic, Professor Walter Rei, and Professor Ioana Bilegan. I would like to sincerely thank them for their support and guidance. Their advice, expertise, and encouragement were essential to the completion of this research. I am especially thankful to Professor Crainic for his valuable guidance, constructive advice, and rigorous feedback, which have greatly shaped the direction of this work. I also thank Professor Rei for his constructive feedback and steady support, which helped to improve the quality of this thesis. I am very grateful to Professor Bilegan, whose patience, continuous guidance, and trust made the research process smoother and more rewarding. It has been an honor and a privilege to work under the supervision of such distinguished professors, whose contributions have had a lasting impact on my academic and personal development.

I would also like to acknowledge the CIRRELT research community for providing an intellectually stimulating and collaborative environment. Exchanges of ideas, seminars, and discussions with my peers have been invaluable opportunities for learning and for broadening my perspectives on the challenges of transportation research.

The financial support of the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Fonds de recherche du Québec (FRQ) is gratefully acknowledged. Their contributions, through research and infrastructure grants, have been essential to the successful completion of this work.

Finally, my deepest appreciation goes to my family. Their unconditional love, patience, and encouragement have been a constant source of strength throughout this journey. This achievement is not only mine but also a reflection of their support.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

RÉSUMÉ

Le transport fluvial par barge constitue une alternative durable et économique au transport routier, offrant des solutions pour réduire la congestion et les émissions. Toutefois, la fréquence et la gravité croissantes des épisodes de sécheresse, exacerbées par les changements climatiques, ainsi que les interventions humaines telles que la gestion des barrages, le dragage et les prélèvements d'eau, provoquent des fluctuations du niveau des eaux dans les rivières et les canaux, compromettant la fiabilité et l'efficacité du transport par barge. Ces variations, incluant les faibles profondeurs, la réduction des niveaux d'eau et les limitations de tirant d'eau navigable, impactent directement les opérations en limitant la capacité de transport. La capacité de chargement d'une barge dépend non seulement de ses caractéristiques physiques, mais aussi du niveau d'eau disponible tout au long de son itinéraire. Des niveaux d'eau insuffisants ou en baisse imposent des restrictions de tirant d'eau qui limitent le volume et le poids des marchandises transportées, créant ainsi des défis opérationnels qui entravent la capacité des transporteurs à répondre efficacement à la demande des expéditeurs tout en maintenant leur rentabilité. Le premier article présente un cadre de planification tactique pour le transport fluvial par consolidation, qui modélise explicitement la relation entre le niveau des eaux et la capacité de chargement des barges. Ce cadre intègre les caractéristiques des navires ainsi que les niveaux d'eau prévus afin d'optimiser l'utilisation des ressources et de maximiser les revenus attendus, tout en assurant la satisfaction de la demande. Des expériences numériques permettent d'analyser l'effet des variations de niveaux d'eau sur la rentabilité, l'efficacité opérationnelle, la satisfaction des expéditeurs et la structure du réseau de services. Dans le deuxième article, nous introduisons le problème de conception de réseau de services programmés sous contrainte de niveau d'eau avec gestion des ressources et des revenus en contexte d'incertitude sur les niveaux d'eau (WL-SSND-RRM), qui traite les défis liés à l'intégration de la variabilité des niveaux d'eau dans la planification tactique des systèmes de transport par barge. Pour tenir compte explicitement de cette incertitude, deux modèles stochastiques alternatifs sont proposés, chacun reposant sur une approche spécifique de gestion des effets liés aux changements aléatoires des niveaux d'eau. Le premier modèle applique des pénalités liées à la demande qui ne pourrait satisfaite sous le plan tactique établi, reflétant les conséquences économiques d'une capacité insuffisante lors de l'exécution du plan. Le second permet une réaffectation dynamique des itinéraires selon les niveaux d'eau observés, ce qui améliore la flexibilité du réseau. Une approche générale de regroupement de scénarios basée sur les décisions est utilisée pour résoudre ces modèles, en proposant des techniques de calcul de bornes alternatives et en calculant l'écart d'optimalité. Comme souligné dans les deux premiers articles, il est essentiel que les modèles de planification intègrent explicitement l'impact de la variabilité des niveaux d'eau sur la capacité de service dans les décisions du côté de l'offre. Tout aussi important est le processus décisionnel du côté de la demande, où la prise en compte de ces facteurs externes est cruciale pour maintenir la rentabilité globale du système. Cela implique le contrôle de l'acceptation des demandes de réservation, reçues de manière continue par le transporteur, et qui forment la base du problème traité dans le troisième article. Nous développons un modèle de type bin-packing qui formalise le problème de réservation selon diverses politiques, le moment de décision, la flexibilité, et l'utilisation de stratégies myopes ou anticipatives. Le modèle est intégré dans un cadre à horizon roulant qui, à chaque étape, considère les demandes actuelles, celles déjà acceptées et, le cas échéant, les demandes futures prévues ainsi que les prévisions mises à jour du niveau des eaux. La performance des stratégies de réservation est évaluée par des expériences numériques approfondies.

ABSTRACT

Barge transportation provides a sustainable and cost-effective alternative to road transport, offering solutions to reduce congestion and lower emissions. However, increasing drought frequency and severity, driven by climate change, along with human interventions such as dam operations, dredging, and water withdrawals, cause fluctuations in water levels in rivers and canals that jeopardize its reliability and efficiency. These variations, including shallow water conditions, reduced water levels, and restricted navigable depths, directly impact barge operations by limiting transportation capacity. A barge's load capacity depends not only on its physical characteristics but also on the available water level along its route. Insufficient or diminished water levels impose draught restrictions that constrain freight volume and weight, creating operational challenges that hinder carriers' ability to meet shipper demands efficiently and maintain profitability. The first paper presents a tactical planning framework for consolidation-based barge transportation that explicitly models the relationship between water levels and barge load capacity. The framework integrates vessel characteristics and predicted water levels to optimize resource utilization and maximize expected carrier revenue while ensuring reliable demand fulfillment. Through computational experiments conducted using a commercial software, we analyze how varying water levels in rivers and canals impact profitability, operational efficiency, shipper satisfaction, and the service network structure. In the second paper, we introduce 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. 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 relies on the application of penalties directly tied to the amount of demand that cannot be routed under the established tactical plan, reflecting the economic consequences of inadequate capacity due to adverse water levels. The second model, instead, allows demand itineraries to be adjusted based on observed water levels, enabling flows to dynamically adapt to fluctuating conditions and enhancing network flexibility. To solve the stochastic models, we apply a general decision-based scenario clustering analysis approach, which enables the computation of a series of alternative bounding techniques aimed at finding feasible solutions and computing the optimality gap. As emphasized in the first two papers, there is a need for planning frameworks that explicitly integrate the impact of water level variability on service capacity into supply-side decision-making. Equally important, however, is the demand-side decision-making process, where accounting for such external factors is essential for maintaining overall system profitability. This involves controlling the acceptance of booking requests, which are continuously received by the carrier and form the basis of the planning problem addressed in the third paper. In this paper, we develop a bin-packing framework that formulates the booking problem under various booking policies, which may differ based on decision timing, flexibility, and the use of myopic or look-ahead policies. To capture the evolution of accepted requests and the operational impact of fluctuating water levels on vessel capacity, the model is embedded in a rolling-horizon framework. At each decision point, the model accounts for current requests, previously accepted requests, and—if applicable—predicted future requests and updated water-level forecasts. The performance of different booking strategies is evaluated through extensive numerical experiments using a commercial solver.

INTRODUCTION

0.1 Background, Motivation &Research Problems

Freight transportation is essential for economic growth, facilitating the smooth flow of goods necessary for industrial production, trade, and consumer markets. As globalization expands and supply chains become more interconnected, the demand for efficient, reliable, and cost-effective transportation systems continues to rise. This growing demand is driven by factors such as urbanization, increased industrial production, and the rapid growth of e-commerce. The scale of freight transportation is vast and continues to expand. For example, the United States, which has one of the most complex freight transportation networks in the world, moved approximately 19.7 billion tons of goods valued at $18.8 trillion in 2022, averaging nearly 54 million tons worth about $52 billion daily TRIP (2023). Projections indicate that by 2045, total freight across all modes—air, vessel, rail, and truck—will reach 25 billion tons, with the value expected to grow to $37 trillion U.S. Bureau of Transportation Statistics (BTS) (2024). These statistics highlight the rapidly increasing demand for efficient and sustainable freight transportation networks. Despite its economic importance, freight transportation presents significant environmental challenges. The transportation sector is responsible for approximately 20% of global energy-related greenhouse gas (GHG) emissions, with freight transport alone contributing between 8% and 10% of global emissions. Road transport, in particular, is a major contributor, accounting for around 15% of total emissions International Energy Agency (IEA) (2023). Extensive reliance on road transport for short and medium-haul freight further exacerbates environmental concerns. According to the European Environment Agency (EEA), road freight accounts for 71% of the EU's total freight emissions, a proportion that has remained largely unchanged despite technological advancements aimed at reducing emissions European Environment Agency (EEA) (2024).

Addressing the environmental impact of freight transportation requires concerted efforts to shift freight movements from road-based systems to more sustainable alternatives. To this end, various international organizations and government bodies have introduced policies aimed at promoting sustainable transportation systems. For example, the European Commission's White Paper on Transport established ambitious goals to significantly reduce greenhouse gas emissions from the transportation sector. The White Paper outlines strategies aimed at shifting at least 30% of road freight over 300 kilometers to rail or waterborne transport by 2030 and more than 50% by 2050 European Commission (2011). These initiatives encourage the adoption of environmentally friendly modes such as rail and inland waterways while simultaneously enhancing vehicle fuel efficiency and promoting technological innovations in clean energy.

Multimodal transport has emerged as a potential solution for reducing freight-related emissions. It involves the integration of different transportation modes such as road, rail, air, and water to provide a seamless, efficient, and environmentally responsible movement of goods. Unlike unimodal transport, which relies on a single mode of transportation throughout the entire logistics process, multimodal transport combines various modes under a single contract but with different operators. A prominent example is the integration of maritime and rail transport to enhance efficiency and reduce costs over long distances. By optimizing the use of each mode's strengths and minimizing their weaknesses, multimodal transport systems can enhance efficiency, reduce costs, and lower environmental impacts. Research suggests that road-rail systems, for example, perform significantly better than all-road transport in terms of environmental impact Aminzadegan *et al*. (2022). According to Rodrigue et Notteboom (2024), the integration of rail and waterborne transport systems, particularly for long-haul movements, can reduce CO2 emissions by up to 75% compared to road-only transport. This potential for emission reduction has made multimodal transport a key focus of policymakers and industry stakeholders aiming to enhance the sustainability of freight transportation.

Intermodal freight transportation, a specialized type of multimodal transportation, aims to efficiently integrate multiple transport modes within a unified network while maintaining a single loading unit, such as a standardized container, throughout the journey. Studies have shown that intermodal transport can offer substantial economic benefits by reducing handling costs and increasing operational efficiency Karam *et al*. (2023). However, its effectiveness heavily depends on the seamless integration of different transport modes and the availability of adequate infrastructure. Intermodal transportation typically comprises three primary segments: pre-haul, long-haul, and end-haul. The pre-haul segment involves moving freight from the shipper to the origin terminal. The long-haul segment refers to the transfer of containers between terminals using modes such as rail, air, or water. Finally, the end-haul segment involves delivering shipments from the destination terminal to the final receiver. A key component of intermodal freight transportation is inland waterway transport, particularly barge transportation. Inland waterways offer a highly energy-efficient and environmentally friendly option for long-haul freight movement. Barges consume 83% less energy per ton-kilometer compared to road transport and generate significantly lower noise pollution. Despite these advantages, barge transportation remains underutilized in many regions. For instance, in the European Union, inland freight transport has been predominantly dominated by road transport, which accounted for 75.5% of total freight movement in 2002. Rail and barge transport accounted for 18.3% and 6.2%, respectively, and by 2012, these figures had barely changed, with road transport at 75.1%, rail at 18.2%, and barge transport at 6.7% Eurostat (2015). This stagnation occurred despite EU targets aimed at shifting a substan-

tial share of freight traffic from road to more sustainable modes such as rail and waterborne transport. The limited utilization of inland waterways highlights the need for further exploration of their potential as integral components of sustainable intermodal transportation systems.

Although multimodal and intermodal transport systems enhance flexibility and reduce environmental impacts by integrating various modes, their inherent complexity also increases vulnerability to disruptions. A comprehensive review of the literature identifies disruptions as a major challenge affecting both operational efficiency and network reliability in such systems Akyüz *et al*. (2023). The effects of disruptions can be assessed using different performance metrics, including the proportion of satisfied demand post-disruption Chen et Miller-Hooks (2012), routing reliability to maximize served demand during disruptions Uddin et Huynh (2016, 2019), increased total travel time He *et al*. (2021), and decreased operational efficiency Ishfaq (2017).

Each mode within multimodal and intermodal systems—whether road, rail, air, or water—is increasingly exposed to a wide range of disruption types that can significantly affect performance and reliability. Disruptions in freight transportation are generally categorized as internal or external. Internal disruptions typically arise from operational inefficiencies, equipment failures, infrastructure breakdowns, or scheduled maintenance. In contrast, external disruptions are caused by factors beyond the operator's control, such as natural disasters, geopolitical instability, economic crises, cyber-attacks, and adverse weather events. The vulnerability of each transport mode depends on its operational characteristics. Road transport is particularly prone to traffic congestion, accidents, roadwork, and severe weather, including snow, ice, and flooding. Rail systems face disruptions from infrastructure bottlenecks, track maintenance, signal failures, and extreme conditions like floods, heatwaves, and landslides. Air transport is susceptible to fog, storms, high winds, and runway congestion, all of which contribute to significant delays and flight cancellations. Moreover, air travel is highly sensitive to events like volcanic ash clouds and severe atmospheric disturbances. Inland waterway transport is uniquely vulnerable to fluctuations in water levels. Inland waterways are directly affected by river flow, rainfall, snowmelt, and temperature. Barge transport, in particular, depends on stable water levels for navigability. Insufficient water depth can limit vessel draught, requiring load reductions that reduce efficiency and increase costs, while excessive water levels may restrict clearance under bridges, increase port congestion, and disrupt scheduled services. Recent statistics underscore the operational and economic impact of these disruptions. In 2022, European inland waterway freight volumes declined by 5.5% compared to the previous year, while transport performance fell by 10.6% to 122 billion

ton-kilometers. The Rhine experienced a 6.8% reduction in cargo traffic, with containerized volumes declining by 12.2%. Particularly striking was the 80% collapse in grain transport capacity along the Middle Danube, caused by persistently low water levels. These constraints also had financial consequences, with dry cargo freight rates on the Rhine increasing by 42.5% due to reduced vessel capacity Eurostat (2023).

Despite the wide range of disruptions affecting various transportation modes, water-related disruptions exhibit a distinct characteristic: predictability. Unlike abrupt events such as equipment failures, accidents, or geopolitical conflicts, seasonal climatic phenomena—including snowfall, droughts, monsoon rains, and river freeze–thaw cycles—tend to occur with predictable regularity. Moreover, advancements in meteorological forecasting, statistical modeling, and artificial intelligence have significantly enhanced the ability to anticipate weather patterns with reasonable accuracy. This predictive capability is particularly relevant in the context of inland waterway transportation, where fluctuations in water levels can be forecasted using hydrological models, real-time monitoring systems, and meteorological data. This allows water-level variability to be explicitly accounted for in planning models, thereby supporting more reliable and informed decision-making.

This study explores the potential of intermodal freight transportation as a sustainable logistics solution, with an emphasis on the role of inland waterways and freight consolidation—the practice of combining shipments from multiple origins into a single transport unit—in improving both economic and environmental performance. As global freight volumes continue to rise, the need for transportation systems that are efficient, reliable, and environmentally responsible has become increasingly critical. Consolidation-based intermodal barge transportation presents a promising avenue for addressing these challenges. By optimizing cargo loads and reducing the number of trips required to move a given volume of goods, this approach can significantly enhance operational efficiency. According to recent estimates, such strategies could reduce CO_2 emissions by up to 72% by 2050 relative to 2015 levels International Transport Forum (ITF) (2021). However, the effectiveness of consolidation-based intermodal barge transportation systems depends heavily on overcoming operational constraints associated with fluctuating water levels. Addressing these constraints requires improved freight planning practices that integrate water-level forecasts into decision-making processes. The motivation of this thesis is to address the planning and booking challenges of consolidation-based barge carriers operating over an entire season, explicitly accounting for water-level variability while mitigating the limitations of consolidation. The first objective is tactical planning. Carriers need effective planning of transportation activities before the beginning of the season—detailing the service network,

resource utilization, and 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 planning relies on demand forecasts and considers different levels of water-level prediction accuracy over short-term horizons (e.g., weekly), which are repeated consistently across the season. The second objective is the design of a booking system that enables carriers to dynamically respond to shipment requests throughout the season, with the aim of accepting only profitable requests while explicitly incorporating water-level fluctuations into the decision-making process.

0.2 Research Methodology & Contributions

A freight transportation system can be considered from two main perspectives: the supply side and the demand side. The demand side originates from shippers who require the movement of goods across various locations, while the supply side is represented by infrastructure providers (e.g., intermodal terminals, locks, and dams) and carriers who offer transportation services operated by vehicles—in our case, vessels.

Supply-side planning is typically structured into three hierarchical levels: strategic, tactical, and operational. Each level addresses distinct planning horizons, ranging from long-term decisions at the strategic level to mid-term decisions at the tactical level and short-term adjustments at the operational level. While each planning level has specific responsibilities, decisions made at higher levels inevitably influence those at lower levels, creating an interconnected planning framework. Strategic planning encompasses high-level policy decisions that establish the fundamental structure of transportation systems. This level involves critical decisions related to network design and fleet size determination. Tactical planning covers the mid-term horizon, typically spanning several months to a year. It focuses on designing detailed transportation plans that determine which services to operate, their scheduling or frequency, and the resources required to efficiently fulfill shipper demand. Unlike strategic planning, which addresses broad structural issues, tactical planning is more concerned with optimizing service schedules, route selection, and resource utilization to enhance efficiency and profitability. Tactical decisions are informed by aggregated demand forecasts. This planning level primarily addresses regular demand, which often accounts for approximately 75–80% of peak daily volume and tends to follow repetitive patterns, such as weekly cycles. Consequently, tactical plans are commonly structured around a standardized schedule, applied repeatedly over a defined planning season. Operational planning, in turn, represents the most granular level and addresses short-term decisions and real-time adjustments required to maintain the feasibility of the predefined tactical plan in

the face of disruptions or demand fluctuations.

Similarly, demand-side planning is organized across strategic, tactical, and operational levels. Strategic demand management aligns with the broader business strategy, defining the scope of service quality options, considering factors such as customer preferences, and identifying target markets. Decisions made at this level include selecting service regions, designing levels of service quality offerings, and defining customer segments intended for subsequent tactical planning. Effective strategic demand management provides the necessary boundaries within which tactical and operational decision-making processes operate. Tactical demand planning translates this strategy into differentiated service offerings and pricing schemes for various customer categories. This is often implemented through Revenue Management (RM) policies, which aim to maximize revenue by aligning service options with anticipated customer willingness-to-pay and demand profiles. In fact, at this level, the fare values (i.e., pricing) for each customer–service combination are explicitly defined. Finally, operational demand planning involves booking acceptance decisions based on available capacity and expected profitability.

Effective freight transportation planning requires strong integration between supply-side and demand-side planning, especially at the tactical level. While strategic planning defines the fundamental structure of transportation systems and aligns with the broader business strategy, tactical planning is used to define the transportation plan to be operated for the upcoming season. Bilegan *et al.* (2022) were the first to propose a tactical planning model that systematically integrates Revenue Management (RM) considerations within a Scheduled Service Network Design with Resource management (SSND-RRM) framework for consolidation-based intermodal barge transportation. In this framework, multiple customer categories are considered based on their business relationships with carriers. Regular shippers have long-term contracts, ensuring predictable demand and access to guaranteed capacity, while spot shippers operate on an ad hoc basis, requesting transportation services without long-term commitments. Carriers offer a range of fare classes, which are priced categories designed to reflect different levels of service quality. These fare classes allow shippers to choose between options that vary in delivery time, flexibility, and cost. For instance, standard delivery involves longer transit times at lower costs, while express delivery offers faster transit times at premium rates. Such fare class structures enable carriers to tailor services to meet the diverse needs of shippers. The aim of the proposed framework is to optimize service schedules, resource utilization, and shipment routing in order to maximize carriers' revenues while accommodating diverse shipper categories and varying service quality requirements.

The current study contributes to the literature by incorporating the effects of water level variability into an SSND-RRM framework—an environmental and infrastructure constraint that directly influences vessel operability and load capacity in inland waterways. High water levels typically allow for full utilization of vessel capacity, enhancing operational efficiency and profitability, but may introduce clearance issues under bridges, restricting larger vessels. Conversely, low water levels reduce the number of containers that can be transported per trip due to the increased risk of grounding, thereby raising transportation costs. This was evident on the Rhine River in 2018 and 2022, when low water levels led to an 11.1% decline in freight volumes due to navigational constraints Federal Statistical Office of Germany (Destatis) (2019).

To address these challenges, we propose a tactical planning framework that explicitly integrates the relationship between water levels and vessel characteristics. Given predicted water levels across the network, the model aims to maximize expected carrier revenue over a tactical horizon (e.g., a season) by establishing a tactical plan detailing the service network, resource utilization, and shipment handling and transport activities for a given schedule length (e.g., a week), executed repeatedly throughout the season. This integration supports adaptive, efficient, and profitable inland waterway operations under variable waterway conditions. We also introduce 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. 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(2-SPSDM model) relies on the application of penalties directly tied to the amount of demand that cannot be routed under the established tactical plan, reflecting the economic consequences of inadequate capacity due to adverse water levels. The second model(2-SPDFA model), instead, allows demand itineraries to be adjusted based on observed water levels, enabling flows to dynamically adapt to fluctuating conditions and enhancing network flexibility. As emphasized, there is a need for planning frameworks that explicitly integrate the impact of water level variability on service capacity into supply-side decision-making. 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.

Building upon this tactical foundation, we then investigate the booking problem for consolidation-based intermodal freight carriers, which governs shipment acceptance and rejection decisions throughout the season. While the tactical plan ensures the availability of services and resources, the booking process dynamically accepts or rejects in response to sequentially arriving non-contractual requests, while honoring all contractual commitments. To the best of our knowledge, this is the first work to explicitly address booking-level control in this setting through a bin-packing framework, thereby linking tactical planning with booking management. The proposed approach proceeds in two steps: (i) identifying feasible service ways for each shipment item based on time–space compatibility, and (ii) applying bin-packing models to assess profitability and determine acceptance decisions. Three models are formulated to capture different booking policies: a myopic bin-packing booking problem (MBBP), a lookahead bin-packing booking problem (LBBP), and an extended lookahead version with water-level constraints (LBBP-WL). Together, these formulations provide a unified framework for analyzing how different informational contexts and environmental conditions shape booking decisions and, ultimately, system performance.

In summary, this study focuses on tactical planning and booking planning within inland waterway freight transportation. To ensure a well-defined research scope, a set of assumptions and limitations are established as follows:

- This study is confined to water-based freight transportation, specifically inland and coastal waterways. While intermodal integration with land-based modes (such as trucks and trains) plays a crucial role in real-world logistics, this research does not explicitly model other transportation modes. Instead, land-based transportation is considered implicitly through time-related constraints at terminals and service connections, such as handling time at terminals.
- The physical network in this study comprises waterways and ports. It is assumed that the waterway network remains stable in terms of width and length, allowing simultaneous vessel travel in both directions without changes in distance. However, water level variations that affect vessel draught and navigability are explicitly considered, alongside capacity constraints at terminals, which limit the number of vessels that can be moored at any given time.
- To simplify the presentation, and because the modeling approach for navigation impact remains similar regardless of the direction of water level changes, this study focuses on situations involving decreased water levels. This focus is justified by the observation that, while occasional water level in-

creases due to excessive rainfall occur, the predominant challenge in logistics and transportation is the long-term trend of declining water levels.

- Water level fluctuations, which directly influence vessel capacity and navigability, are explicitly modeled in both tactical and operational planning. The study assumes that water level forecasts are available with different degrees of accuracy.

- The study assumes that a single carrier operates scheduled services using a predefined fleet of vessels. Each vessel type is characterized by technical specifications such as length, speed, draught (both maximum and minimum), and cargo capacity in terms of tonnage and twenty-foot equivalent units (TEUs). A predefined set of potential service schedules specifies fixed arrival and departure times at terminals along a designated route consisting of an origin, a destination, and intermediate terminals. The cost structure incorporates fixed costs associated with service selection, variable transportation costs, and handling costs.

- This study does not focus on demand forecasting. Instead, it assumes that demand-related information—such as origin, destination, arrival time, due time, and cargo dimensions—is available for decision-making. Forecasted transport requests, demand patterns, and shipper requirements are treated as given inputs rather than variables requiring prediction.

0.3 Structure of the thesis

This thesis is composed of three chapters, organized as follows: Chapter I introduces the integration of water-level constraints into tactical planning for a consolidation-based intermodal barge transportation system. Chapter II presents the Water-Level-Constrained Scheduled Service Network Design with Resource and Revenue Management (WL-SSND-RRM) framework, which addresses the challenges of incorporating water-level variability into tactical service network planning. Chapter III introduces comprehensive booking-level decision problem for consolidation- based intermodal freight carriers operating over fixed service networks, where sequential accept/reject decisions must be made under incomplete information, accounting for both contractual shipment commitments and dynamic non-contract-shipment requests. We propose a bin packing–based modelling framework to accept profitable requests.

CHAPTER 1
TACTICAL NETWORK PLANNING FOR INTERMODAL BARGE TRANSPORTATION CONSIDERING VARYING WATER LEVELS

Barge transportation provides a sustainable and cost-effective alternative to road transport, offering solutions to reduce congestion and lower emissions. However, increasing drought frequency and severity, driven by climate change, along with human interventions such as dam operations, dredging, and water withdrawals, cause fluctuations in water levels in rivers and canals that jeopardize its reliability and efficiency. These variations, including shallow water conditions, reduced water levels, and restricted navigable depths, directly impact barge operations by limiting transportation capacity. A barge's load capacity depends not only on its physical characteristics but also on the available water level along its route. Insufficient or diminished water levels impose draught restrictions that constrain freight volume and weight, creating operational challenges that hinder carriers' ability to meet shipper demands efficiently and maintain profitability. This chapter presents a tactical planning framework for consolidation-based barge transportation that explicitly models the relationship between water levels and barge load capacity. The framework integrates vessel characteristics and predicted water levels to optimize resource utilization and maximize expected carrier revenue while ensuring reliable demand fulfillment. Through computational experiments conducted using a commercial software, we analyze how varying water levels in rivers and canals impact profitability, operational efficiency, shipper satisfaction, and the service network structure.

1.1 Introduction

Inland waterways, including rivers and canals, are crucial for freight transport. In the United States, they handled approximately 495 billion tonne-kilometers of cargo annually from 2004 to 2022 CEIC Data (2023). Similarly, European Union (EU) waterways transported 469 million tons of goods in 2023 Eurostat (2023). Globally, barges play a critical role in regions with extensive waterway networks, offering a highly efficient and environmentally sustainable mode of transportation. A single barge can replace 40 to 60 trucks, reducing fuel consumption by 75% and greenhouse gas emissions by 80% compared to road transport. Barges also require significantly less energy per ton-kilometer, consuming only 17% of road transport energy and 50% of rail transport energy. Due to these advantages, barges are increasingly integrated into intermodal freight transportation systems, where different modes—such as barges, trucks, trains, and airplanes-are integrated to form an efficient and flexible transportation network. This integration is particularly advanta-

geous in regions with well-developed inland waterway infrastructure.

Within these systems, consolidation-based inland waterway carriers operate scheduled services across terminals connected by navigable waterways. Each service is operated by a vessel and follows a fixed route—including origin, destination, and potentially intermediate stops—with specified departure and arrival times. It is characterized by the vessel's physical attributes (e.g., draught, length, load capacity) and operational features (e.g., speed, cost). Carriers offer these services in response to transportation demands from shippers, who request that shipments be picked up from the origin terminal no earlier than their availability time and delivered to the destination terminal no later than the due time, often with specific service quality requirements (e.g., standard or express). A key challenge in this context is tactical planning, which involves jointly selecting a cost-efficient set of services and schedules while optimizing shipment itineraries that define routing and operations (e.g., transfers, consolidation). Bilegan *et al.* (2022) provide an in-depth discussion of tactical planning in intermodal freight transportation for inland waterway transport (IWT) carriers. They propose a new Scheduled Service Network Design model with integrated Resource and Revenue Management (SSND-RRM) to optimize service schedules, resource utilization, and shipment routing, aiming to maximize carriers' revenues while accommodating diverse shipper categories and service quality requirements.

The current study contributes to the literature by incorporating the effects of water level variability into an SSND-RRM formulation—an environmental and infrastructure constraint that directly influences vessel operability and load capacity in inland waterways. High water levels typically allow for full utilization of vessel capacity, enhancing operational efficiency and profitability, but may introduce clearance issues under bridges, restricting larger vessels. Conversely, low water levels reduce the number of containers that can be transported per trip due to the increased risk of grounding, thereby raising transportation costs. This was evident on the Rhine River in 2018 and 2022, when low water levels led to an 11.1% decline in freight volumes due to navigational constraints Federal Statistical Office of Germany (Destatis) (2019). To address these challenges, we propose a tactical planning framework that explicitly integrates the relationship between water levels and vessel characteristics. Given predicted water levels across the network, the model aims to maximize expected carrier revenue over a tactical horizon (e.g., a season) by establishing a tactical plan detailing the service network, resource utilization, and shipment handling and transport activities for a given schedule length (e.g., a week), executed repeatedly throughout the season. This integration supports adaptive, efficient, and profitable IWT operations under variable waterway conditions.

The remainder of the paper is structured as follows: Section 2 discusses water levels and vessel characteristics; Section 3 presents the mathematical model formulation; Section 4 reports experimental results and analyses; and Section 5 concludes with a summary of findings and directions for future research.

1.2 Water Levels and vessel characteristics

On one hand, the variability of waterway depth—typically measured at critical points along the waterway as the vertical distance between the water surface and the riverbed or canal bottom—represents a key operational constraint in IWT. Water levels rarely remain constant throughout an entire waterway and often vary from segment to segment due to both natural factors (e.g., geographic location, rainfall, droughts, snowmelt, sedimentation) and human interventions (e.g., dam operations, dredging, or water withdrawals for agricultural and urban use). For instance, upstream dam releases may temporarily raise water levels in certain sections while reducing downstream flow, resulting in shallower segments. Similarly, heavy rainfall can improve navigability, whereas prolonged drought or sediment buildup may decrease it. These variations are typically uneven, meaning that different sections of a river or canal may exhibit significantly different water levels at the same time.

On the other hand, another key operational constraint lies in the physical characteristics of vessels navigating inland waterways. One critical parameter is the vessel's draught—defined as the vertical distance between the waterline and the lowest point of the hull—which indicates how deeply the vessel is submerged and is directly influenced by the weight of the containers being transported. Each vessel has two critical draught thresholds: the maximum draught, representing the depth reached when the vessel is fully loaded, and the minimum draught, corresponding to the depth when it is empty Prandtstetter *et al*. (2023). These values define the vessel's safe loading range, or loading capacity of each vessel, typically measured in weight (tons), volume (TEUs), or both. A vessel's actual draught, which varies with the number and weight of containers, falls within this defined range. As more containers are loaded, the actual draught increases accordingly, approaching the maximum value. The rate at which this increase occurs is determined by the vessel's load–draught coefficient, which quantifies the sensitivity of the draught to added weight.

Therefore, for safe and efficient operation of scheduled services in inland waterway transportation, it is essential to consider the relationship between water levels and vessel characteristics. At every segment of the service route, the vessel's draught must remain less than or equal to the available water level. This constraint becomes particularly critical during periods of low water levels, as it limits load capacity and

increases the risk of grounding. Conversely, when water levels rise, although submerged clearance improves and allows for heavier loads, another constraint emerges. As the water surface elevates, the entire vessel is lifted, decreasing the vertical distance between the vessel's highest point and overhead structures such as bridges. This may restrict the passage of larger vessels with tall superstructures.

These constraints highlight the importance of adaptive tactical planning that accounts for spatial water level variability and vessel characteristics, as they jointly impact transport capacity, service feasibility, and operational reliability. Integrating these factors into tactical planning frameworks is essential for efficient and reliable IWT operations, ultimately contributing to smooth operations of intermodal transportation networks. To simplify the analysis, the current study focuses primarily on submerged clearance rather than bridge clearance. This choice is supported by several key considerations: first, submerged clearance is the predominant navigational constraint, as most inland waterways are not critically restricted by bridge heights; second, declining water levels—driven by climate change and increasing water withdrawals—occur more frequently and have greater operational impact; and third, the modeling approach is conceptually similar for both submerged and bridge clearance. Therefore, focusing on submerged clearance enables a more relevant and streamlined analysis.

1.3 Notations and mathematical formulation

1.3.1 Notations

Let us consider a physical network $\mathcal{G}^{\mathsf{ph}} = (\mathcal{N}^{\mathsf{ph}}, \mathcal{A}^{\mathsf{ph}})$, where $\mathcal{N}^{\mathsf{ph}}$ denotes the set of geographical ports, each having a specific berthing capacity Q_η measured in length units, and $\mathcal{A}^{\mathsf{ph}}$ represents the set of physical links. Each physical link $a \in \mathcal{A}^{\mathsf{ph}}$ is characterized by a specific water level p_a at each time period, which is typically represented by a point estimate derived statistically from historical data. This point estimate, often calculated as the mode or average, reflects the most likely water level along each physical link over the planning horizon.

We define the set of shipping demands $\mathcal{K} = \mathcal{K}^{\mathsf{R}} \cup \mathcal{K}^{\mathsf{P}} \cup \mathcal{K}^{\mathsf{F}}$, where $\mathcal{K}^{\mathsf{R}}$ (regular shippers) includes demands that must be fully served, $\mathcal{K}^{\mathsf{P}}$ (partial-spot shippers) includes demands that may be partially met, and $\mathcal{K}^{\mathsf{F}}$ (full-spot shippers) includes demands that can either be fully met or not met at all. Each demand $k \in \mathcal{K}$ is defined by its volume $d(k)$ (in TEUs), origin $O(k)$, destination $D(k)$, release time $\alpha(k)$, and due time $\beta(k)$. The volume is converted to weight using a factor ω. Demands can be standard or express ($class(k)$), with

express deliveries being charged higher fares ($\phi(k)$). Container type $\gamma(k) \in \Gamma$ influences handling, storage, and transportation costs.

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 $\text{cap}_w(v)$ in tonnage and $\text{cap}_{vol}(v)$ 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 and by its schedule. 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 $\mathcal{N}^{\text{ph}}(\sigma) = \{\eta_i(\sigma) \mid i = 0, \ldots, n(\sigma)\}$. 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 $i = 1, \ldots, n(\sigma)$. Network operating costs include holding costs ($h(\eta, v)$) for idle vessels at terminals, fixed service setup costs ($f(\sigma)$), and penalties ($\mu(v)$) for unused vessels. Demand fulfillment costs cover transportation, holding, and handling. Transportation costs depend on container type and vessel type ($c_i(\gamma(k), v(\sigma))$), terminal storage costs are represented by $c(\eta, \gamma(k))$, and handling costs by $\kappa(\eta_i(\sigma), \gamma(k))$.

1.3.2 Mathematical formulation

The 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, $t \in \{0, 1, \ldots, T-1\}$. The node set $\mathcal{N}$ is defined as $\{(\eta, t) \mid \eta \in \mathcal{N}^{\text{ph}}, t = 0, \ldots, T-1\}$, 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: $\mathcal{A}^M = \{((\eta, t), (\eta', t')) \mid \eta, \eta' \in \mathcal{N}^{\text{ph}}, t, t' \in \{0, \ldots, T-1\}, t < t'\}$. 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: $\mathcal{A}^H = \{((\eta, t), (\eta, t+1)) \mid \eta \in \mathcal{N}^{\text{ph}}, t \in \{0, \ldots, T-1\}\}$.

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 $l_i(\sigma) = \{(\eta_{i-1}(\sigma), \eta_i(\sigma)) \mid i = 1, \ldots, n(\sigma), \sigma \in \Sigma\}$, 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 $a_{\eta t} \in \mathcal{A}^H$. Having established the time-space network, we proceed to formulate the SSND-RRM model by first introducing the decision variables:

- $y(\sigma) \in \{0, 1\}$: 1 if transportation service σ is selected, 0 otherwise.
- $\xi(k) \in [0, 1]$: percentage of the volume of partial-spot shipper demand $k \in \mathcal{K}^{\mathsf{P}}$ that is selected and will be serviced.
- $\zeta(k) \in \{0, 1\}$: 1 if full-spot shipper demand $k \in \mathcal{K}^{\mathsf{F}}$ is serviced, 0 otherwise.
- $z(v, a_{\eta_t}) \in \mathbb{Z} \geq 0$: number of temporarily idle vessels of type v waiting at holding arc $a\eta_t$ 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, it $B(v)$ remains constant across all time periods, although vessels may either be moving or idle in ports at any given period.
- $x(k, a_{l_i}(\sigma)) \geq 0$: volume of demand $k \in \mathcal{K}$ transported by service σ on its i^{th} leg.
- $x^{\text{out}}(k, a_{l_i}(\sigma)) \geq 0$: volume of demand $k \in \mathcal{K}$ to be unloaded from leg i of service σ.
- $x^{\text{in}}(k, a_{l_i}(\sigma)) \geq 0$: volume of demand $k \in \mathcal{K}$ to be loaded onto leg i of service σ.
- $x^{\text{hold}}(k, a_{\eta_t}) \geq 0$: volume of demand $k \in \mathcal{K}$ to be held at terminal η during the time period $(t, t+1)$.

$$\max \sum_{k\in\mathcal{K}^{\mathsf{R}}} \phi(k)d(k) + \sum_{k\in\mathcal{K}^{\mathsf{P}}} \phi(k)\xi(k)d(k) + \sum_{k\in\mathcal{K}^{\mathsf{F}}} \phi(k)\zeta(k)d(k) - \sum_{v\in\mathcal{V}} \mu(v)(F_v - B(v)) - \sum_{\sigma\in\Sigma} f(\sigma)y(\sigma)$$
$$- \sum_{a_{\eta_t}\in\mathcal{A}^H} \sum_{v\in\mathcal{V}} h(\eta, v)z(v, a_{\eta_t}) - \sum_{a_{l_i}(\sigma)\in\mathcal{A}^M} \sum_{k\in\mathcal{K}} c_i(\gamma(k), v(\sigma))x(k, a_{l_i}(\sigma)) - \sum_{a_{\eta_t}\in\mathcal{A}^H} \sum_{k\in\mathcal{K}} c(\eta, \gamma(k))x^{\text{hold}}(k, a_{\eta_t})$$
$$- \sum_{a_{l_i}(\sigma)\in\mathcal{A}^M} \sum_{k\in\mathcal{K}} \kappa(\eta_i(\sigma), \gamma(k)) \left(x^{\text{in}}(k, a_{l_i}(\sigma)) + x^{\text{out}}(k, a_{l_i}(\sigma))\right) \tag{1.1}$$

Subject to

$$x^{\text{hold}}(k, a_{\eta_t}) + \sum_{a_{l_i}(\sigma)\in\mathcal{A}^M, (\eta_{i-1}=\eta, \alpha(\eta_{i-1})=t)} x^{\text{in}}(k, a_{l_i}(\sigma)) = \begin{cases} d(k), & \forall k \in \mathcal{K}^{\mathsf{R}}, \eta = O(k), t = \alpha(k) \\ \xi(k)d(k), & \forall k \in \mathcal{K}^{\mathsf{P}}, \eta = O(k), t = \alpha(k) \\ \zeta(k)d(k), & \forall k \in \mathcal{K}^{\mathsf{F}}, \eta = O(k), t = \alpha(k) \end{cases} \tag{1.2}$$

$$\sum_{\alpha(k)<t\leq\beta(k)} \sum_{a_{l_i}(\sigma)\in\mathcal{A}^M, (\eta_i=\eta, \beta(\eta_i)=t)} x^{\text{out}}(k, a_{l_i}(\sigma)) = \begin{cases} d(k), & \forall k \in \mathcal{K}^{\mathsf{R}}, \eta = D(k) \\ \xi(k)d(k), & \forall k \in \mathcal{K}^{\mathsf{P}}, \eta = D(k) \\ \zeta(k)d(k), & \forall k \in \mathcal{K}^{\mathsf{F}}, \eta = D(k) \end{cases} \tag{1.3}$$

$$\mathrm{x}^{\text{hold}}(k, a_{\eta_{t-1}}) + \sum_{a_{l_i}(\sigma)\in\mathcal{A}^M (\eta_i=\eta, \beta(\eta_i)=t)} \mathrm{x}^{\text{out}}(k, a_{l_i}(\sigma))$$
$$-\mathrm{x}^{\text{hold}}(k, a_{\eta_t}) - \sum_{a_{l_i}(\sigma)\in\mathcal{A}^M ((\eta_{i-1}=\eta, \alpha(\eta_{i-1})=t)} \mathrm{x}^{\text{in}}(k, a_{l_i}(\sigma)) = 0 \tag{1.4}$$
$$\forall (\eta, t) \neq (O(k), \alpha(k)) \forall \eta \neq D(k), \forall k \in \mathcal{K}$$

$$x^{\text{in}}(k, a_{l_i}(\sigma)) - x(k, a_{l_i}(\sigma)) = 0, \quad \forall a_{l_i}(\sigma) \in \mathcal{A}^M,\ \eta_{i-1}(\sigma) = O(\sigma),\ k \in \mathcal{K} \tag{1.5}$$

$$x(k, a_{l_i}(\sigma)) - x^{\text{out}}(k, a_{l_i}(\sigma)) = 0, \quad \forall a_{l_i}(\sigma) \in \mathcal{A}^M,\ \eta_i(\sigma) = D(\sigma),\ k \in \mathcal{K} \tag{1.6}$$

$$\begin{aligned} x(k, a_{l_{i-1}}(\sigma)) - \mathrm{x}^{\text{out}}(k, a_{l_{i-1}}(\sigma)) + \mathrm{x}^{\text{in}}(k, a_{l_i}(\sigma)) - x(k, a_{l_i}(\sigma)) = 0, \\ \forall \sigma \in \Sigma,\ \eta_{i-1} \neq O(\sigma), \eta_i \neq D(\sigma),\ k \in \mathcal{K} \end{aligned} \tag{1.7}$$

$$\omega \sum_{k \in \mathcal{K}} x(k, a_{l_i}(\sigma)) \leq cap_w(v(\sigma))y(\sigma), \quad \forall \sigma \in \Sigma, a_{l_i}(\sigma) \in \mathcal{A}^M \tag{1.8}$$

$$\sum_{k \in \mathcal{K}} x(k, a_{l_i}(\sigma)) \leq cap_{vol}(v(\sigma))y(\sigma), \quad \forall \sigma \in \Sigma, a_{l_i}(\sigma) \in \mathcal{A}^M \tag{1.9}$$

$$\theta(v(\sigma))\omega \sum_{k \in \mathcal{K}} x(k, a_{l_i}(\sigma)) + dh^-(v(\sigma)) \leq p_{a_{l_i}(\sigma)}, \quad \forall \sigma \in \Sigma, a_{l_i}(\sigma) \in \mathcal{A}^M \tag{1.10}$$

$$B(v) = \sum_{\eta \in \mathcal{N}^{\text{ph}}} z(v, a_{\eta_0}) + \sum_{\sigma \in \Lambda_{0l}} y(\sigma), \quad \forall v \in \mathcal{V} \tag{1.11}$$

$$B(v) \leq F_v, \quad \forall v \in \mathcal{V} \tag{1.12}$$

$$\sum_{\sigma \in \Sigma^-_{\eta t v}} y(\sigma) + z(v, a_{\eta_{t-1}}) = \sum_{\sigma \in \Sigma^+_{\eta t v}} y(\sigma) + z(v, a_{\eta_t}), \quad \forall v \in \mathcal{V},\ a_{\eta_{t-1}}, a_{\eta_t} \in \mathcal{A}^H \tag{1.13}$$

$$\sum_{v \in \mathcal{V}} len(v)(\sum_{\sigma \in \Sigma^-_{\eta t v}} y_\sigma + z(v, a_{\eta_{t-1}})) \leq Q_{a_{\eta t}}, \quad \forall a_{\eta t} \in \mathcal{A}^H \tag{1.14}$$

$$y(\sigma) \in \{0,1\} \ \forall \sigma \in \Sigma \tag{1.15}$$

$$\xi(k) \in [0,1] \ \forall k \in \mathcal{K}^{\mathsf{P}} \tag{1.16}$$

$$\zeta(k) \in \{0,1\} \forall k \in \mathcal{K}^{\mathsf{F}} \tag{1.17}$$

$$z(v, a_{\eta_t}) \geq 0 \ \forall v \in \mathcal{V}, \ a_{\eta_t} \in \mathcal{A}^H \tag{1.18}$$

$$B(v) \geq 0, integer \ \forall v \in \mathcal{V} \tag{1.19}$$

$$x(k, a_{l_i}(\sigma)) \geq 0 \ \forall k \in \mathcal{K}, \ a_{l_i}(\sigma) \in \mathcal{A}^M \tag{1.20}$$

$$x^{\text{out}}(k, a_{l_i}(\sigma)) \geq 0 \quad \forall k \in \mathcal{K}, \ a_{l_i}(\sigma) \in \mathcal{A}^M \tag{1.21}$$

$$x^{\text{in}}(k, a_{l_i}(\sigma)) \geq 0 \quad \forall k \in \mathcal{K}, \ a_{l_i}(\sigma) \in \mathcal{A}^M \tag{1.22}$$

$$x^{\text{hold}}(k, a_{\eta_t}) \geq 0 \quad \forall k \in \mathcal{K}, \ a_{\eta_t} \in \mathcal{A}^H \tag{1.23}$$

The objective function 1.1 maximizes net profit by considering revenue from servicing regular, partial-spot, and full-spot shippers, along with costs related to unused vessels, service setup, vessel idling, container transportation, and handling at terminals. Equations 1.2, 1.3, and 1.4 are flow-conservation constraints for containers of all shipper types, at their particular origins, destinations, and intermediary nodes, respectively. Similarly, Equations 1.5, 1.6 and 1.7 enforce the conservation of container flows, for all shipper types, on each service at its origin, destination and intermediary stops, respectively. 1.8 and 1.9 guarantee that the weight and volume of all demands $k \in \mathcal{K}$ transported by service σ on its leg i do not exceed the nominal capacity of vessel v performing service σ. Constraint 1.10 assures that the draught of vessels executing the i^{th} leg of service σ is less than or equal to the water level on that leg. Equation 1.11 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 $\Lambda_{0l} = \{\sigma \in \Sigma, v(\sigma) = v | (\alpha_{n(\sigma)} \mod T) < \beta_0(\sigma) \quad \text{and} \quad \beta_0(\sigma) \geq 0\} \subseteq \Sigma$ containing all services, of the appropriate vessel type, that operate one of its legs during the first period. Constraints 1.12 enforce the fleet size for each vessel type, while Equations 1.13 are the so-called design-balance constraints, enforcing the vehicle-flow conservation at terminals (the number of services and idle vessels entering a node equals the number exiting the node), where sets $\Sigma^-_{\eta t v} = \{\sigma \in \Sigma \mid \eta_i(\sigma) = \eta, \beta(\eta_i(\sigma)) = t, v(\sigma) = v, \text{ for } i = 1, \ldots, n(\sigma)\}$

and $\Sigma^{+}_{\eta t v} = \{\sigma \in \Sigma \mid \eta_i(\sigma) = \eta, \alpha(\eta_i(\sigma)) = t, v(\sigma) = v, \text{ for } i = 0\}$ group the services with a vessel type v that arrive at their destination or depart from their origin, respectively. Finally, Constraints 1.14 enforce the terminal berthing capacity at each time period. Decision-variable domains are defined by Constraints 1.15 - 1.23.

1.4 Experimental results

We generate test instances for two network topologies: a linear network with four connected terminals and a star network with five terminals linked to a central hub. Water levels are considered at critical points along the waterways, ranging between 150 cm and 350 cm, reflecting natural variations as described by Christodoulou *et al*. (2020). The linear and star networks include 20 and 71 routes, respectively. Two types of vessels are used: small and large, with nominal capacities of 20 and 50 TEUs, respectively, and with large vessels costing twice as much as small vessels. Services are scheduled with nine distinct start times over a 14-period weekly cycle, resulting in 360 potential services for the linear network (20 routes $\times$ 9 start times $\times$ 2 vessel types) and 1278 for the star network (71 routes $\times$ 9 start times $\times$ 2 vessel types). Four demand instances are generated for each network, with demand volumes randomly ranging between 5 and 25 TEUs. The naming convention for test instances comprises two parts: the first part specifies the network topology, where "L" denotes the linear network and "S" denotes the star network. The second part indicates the number of demand requests, ensuring a diverse range of transportation needs. In the following, we present the analysis and results obtained using the described test instances. All implementations were carried out 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 256GB of memory. The results from the computational experiments, summarized in Tables 1.1 to 1.4, compare solutions with and without water level constraints for both linear and star networks. Tables 1.1 and 1.2 show results for the linear network, while Tables 1.3 and 1.4 present results for the star network. Performance indicators such as profitability, service costs, demand flow costs, capacity utilization (measured as the ratio of total volume-kilometers moved to total capacity-kilometers operated), and shipper satisfaction (the proportion of requested demand volume accepted for transportation), along with computational time (in seconds) reported in the last column, are evaluated.

In terms of profitability, excluding water level constraints consistently yields higher results. In the linear network, the impact is especially pronounced. For instance, profit increases from 184,827.84 to 222,641.00 in TestL-124, and from 191,580.54 to 230,475.00 in TestL-136. On average, the linear network sees a 20.4%

increase in profitability when water levels are not considered. In the star network, the gain is more modest, with an average increase of approximately 0.5%. The most significant improvement occurs in TestS-170, where profit increases from 272,659.51 to 273,399.00 units.

Table 1.1: Performance evaluation of solutions **with water levels** for the linear network.

Test Name	Profit	Demand Cost	Service Cost	Satisfaction (%)	Cap. Usage (%)	Time (s)
TestL-73	109,523.79	31,426.21	2,240.00	57.32	59.27	21.82
TestL-89	136,297.30	36,442.81	3,080.00	53.52	45.56	25.29
TestL-124	184,827.84	45,461.85	3,220.00	37.61	48.23	41.08
TestL-136	191,580.54	51,029.85	3,080.00	41.57	56.63	33.28

Table 1.2: Performance evaluation of solutions **without water levels** for the linear network.

Test Name	Profit	Demand Cost	Service Cost	Satisfaction (%)	Cap. Usage (%)	Time (s)
TestL-73	110,318.00	30,402.00	2,170.00	61.02	58.29	30.92
TestL-89	150,407.00	40,543.00	2,940.00	63.41	53.62	42.35
TestL-124	222,641.00	56,074.00	3,570.00	68.81	61.42	55.65
TestL-136	230,475.00	63,030.00	3,780.00	64.07	57.67	55.02

When water level constraints are considered, service costs tend to increase due to a higher number of services and greater reliance on small vessels. Given their lower nominal capacity compared to large vessels, small vessels offer fewer consolidation opportunities, resulting in limited cost savings. This effect is observed in TestL-73 and TestL-89, where service costs are higher under water level constraints (2,240.00 vs. 2,170.00; 3,080.00 vs. 2,940.00), along with an increase in the number of services and the share of small vessels. Specifically, TestL-73 involves 30 services (2 large, 28 small) compared to 23 services (8 large, 15 small) when water level constraints are not considered. Similarly, in TestL-89, the fleet composition shifts from 8 large and 28 small vessels to 12 large and 18 small. When water level constraints are not considered, the total number of services generally decreases and the share of large vessels increases, enhancing consolidation opportunities and allowing for greater potential cost savings. However, in some cases, such as TestL-124 and TestL-136, the increased use of large vessels results in higher fixed costs, which are not fully compensated by the consolidation gains. As a result, total service costs increase—from 3,220.00 to 3,570.00 in TestL-124, and from 3,080.00 to 3,780.00 in TestL-136. In TestL-124, the number of large vessels increases from 8 (with constraints) to 25 (without), while small vessels decrease from 30 to 1. Similarly, in TestL-136, the fleet shifts

from 8 large and 28 small vessels to 23 large and 8 small vessels.

Table 1.3: Performance evaluation of solutions **with water levels** for the star network.

Test Name	Profit	Demand Cost	Service Cost	Satisfaction (%)	Cap. Usage (%)	Time (s)
TestS-127	183,972.38	77,457.62	2,170.00	24.95	40.17	193.55
TestS-142	206,095.10	83,814.90	2,590.00	24.82	35.70	235.78
TestS-170	272,659.51	105,167.53	2,170.00	18.76	41.39	323.47
TestS-228	329,126.75	134,933.25	2,940.00	20.95	43.34	858.60

Table 1.4: Performance evaluation of solutions **without water levels** for the star network.

Test Name	Profit	Demand Cost	Service Cost	Satisfaction (%)	Cap. Usage (%)	Time (s)
TestS-127	184,194.00	77,446.00	1,960.00	23.70	39.91	202.83
TestS-142	207,722.00	84,368.00	2,310.00	25.99	43.55	249.82
TestS-170	273,399.00	105,016.00	2,380.00	22.62	47.05	267.64
TestS-228	329,330.00	134,690.00	3,080.00	24.41	46.96	798.49

A similar pattern is observed in the star network. When water levels are considered, more services are activated to compensate for the reduced vessel capacity. For example, in TestS-127, the system operates 5 large and 21 small vessels (total: 26), compared to 4 large and 20 small (total: 24) without water level constraints. In TestS-142, the fleet decreases from 29 vessels (8 large and 21 small) to 27 (6 large and 21 small) when constraints are removed, resulting in a reduction in service cost—from 2,590.00 (with water levels) to 2,310.00 (without). However, in TestS-170 and TestS-228, service costs increase when water level constraints are removed. In TestS-170, the cost rises from 2,170.00 (with water levels) to 2,380.00 (without), and in TestS-228, from 2,940.00 to 3,080.00. This increase occurs despite the total number of services remaining constant (28 in both cases), as the system shifts toward a greater reliance on large vessels in the absence of navigational limitations. Therefore, water level constraints impact service costs by influencing both the number of services and the types of vessels deployed, based on demand and navigability.

Demand flow costs generally decrease when water level constraints are applied, but this reduction reflects unmet demand rather than improved efficiency. In the linear network, demand flow costs drop significantly (e.g., TestL-89: 40,543.00 to 36,442.81), but lower capacity usage and decreased shipper satisfaction indicate poor utilization of vessel capacity. For instance, shipper satisfaction drops from 63.41% to 53.52% in TestL-89. Only TestL-73 shows stable performance with a slight cost increase but relatively consistent satis-

faction and capacity usage. In the star network, demand flow costs and shipper satisfaction remain stable with or without water level constraints, suggesting resilience due to the centralized structure where all nodes connect to a central hub. The average demand flow cost is similar (100,343.83 with constraints vs. 100,379.5 without), and satisfaction only slightly varies (22.37% vs. 24.18%), indicating consistent performance despite navigational challenges.

1.5 Conclusions and perspectives

This study introduces varying water levels as a critical environmental and infrastructural constraint in the tactical planning of consolidation-based inland waterway carriers. By enhancing the existing methodology, we propose a framework that explicitly models the relationship between vessel load capacity and water levels, ensuring service feasibility only when a vessel's draught remains less than or equal to the available water level.

Experimental tests illustrate that the proposed model offers a more accurate and resilient approach to service network design, resource utilization, and demand satisfaction. Comparative analysis across linear and star network topologies shows that neglecting water level constraints leads to overestimated profitability and shipper satisfaction. Incorporating water levels encourages the use of smaller, adaptable vessels, increased service selection, and capacity deployment adjustments aligned with actual navigability conditions. Additionally, the analysis reveals that linear networks are more vulnerable to water level fluctuations due to limited routing options, where each node is directly connected only to its neighbors. In contrast, star networks benefit from a centralized structure that provides multiple routing alternatives through a central hub, enhancing flexibility and robustness against disruptions.

Future research could improve the model by incorporating stochastic water levels, addressing uncertainties in environmental and infrastructure conditions, and considering how water levels affect service travel times, including delays and associated penalty costs for late pickups and deliveries. This enhancement would provide more realistic insights and support better decision-making in inland waterway freight transportation. Additionally, a real case study would help validate the model's practical relevance and provide insights for managerial decision-making.

CHAPTER 2
SERVICE NETWORK DESIGN WITH UNCERTAINTY ON WATER LEVELS FOR INTERMODAL RIVER TRANSPORT

Barge transport presents a sustainable alternative to road transport, offering the potential to alleviate congestion and 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 also 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.

2.1 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 et Hewitt, 2021; Crainic et Rei, 2024). 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 (Prandtstetter *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 et 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 2.2 presents the barge transportation system, including the physical network, supply side, and demand side. Section 2.3 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 2.4 outlines the problem setting and introduces the proposed stochastic WL-SSND-RRM formulations for the barge transportation system. Section 2.5 explains the solution approach adopted in this research. Section 2.6 reports the experimental results. Finally, Section 2.7 summarizes the main findings and outlines future research directions.

2.2 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 2.2.1. The key stakeholders in this system are carriers and shippers, representing the supply and demand sides, respectively, which will be further explored in Sections 2.2.2 and 2.2.3.

2.2.1 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 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.

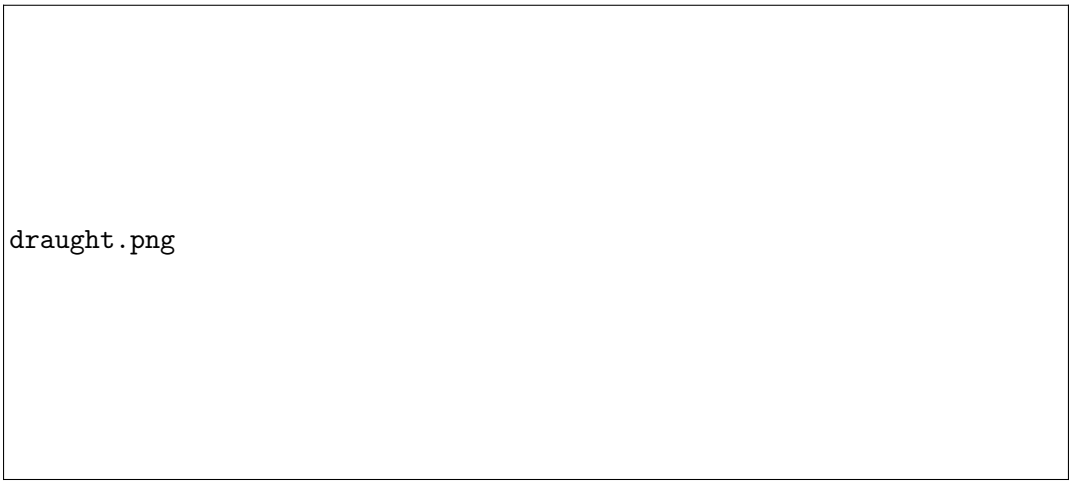

Figure 2.1: Schematic view of the most relevant measurements related to vessel's draught and water level

2.2.2 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 et 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 Figure2.1). 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 2.1. The actual draught can be described as an incremental function, where each added unit of load results in a proportional increase in draught (**?**). 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.

2.2.3 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 et 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.

2.3 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 et Hewitt, 2021; Crainic et Rei, 2024). 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 et Kim, 2007; Agamez-Arias et 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 et 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 et 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 strategies—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 et Hewitt, 2021) and (Crainic et Rei, 2024), 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 et 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.*, 2021b). 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.

2.4 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 2.4.1. This is followed by a detailed explanation of the specific aspects and mathematical formulations for each variant in Section 2.4.2.

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 $\mathcal{N}^{\text{ph}}$ denotes the set of geographical ports and $\mathcal{A}^{\text{ph}}$ 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 $\mathcal{K}^{\text{R}}$ (regular shippers) includes demands that must be fully served, $\mathcal{K}^{\text{P}}$ (partial-spot shippers) includes demands that may be partially met, and $\mathcal{K}^{\text{F}}$ (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 $\text{cap}_w(v)$ in tonnage and $\text{cap}_{vol}(v)$ in TEUs. Additional characteristics include length $\text{len}(v)$, 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 $\mathcal{N}^{\text{ph}}(\sigma) = \{\eta_i(\sigma) \mid i = 0, \ldots, n(\sigma)\}$. 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 $i = 1, \ldots, n(\sigma)$.

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, $t \in \{0, 1, \ldots, T-1\}$. The node set $\mathcal{N}$ is defined as $\{(\eta, t) \mid \eta \in \mathcal{N}^{\text{ph}}, t = 0, \ldots, T-1\}$, 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: $\mathcal{A}^M = \{((\eta, t), (\eta', t')) \mid \eta, \eta' \in \mathcal{N}^{\text{ph}}, t, t' \in \{0, \ldots, T-1\}, t < t'\}$ 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: $\mathcal{A}^H = \{((\eta, t), (\eta, t+1)) \mid \eta \in \mathcal{N}^{\text{ph}}, t \in \{0, \ldots, T-1\}\}$. 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 $l_i(\sigma) = \{(\eta_{i-1}(\sigma), \eta_i(\sigma)) \mid i = 1, \ldots, n(\sigma), \sigma \in \Sigma\}$, 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 $k \in \mathcal{K}$.

2.4.2 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 realizations. 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.

2.4.2.1 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.

  - $x^{\textbf{out}}(k, a_{l_i}(\sigma))^s \geq 0$: Volume of demand $k \in \mathcal{K}$ unloaded from the i^{th} leg of service σ under scenario s.

  - $x^{\textbf{in}}(k, a_{l_i}(\sigma))^s \geq 0$: Volume of demand $k \in \mathcal{K}$ loaded onto the i^{th} leg of service σ under scenario s.

  - $x^{\textbf{hold}}(k, a_{\eta_t})^s \geq 0$: 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:

$$\begin{aligned} \max \quad & \sum_{k \in \mathcal{K}^{\mathsf{R}}} \phi(k) d(k) + \sum_{k \in \mathcal{K}^{\mathsf{P}}} \phi(k) \xi(k) d(k) + \sum_{k \in \mathcal{K}^{\mathsf{F}}} \phi(k) \zeta(k) d(k) \\ & - \sum_{v \in \mathcal{V}} \mu(v)(F_v - B(v)) - \sum_{\sigma \in \Sigma} f(\sigma) y(\sigma) - \sum_{a_{\eta_t} \in \mathcal{A}^H} \sum_{v \in \mathcal{V}} h(\eta, v) z(v, a_{\eta_t}) - \sum_{s \in S} \rho^s G^s(\boldsymbol{\chi}) \end{aligned} \tag{2.1}$$

Subject to

$$B(v) = \sum_{\eta \in \mathcal{N}^{\mathsf{ph}}} z(v, a_{\eta_0}) + \sum_{\sigma \in \Lambda_{0l}} y(\sigma), \ \ \forall v \in \mathcal{V} \tag{2.2}$$

$$B(v) \leq F_v, \ \ \forall v \in \mathcal{V} \tag{2.3}$$

$$\sum_{\sigma \in \Sigma^-_{\eta t v}} y(\sigma) + z(v, a_{\eta_{t-1}}) = \sum_{\sigma \in \Sigma^+_{\eta t v}} y(\sigma) + z(v, a_{\eta_t}), \ \ \forall v \in \mathcal{V}, \ a_{\eta_{t-1}}, a_{\eta_t} \in \mathcal{A}^H \tag{2.4}$$

$$\sum_{v\in\mathcal{V}} len(v)(\sum_{\sigma\in\Sigma^-_{\eta tv}} y_\sigma + z(v,a_{\eta_{t-1}})) \le Q_{a_{\eta t}},\ \ \forall a_{\eta t}\in\mathcal{A}^H \tag{2.5}$$

$$y(\sigma)\in\{0,1\}\ \ \forall\sigma\in\Sigma \tag{2.6}$$

$$\xi(k)\in[0,1]\ \ \forall k\in\mathcal{K}^{\mathsf{P}} \tag{2.7}$$

$$\zeta(k)\in\{0,1\}\forall k\in\mathcal{K}^{\mathsf{F}} \tag{2.8}$$

$$z(v,a_{\eta_t})\ge 0\ \ \forall v\in\mathcal{V},\ \ a_{\eta_t}\in\mathcal{A}^H \tag{2.9}$$

$$B(v)\ge 0, \text{integer}\ \ \forall v\in\mathcal{V} \tag{2.10}$$

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.

$$\begin{aligned} G^s(\boldsymbol{\chi}) = \min & \sum_{a_{l_i}(\sigma)\in\mathcal{A}^M}\sum_{k\in\mathcal{K}} c_i(\gamma(k),v(\sigma))\,x(k,a_{l_i}(\sigma))^s + \sum_{a_{\eta_t}\in\mathcal{A}^H}\sum_{k\in\mathcal{K}} c(\eta,\gamma(k))\,\mathrm{x}^{\mathrm{hold}}(k,a_{\eta_t})^s \\ & + \sum_{a_{l_i}(\sigma)\in\mathcal{A}^M}\sum_{k\in\mathcal{K}} \kappa(\eta_i(\sigma),\gamma(k))\left(\mathrm{x}^{\mathrm{in}}(k,a_{l_i}(\sigma))^s + \mathrm{x}^{\mathrm{out}}(k,a_{l_i}(\sigma))^s\right) \end{aligned} \tag{2.11}$$

Subject to

$$\mathrm{x}^{\mathrm{hold}}(k,a_{\eta_t})^s + \sum_{a_{l_i}(\sigma)\in\mathcal{A}^M,(\eta_{i-1}=\eta,\alpha(\eta_{i-1})=t)} \mathrm{x}^{\mathrm{in}}(k,a_{l_i}(\sigma))^s = \begin{cases} d(k), & \forall k\in\mathcal{K}^{\mathsf{R}},\eta=O(k),t=\alpha(k) \\ \xi(k)d(k), & \forall k\in\mathcal{K}^{\mathsf{P}},\eta=O(k),t=\alpha(k) \\ \zeta(k)d(k), & \forall k\in\mathcal{K}^{\mathsf{F}},\eta=O(k),t=\alpha(k) \end{cases} \tag{2.12}$$

$$\sum_{\alpha(k)<t\le\beta(k)}\ \sum_{a_{l_i}(\sigma)\in\mathcal{A}^M,(\eta_i=\eta,\beta(\eta_i)=t)} \mathrm{x}^{\mathrm{out}}(k,a_{l_i}(\sigma))^s = \begin{cases} d(k), & \forall k\in\mathcal{K}^{\mathsf{R}},\eta=D(k), \\ \xi(k)d(k), & \forall k\in\mathcal{K}^{\mathsf{P}},\eta=D(k) \\ \zeta(k)d(k), & \forall k\in\mathcal{K}^{\mathsf{F}},\eta=D(k) \end{cases} \tag{2.13}$$

$$\mathrm{x}^{\mathrm{hold}}(k, a_{\eta_{t-1}})^s + \sum_{a_{l_i}(\sigma)\in\mathcal{A}^M(\eta_i=\eta,\beta(\eta_i)=t)} \mathrm{x}^{\mathrm{out}}(k, a_{l_i}(\sigma))^s - \mathrm{x}^{\mathrm{hold}}(k, a_{\eta_t})^s - \sum_{a_{l_i}(\sigma)\in\mathcal{A}^M((\eta_{i-1}=\eta,\alpha(\eta_{i-1})=t)} \mathrm{x}^{\mathrm{in}}(k, a_{l_i}(\sigma))^s = 0 \quad \forall(\eta,t) \neq (O(k),\alpha(k))\forall\eta \neq D(k), \forall k \in \mathcal{K} \tag{2.14}$$

$$\mathrm{x}^{\mathrm{in}}(k, a_{l_i}(\sigma))^s - x(k, a_{l_i}(\sigma))^s = 0, \forall a_{l_i}(\sigma) \in \mathcal{A}^M,\ \eta_{i-1}(\sigma) = O(\sigma),\ k \in \mathcal{K} \tag{2.15}$$

$$x(k, a_{l_i}(\sigma))^s - \mathrm{x}^{\mathrm{out}}(k, a_{l_i}(\sigma))^s = 0,\ \ \forall a_{l_i}(\sigma) \in \mathcal{A}^M,\ \eta_i(\sigma) = D(\sigma),\ k \in \mathcal{K} \tag{2.16}$$

$$x(k, a_{l_{i-1}}(\sigma))^s - \mathrm{x}^{\mathrm{out}}(k, a_{l_{i-1}}(\sigma))^s + \mathrm{x}^{\mathrm{in}}(k, a_{l_i}(\sigma))^s - x(k, a_{l_i}(\sigma))^s = 0, \quad \forall\sigma \in \Sigma,\ \eta_{i-1} \neq O(\sigma), \eta_i \neq D(\sigma),\ k \in \mathcal{K} \tag{2.17}$$

$$\omega \sum_{k\in\mathcal{K}} x(k, a_{l_i}(\sigma))^s \leq cap_w(v(\sigma))y(\sigma),\ \ \forall\sigma \in \Sigma, a_{l_i}(\sigma) \in \mathcal{A}^M \tag{2.18}$$

$$\sum_{k\in\mathcal{K}} x(k, a_{l_i}(\sigma))^s \leq cap_{vol}(v(\sigma))y(\sigma),\ \ \forall\sigma \in \Sigma, a_{l_i}(\sigma) \in \mathcal{A}^M \tag{2.19}$$

$$\theta(v(\sigma))\omega \sum_{k\in\mathcal{K}} x(k, a_{l_i}(\sigma))^s + dh^-(v(\sigma)) \leq p'^s_{a_{l_i}(\sigma)} \forall\sigma \in \Sigma, a_{l_i}(\sigma) \in \mathcal{A}^M \tag{2.20}$$

$$x(k, a_{l_i}(\sigma))^s \geq 0 \quad \forall k \in \mathcal{K},\ a_{l_i}(\sigma) \in \mathcal{A}^M \tag{2.21}$$

$$\mathrm{x}^{\mathrm{out}}(k, a_{l_i}(\sigma))^s \geq 0 \quad \forall k \in \mathcal{K},\ a_{l_i}(\sigma) \in \mathcal{A}^M \tag{2.22}$$

$$\mathrm{x}^{\mathrm{in}}(k, a_{l_i}(\sigma))^s \geq 0 \quad \forall k \in \mathcal{K},\ a_{l_i}(\sigma) \in \mathcal{A}^M \tag{2.23}$$

$$\mathrm{x}^{\mathrm{hold}}(k, a_{\eta_t})^s \geq 0 \quad \forall k \in \mathcal{K},\ a_{\eta_t} \in \mathcal{A}^H \tag{2.24}$$

The mathematical formulation is presented in Eqs. (2.1)-(2.24). The objective function (Eq. (2.1)) 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. (2.11). Constraints (2.2)-(2.5) ensure the feasibility of the first-stage decisions, while constraints (2.6)-(2.10) define the boundaries for these decisions.Equation (2.2) 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 $\Lambda_{0l} = \{\sigma \in \Sigma, v(\sigma) = v | (\alpha_{n(\sigma)} \mod T) < \beta_0(\sigma) \text{ and } \beta_0(\sigma) \geq 0\} \subseteq \Sigma$ containing all services, of the appropriate vessel type, that operate one of its legs during the first period. Constraints (2.3) enforce the fleet size for each vessel type, while Equations (2.4) 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 $\Sigma^+_{\eta tv}$

$$\Sigma^-_{\eta tv} = \{\sigma \in \sigma \mid D(\sigma) = \eta, \beta_n(\sigma) = t, v(\sigma) = v\} \tag{2.25}$$

$$\Sigma^+_{\eta tv} = \{\sigma \in \Sigma \mid O(\sigma) = \eta, \alpha_0(\sigma) = t, v(\sigma) = v\} \tag{2.26}$$

group the services with a vessel type v that arrive at their destination or depart from their origin, respectively. Finally, Constraints (2.5) enforce the terminal berthing capacity at each time period. Constraints (2.12)–(2.20) 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 (2.12), (2.13), and (2.14), respectively. Similarly, container flow is conserved along each service route at its origin, destination, and intermediate legs via (2.15), (2.16), and (2.17). Eqs. (2.18) and (2.19) 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 (2.20) ensures that the vessel's draught does not exceed the available water level under each scenario $s \in S$. Finally, constraints (2.21)–(2.24) define the non-negativity domains of the associated recourse variables.

2.4.2.2 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:

$$
\begin{aligned}
\max \quad & \sum_{k\in\mathcal{K}^{\mathsf{R}}} \phi(k)d(k) + \sum_{k\in\mathcal{K}^{\mathsf{P}}} \phi(k)\xi(k)d(k) + \sum_{k\in\mathcal{K}^{\mathsf{F}}} \phi(k)\zeta(k)d(k) \\
& - \sum_{v\in\mathcal{V}} \mu(v)(F_v - B(v)) - \sum_{\sigma\in\Sigma} f(\sigma)y(\sigma) - \sum_{a_{\eta_t}\in\mathcal{A}^H} \sum_{v\in\mathcal{V}} h(\eta, v)z(v, a_{\eta_t}) \\
& - \sum_{a_{l_i}(\sigma)\in\mathcal{A}^M} \sum_{k\in\mathcal{K}} c_i(\gamma(k), v(\sigma))x(k, a_{l_i}(\sigma)) - \sum_{a_{\eta_t}\in\mathcal{A}^H} \sum_{k\in\mathcal{K}} c(\eta, \gamma(k))\mathrm{x}^{\mathrm{hold}}(k, a_{\eta_t}) \\
& - \sum_{a_{l_i}(\sigma)\in\mathcal{A}^M} \sum_{k\in\mathcal{K}} \kappa(\eta_i(\sigma), \gamma(k)) \left(\mathrm{x}^{\mathrm{in}}(k, a_{l_i}(\sigma)) + \mathrm{x}^{\mathrm{out}}(k, a_{l_i}(\sigma))\right) - \sum_{s\in S} \rho^s G^s(\boldsymbol{\chi})
\end{aligned}
\tag{2.27}
$$

Subject to

constraints defined in Eqs. (2.2)–(2.10), (2.12)–(2.19), and (2.21)–(2.24).

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.

$$
G^s(\boldsymbol{\chi}) = \min\ b(k) \sum_{a_{l_i}(\sigma)\in\mathcal{A}^M} \sum_{k\in\mathcal{K}} \lambda(k, a_{l_i}(\sigma))^s \tag{2.28}
$$

Subject to:

$$
\theta(v(\sigma))\,\omega \sum_{k\in\mathcal{K}} (x(k, a_{l_i}(\sigma)) - \lambda(k, a_{l_i}(\sigma))^s) + dh^-(v(\sigma)) \le p'^s_{a_{l_i}(\sigma)}, \quad \forall\sigma\in\Sigma,\ a_{l_i}(\sigma)\in\mathcal{A}^M \tag{2.29}
$$

$$
\lambda(k, a_{l_i}(\sigma))^s \le x(k, a_{l_i}(\sigma)), \quad \forall k\in\mathcal{K},\ a_{l_i}(\sigma)\in\mathcal{A}^M \tag{2.30}
$$

$$
\lambda(k, a_{l_i}(\sigma))^s \ge 0, \quad \forall k\in\mathcal{K},\ a_{l_i}(\sigma)\in\mathcal{A}^M \tag{2.31}
$$

The 2-SPSDM model, formulated in Eqs. (2.27)–(2.31), 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 de-

mand across all scenarios, each corresponding to a potential water-level realization (Eq. (2.28)). The model constraints ensure the feasibility of the first-stage decisions. These include flow conservation constraints (Eqs. (2.12), (2.13), (2.14)), flow balance across services (Eqs. (2.15), (2.16), (2.17)), service capacity limitations (Eqs. (2.18), (2.19)), and constraints on fleet size, fleet balance, and terminal berthing capacity (Eqs. (2.2), (2.3), (2.4), (2.5)). 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 (2.29), (2.30), and (2.31). Specifically, constraint (2.29) ensures that the vessel's draught does not exceed the available water level for each scenario $s \in S$. Constraints (2.30) and (2.31) define the feasible region for second-stage decisions.

2.5 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.

2.5.1 Step 1: Scenario Clustering Analysis

Given scenario set $S = \{s_1, s_2, \ldots, s_n\}$, 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 2.4.1 for the 2-SPSDM model and

in Section 2.4.2 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 $g : \mathcal{F}_{\mathcal{X}} \times S \to \mathbb{R}^{+}$. 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 $\theta(s_i \mid s_j) = g(\chi^*_{s_j}; s_j) - g(\chi^*_{s_i}; s_j)$, where $\theta(s_i \mid s_j)$ 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 $\theta(s_i \mid s_j)$ reflects the relative performance of the decisions across different scenarios. However, the opportunity cost can be *asymmetric* ($\theta(s_i \mid s_j) \neq \theta(s_j \mid s_i)$), 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: $D(s_i, s_j) = \theta(s_i \mid s_j) + \theta(s_j \mid s_i)$. 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 $C_k^{\min},\ k = 1, \ldots, K^{\min}$, 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 $C_k^{\max},\ k = 1, \ldots, K^{\max}$, 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.

2.5.2 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 $C_k^{\min},\ k = 1, \ldots, K^{\min}$. 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 $\bar{\chi} = \arg\max_{\chi \in \mathcal{F}_\mathcal{X}} \sum_{M_k \in M} g(\chi, M_k) \times p_{M_k}$. 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 $C_k^{\max},\ k = 1, \ldots, K^{\max}$. 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: $\bar{\chi}_{C_k^{\max}} = \arg\max_{\chi \in \mathcal{F}_\mathcal{X}} \sum_{s_i \in C_k^{\max}} g(\chi, s_i) \times \frac{\rho^{s_i}}{P_{C_k^{\max}}}$. The upper bound is calcu-

lated 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 $\min_{k=1,\ldots,K^{\max}} \sum_{s_i \in C_k^{\max}} g(\bar{\chi}_{C_k^{\max}}; s_i) \times \rho^{s_i}$.

2.6 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 2.6.1 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 2.6.2. We also investigate several critical managerial questions in the remainder of the section. Section 2.6.3 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 2.6.4 assesses the effectiveness of these alternative recourse strategies. Finally, Section 2.6.5 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.

2.6.1 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 2.2a) 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 conditions, 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.

(a) Physical network (b) Correlation matrix

Figure 2.2: Shipment-request ways impact

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.

2.6.2 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 2.5 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 2.1 and Table 2.2 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.

The performance evaluation for the 2-SPSDM model (Table 2.1) shows that the Cluster-based approach is

Table 2.1: 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

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 2.2 reports the performance evaluation of the proposed solution approach for the 2-SPDFA model

Table 2.2: 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

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 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.

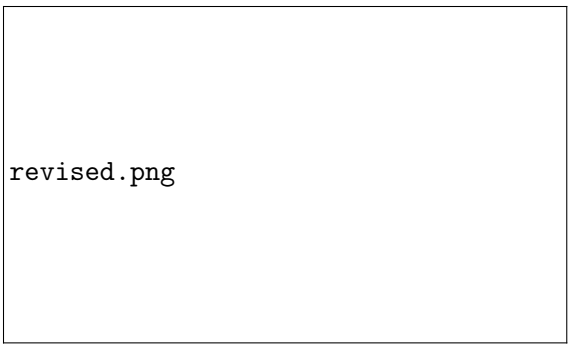

(a) Comparison of SS and EVS for 2-SPSDM model

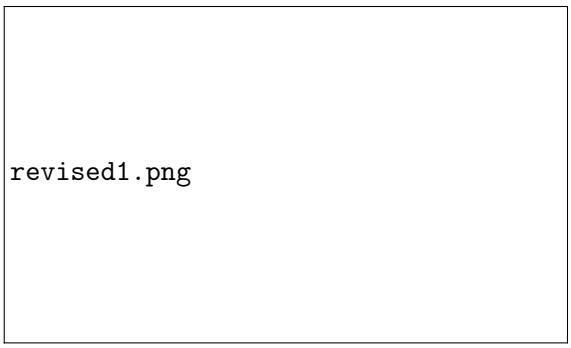

(b) Comparison of SS and EVS for 2-SPDFA model

Figure 2.3: Value of the Stochastic Solution

2.6.3 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 2.3a and Figure 2.3b 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 2.3a, 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 2.3b, 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.

2.6.4 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 2.6.1. 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 2.3: Comparison tests

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

In Table 2.3, 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 $\Delta 4$ 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.

$\Delta 3$ 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 $\Delta 4$, 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.

2.6.5 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.

2.6.5.1 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.

2.6.5.2 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.

Figure 2.4 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 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

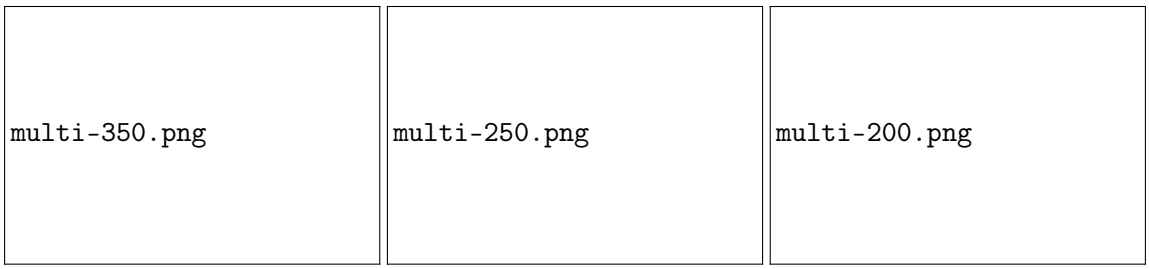

(a) Water level = 350 cm (b) Water level = 250 cm (c) Water level = 200 cm

Figure 2.4: Multi-path usage at different water levels: 350 cm, 250 cm, and 200 cm.

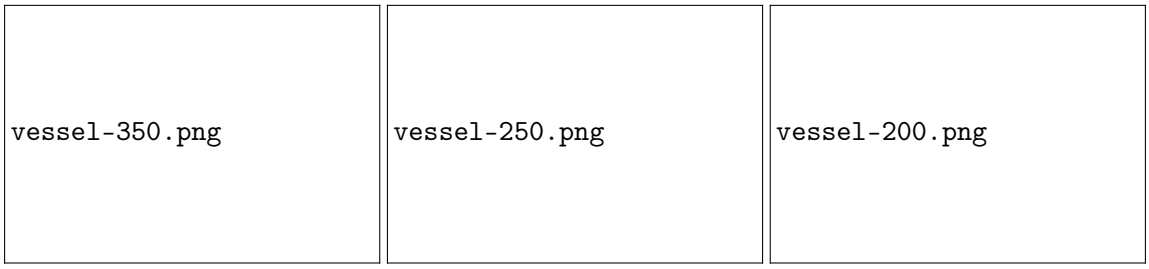

(a) Water level = 350 cm (b) Water level = 250 cm (c) Water level = 200 cm

Figure 2.5: Vessel usage at different water levels: 350 cm, 250 cm, and 200 cm.

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 2.5 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 2.4).

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 2.4.

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 2.4: 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 demand 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.

2.7 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 de-

signed 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.

2.8 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 2.5: 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 2.5 presents the optimal results obtained from the 2-SPSDM model, indicating that as the number

Table 2.6: 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

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 2.6, 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%.

CHAPTER 3
EFFECTIVE POLICIES FOR ADDRESSING THE BOOKING PROBLEM IN INTERMODAL BARGE TRANSPORTATION

This paper addresses a booking problem in intermodal barge transportation, where a carrier must decide whether to accept or reject sequentially arriving shipment requests in order to maximize profit using fixed, vessel-supported services with predefined schedules and routes. In addition to these services, the carrier may rely on outsourced services with higher costs, provided that accepting the request remains profitable, i.e., cost-efficient relative to the revenue generated. The problem is inherently dynamic, as requests arrive continuously over time, and is further complicated by uncertainty about the potential arrival of more profitable requests in the future. In inland waterway systems, this complexity is further amplified by fluctuating water levels, which directly affect vessel capacity and service feasibility over time. We develop a bin-packing-based framework to model this booking problem under various decision policies that differ in decision timing and in their ability to anticipate future shipment requests. In this context, we extend the classical bin packing problem to a time–space network setting, where "bins" represent scheduled, vessel-supported services defined not only by capacity and cost but also by route, departure time, and travel duration. "Items" represent shipment requests characterized by origin–destination, availability and due times, and economic value. Assignment decisions must satisfy spatial–temporal compatibility and profitability, whether using regular services or outsourcing alternatives. To capture the evolution of accepted requests and the operational impact of fluctuating water levels on vessel capacity, the model is embedded into a rolling-horizon framework. At each decision point, the model accounts for current requests, previously accepted requests, and, if applicable, predicted future requests and updated water-level forecasts. The performance of different booking policies is evaluated through extensive numerical experiments using a commercial solver.

3.1 Introduction

Intermodal freight transportation refers to the movement of freight using a combination of at least two distinct transportation modes, with transfers between modes occurring at designated intermodal terminals (such as seaports, inland ports, or rail yards) without the need to handle the freight itself (Bektaş et Crainic, 2008; SteadieSeifi *et al.*, 2014). Throughout the journey, each shipment typically remains within a single loading unit, most commonly a standardized container, to ensure continuity, reduce transshipment costs,

and facilitate interoperability across modes. Intermodal freight transportation systems are designed to capitalize on the unique advantages of each mode along different segments of the network: for example, road transport offers accessibility and flexibility for local pickup and delivery; rail provides cost efficiency and scalability for long-haul transport; and inland waterway transport offers environmental benefits and efficiency for medium-distance segments between inland terminals. A significant share of inland and medium- to long-distance intermodal freight is managed by consolidation-based carriers, which group shipments from multiple shippers into shared transportation services. This consolidation strategy enables carriers to leverage economies of scale, optimize capacity utilization, and reduce per-unit transportation costs. These carriers typically operate along fixed service routes with regular schedules defined over short-term planning cycles (e.g., weekly), which are repeated consistently over a longer planning horizon (e.g., a season) (Crainic et Rei, 2024). For instance, a barge-based consolidation carrier may offer weekly round-trip services connecting inland terminals such as Antwerp, Rotterdam, and Strasbourg, with predefined departure and arrival times at each terminal that remain fixed and are repeated weekly throughout the season.

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

However, because shipments are not yet physically available at the time of booking, they cannot be di-

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

One of the key challenges faced by carriers, yet insufficiently addressed in the existing literature, is the lack of effective decision-support tools to guide sequential acceptance decisions in dynamic booking environments. Existing methodologies for determining optimal shipment itineraries, most often used in operational planning, such as multi-commodity network flow models, are not well suited to this context, as they require detailed shipment itineraries. These models must specify exactly which terminals each shipment passes through, including the departure and arrival times at each visited terminal, the specific times at which loading and unloading occur, and the consolidation of shipments along the journey. Time–space networks are the modelling instruments typically used to capture this level of detail. To represent the temporal dimension, the network is divided into many small time intervals, creating a separate copy of each terminal for every discretized time step. This approach causes the model to become extremely large, often containing thousands of nodes and hundreds of thousands of arcs (i.e., arcs between different terminals representing movement over time and arcs between the same terminal representing loading, unloading, and holding activities over time). In large operational networks, adding separate decision variables for the routing of both contract-shipment requests (which is necessary because, when accepting non-contract-shipment requests, the carrier must remain committed to fulfilling all seasonal obligations for contract-shipment requests) and non-contract-shipment requests further increases the size, resulting in mixed-integer programming problems that are too large to solve within reasonable time.

This paper addresses this gap by introducing a novel framework that enables carriers to make informed booking decisions(i.e., to accept only profitable non-contract-shipment requests) based on conceptual and quantitative evaluations that estimate each request's likely shipment itinerary and its contribution to overall

profitability. In this framework, the objective is not to establish a detailed load plan at the time of booking, but rather to quickly assess origin–destination and time compatibility between shipments and scheduled services, along with an estimation of profitability. The proposed approach is designed for speed, responsiveness, and scalability, offering a practical alternative that is well aligned with the requirements of freight booking systems, where immediate decisions are often necessary.

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

The contributions of this paper are as follows: (i) We introduce a comprehensive booking-level decision problem for consolidation-based intermodal freight carriers operating over fixed service networks, where sequential accept/reject decisions must be made under incomplete information, accounting for both contractual shipment commitments and dynamic non-contract-shipment requests. (ii) We propose a bin packing–based modelling framework specifically tailored to the booking context. The framework is designed for rapid decision-making and scalability, making it particularly suitable for environments where immediate responses are required. (iii) We design booking policies reflecting different levels of predictability, including a myopic policy and a lookahead policy. (iv) While the proposed modelling framework is general, we apply it to inland barge transportation and evaluate the policies through comparative analyses that measure their impact on decision quality (e.g., acceptance rates) and overall profitability. (v) We conduct extensive compu-

tational experiments on diverse instances, analyzing the impact of the scheduled service network structure on the booking-level decision solutions. The analysis also examines how booking behavior differences (e.g., early vs. late booking tendencies), changes in the prediction confidence parameter, and environmental variations (such as smooth versus abrupt water-level fluctuation patterns) can affect the system's overall performance.

The remainder of the paper is organized as follows. Section 3.2 presents a general overview of decision levels in freight transportation systems, with a focus on consolidation-based carriers, and reviews the literature related to acceptance and rejection decisions in such systems. Section 3.3 describes the problem setting in detail. Section 3.4 introduces the proposed methodology and provides the mathematical model formulation. Section 3.5 reports and discusses the experimental results. Finally, Section 3.6 concludes the paper by summarizing the key findings and suggesting directions for future research.

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

The aim of this section is to position our study within the broader context of freight transportation planning and the different decision levels involved. Sub-section 3.2.1 provides an overview of these decision levels, with particular emphasis on their relevance to consolidation-based carriers. Sub-section 3.2.2 reviews the literature on acceptance and rejection, also referred to as demand control decisions in freight transportation.

3.2.1 Decision Levels in Freight Transportation

In a freight transportation system, particularly within the context of consolidation-based intermodal transport, one side consists of shippers who require the movement of freight between various origins and destinations. On the other side are carriers, particularly consolidation-based carriers for medium- and long-distance services, who operate scheduled services over a fixed network of terminals connected by predefined routes. Each service corresponds to a resource that follows a fixed route with an origin, a destination, and possibly intermediate stops, all with specified departure and arrival times. The physical and operational characteristics of these resources, such as load capacity and speed, determine the service's capacity and quality (e.g., express or standard). Effective planning is therefore required at multiple, interconnected decision levels, including tactical, booking, and operational, which together guide the carrier toward its ultimate objective of maximizing profitability.

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

At the booking level, the carrier begins interacting with actual demand in the form of booking requests from non-contractual shippers arriving throughout the season. These must be considered alongside the obligation to fulfill the shipment requests of contractual shippers, whose shipments (specified in long-term agreements) will become available at predetermined times and origin–destination (O–D) terminals during the season. Even under the assumption that contractual shipments will be on time and on weight, the booking level does not rely on their preplanned itineraries from the tactical planning level. Instead, contractual shipments are re-routed at the booking level only to allow flow redistributions that can make room for newly arriving non-contractual requests and, when necessary, to permit outsourcing of part of the contractual

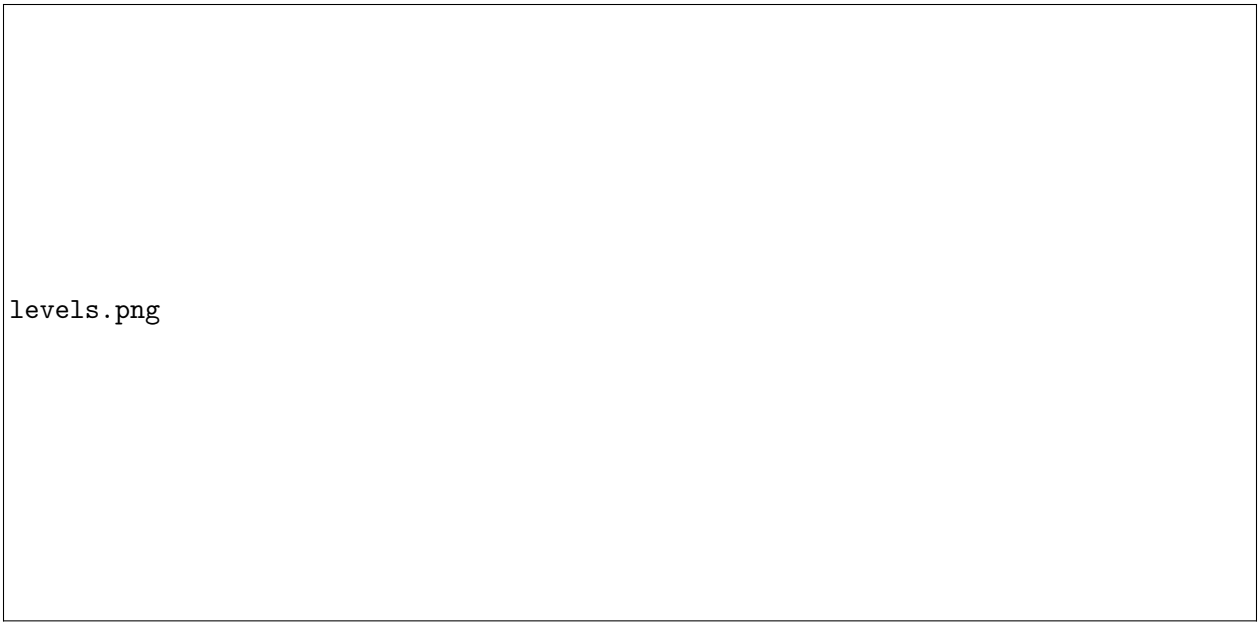

Figure 3.1: Multi-level decision-making framework

demand. Any acceptance of a non-contract-shipment request is therefore evaluated under the constraint that all contractual commitments will be satisfied, either by using internal service capacity or through outsourcing, even though no actual allocation decisions are made at this stage. Each non-contract-shipment request is evaluated in two steps before an accept/reject decision is made. First, feasibility is assessed by verifying space and time compatibility within the tactical plan. This involves identifying *shipment–service ways*—sequences of scheduled services with proper time and space matching—that can transport the shipment from its origin to its destination within its availability and due dates. Second, quantitative estimation of the operational cost of routing the non-contract-shipment request along these feasible shipment–service ways, possibly including potential outsourcing is conducted. The profitable requests are then accepted while remaining committed to all seasonal obligations for contract-shipment requests.

The decision to accept or reject a non-contract-shipment request follows predefined policies that govern how evaluations are conducted as requests are received during the season. These policies can be characterized by two main components:

- Timing of the decision-making: The evaluation may occur immediately upon non-contract-shipment request submission (order-based) or after a controlled delay that allows multiple non-contract-shipment requests to be assessed together (batch-based).

- Informational context used: The evaluation may rely solely on current and previously accepted non-contract-shipment requests (myopic approach) or may also incorporate forecasts of future non-contract-shipment requests (lookahead approach). In this case, such forecast requests are considered if their availability date occurs before the due date of the current request, even if their own due date extends beyond it.

By combining these two dimensions, different booking policies emerge, each with distinct implications for how non-contract-shipment requests are evaluated. Booking decisions are typically confirmed immediately or after a short controlled delay, before the shipment actually becomes available at the origin terminal.

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

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

3.2.2 Literature Review on Acceptance and Rejection Decisions

As discussed earlier, decision-making in freight transportation systems spans three interconnected levels, with the role of shipment routing differing at each stage. At the tactical level, shipment routing sup-

ports main tactical planning decisions such as selecting cost-efficient services and defining their schedules—referred to as the scheduled service network design. The booking level relies on this tactical plan but does not yet make routing implementation decisions; instead, it uses shipment routing as a decision-support tool to evaluate the time-space compatibility of each non-contracted shipment request with the selected service network and to estimate its expected profitability. At this level, accepted non-contract-shipment requests, together with all contract-shipment requests, are assumed to be accommodated later, but no specific resource-supported service capacity is yet assigned. The operational level then makes routing implementation decisions by allocating the capacity of available resource-supported services to shipments that are available for transportation operations. The existing literature has primarily addressed shipment routing in the context of tactical or operational planning, with less emphasis on its role in booking-level decision-making and resource allocation mechanisms within a revenue management framework. In the latter context—referred to here as an acceptance/rejection-based resource allocation problem—shipment routing functions as a decision-support mechanism for allocating limited resource-supported services, such as transportation services on single-leg or networked routes, with the objective of maximizing carrier revenue through optimal use of available capacity (Meng *et al.*, 2019). The remainder of this section focuses on acceptance/rejection-based resource allocation problems, as they closely resemble the booking-level decision-making process. Both involve evaluating each incoming request and deciding whether to accept or reject it based on feasibility and profitability considerations. The key distinction lies in the role of shipment routing: in booking problems, routing is used to assess both the feasibility and profitability of a request, whereas in acceptance/rejection-based resource allocation problems, routing determines the specific shipment itinerary through the scheduled service network, and the residual capacities of those services are updated once a shipment is accepted.

The literature on the acceptance/rejection-based resource allocation problem highlights two primary strategic approaches: booking limits and bid-price controls. Booking limits define the maximum number of shipments that may be accepted and typically reflect a first-come, first-served strategy, where requests are accepted sequentially until the limit is reached. In contrast, the bid-price strategy captures the opportunity cost of allocating capacity to a current request versus reserving it for a potentially more profitable future request (Littlewood, 1972). Several studies have compared these strategies. (Kapetanović *et al.*, 2018) analyze a bid-price-based dynamic acceptance/rejection strategy applied to a network train service. A dynamic programming formulation and deterministic approximations are used to demonstrate the potential revenue benefits of bid-price strategies compared to the traditional first-come, first-served (FCFS) strategy,

particularly under high-demand situations. Similarly, (Bilegan *et al.*, 2015) and (Wang *et al.*, 2016) use a bid-price-based dynamic acceptance/rejection strategy for rail and intermodal freight transportation, and show advantages in comparison with FCFS. Both papers propose a probabilistic mixed-integer program on a time–space network, incorporating the probability distributions of future requests (considering volume uncertainty) in a rolling horizon framework to handle dynamic request arrivals. (Wang *et al.*, 2017) formulate a dynamic and stochastic resource allocation problem in single-leg intermodal freight transportation as a Markov decision process, where multiple shipper arrivals with random shipment volumes are considered at each decision epoch, thereby improving capacity management and planning under uncertainty. In addition to the aforementioned strategies that highlight the importance of informational context—such as the contrast between FCFS, which bases decisions solely on current requests, and bid-price approaches, which incorporate expectations of future demand—some studies have emphasized the value of postponing decisions as a means to improve resource allocation outcomes. For instance, (Lee *et al.*, 2009) study a dynamic and stochastic resource allocation problem in single-leg maritime transportation and propose a stochastic dynamic programming model. The paper considers both contract-shipment requests and non-contract-shipment requests and introduces flexibility by allowing the carrier to postpone contractual shipment deliveries to better manage available capacity.

Service disruptions have not been extensively addressed in the existing literature on resource allocation models. Early studies such as (Wang, 2016a,b) were among the first to introduce resource allocation problems under service disruptions and random resource capacities. However, both works focus on static settings, dealing only with single-period resource allocation problems. The dynamic extension is proposed by (Wang, 2017), who formulate the multi-period version as a dynamic programming model. Their study shows that accounting for disruption risk enhances the robustness of booking and allocation decisions. The authors consider two distinct cases: under uniform resource consumption rates, they characterize the monotonicity of the optimal solution and propose a tailored backward induction algorithm. In contrast, for general consumption rates, they demonstrate through a counterexample that the optimal solution may not be monotone, and provide an alternative backward induction algorithm to solve the model. In the context of global synchromodal transportation, (Guo *et al.*, 2021) investigate a dynamic and stochastic shipment matching problem, in which a platform makes real-time decisions to accept or reject shipment requests and assign them to multimodal services (ship, barge, rail, truck), while facing uncertainties in both demand and travel times. The problem is formulated using a hybrid stochastic model, combining sample average approximation to represent the stochastic information of future requests and chance-constrained program-

ming to handle uncertain travel times. The model is embedded within a rolling horizon framework, where decisions are made over time as new information becomes available. Their model explicitly incorporates service disruptions and infeasible transshipments into the planning process, with the goal of maximizing total profit while minimizing infeasibility and delivery delays. In intermodal barge transportation systems, (Cui *et al.*, 2024) adopt a bid-price-based strategy and introduce a mixed-integer programming model on a time–space network that incorporates probability distributions of future shipment volumes. The model is also embedded within a rolling horizon framework to evaluate whether to accept or reject a current non-contract-shipment request in order to maximize expected revenue. In addition, the model allows rerouting of accepted but undelivered shipments to enhance capacity utilization. Importantly, the study addresses the impact of environmental disruptions—such as weather events and fluctuating water levels—by including both rerouting options and penalty costs for outsourced demand resulting from capacity shortfalls.

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

3.3 Problem Setting

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

Within this context, the booking system involves multiple shippers—including producers, wholesalers, and retailers—who request the transportation of their shipments across the network. These shipment requests are categorized as either contractual or non-contractual. Contractual shipments originate from long-term

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

Building on this system description,the booking problem is defined over a finite booking length (e.g., several days), within which booking decisions are made at specific decision points determined by the adopted policy. Under a request-based policy, each arriving request triggers an immediate decision, so that every request constitutes a decision point. Under a delay-based policy, the booking length is partitioned into fixed time intervals (e.g., half-day periods), during which multiple non-contractual requests accumulate and are jointly evaluated at the end of each interval. The request-based policy is not considered in this study, as request-by-request decision-making is incompatible with the operational characteristics of barge transportation systems. For instance, in contrast to road or air transport, where services typically operate at high frequency—such as daily or even hourly—barge services operate at low frequency, often on a weekly basis. Furthermore, the need for shipment consolidation requires that multiple requests be evaluated together over a defined decision interval, making batch-based evaluation more appropriate.

The composition of the item set and the associated service ways considered at each decision point depends on the adopted booking policy, which is defined by the informational context available. Under a myopic policy, only current and past information is used: the item set includes (i) all contractual items, which must be satisfied due to contractual obligations, and previously accepted non-contractual items, and (ii) the current batch of available non-contractual items. Under a lookahead policy, the informational scope is extended

to also include (iii) predicted non-contractual items that are expected to become available within the time window defined by the current batch—namely, the interval between the earliest availability time and the latest due time among current items.

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

This setting is modeled as a bin-packing-based framework, where shipment requests are represented as items to be packed into item-feasible service ways that act as bins with limited capacity. Spot-market options are also modeled as auxiliary bins, incurring higher outsourcing costs. Our methodology builds upon the classical bin-packing framework, extending it to address the specific requirements of the booking problem in consolidation-based freight transportation. Unlike most bin-packing formulations that operate over a single decision period, our models consider a multi-period planning horizon to capture the temporal dynamics inherent in real-world freight operations. To the best of our knowledge, time-window constraints—critical for aligning with shipper delivery requirements—are rarely incorporated in the bin-packing literature. Our work extends the study of Crainic *et al.* (2021a), which applied a multi-period bin-packing approach in a single-corridor setting, by generalizing it to a network context. Furthermore, the framework explicitly incorporates item profits and distinguishes between items, providing a richer and more realistic representation of demand heterogeneity.

The primary objective of the assignment model is to evaluate profitability rather than determine the physical arrangement of items—a focus more relevant to operational-level planning. In our case, item "size" is approximated solely by weight, which is consistent with the operational reality of inland waterway transport systems where service capacity is strongly influenced by water levels. Specifically, a service-way is considered feasible only if the total weight of assigned items remains within the service's adjusted capacity, which depends on the water level due to constraints such as vessel draught limitations and grounding

risk. Therefore, we do not adopt a full three-dimensional bin-packing formulation. This abstraction allows us to preserve computational tractability while still capturing the critical interaction between weight-based capacity and environmental limitations in barge transportation networks.

3.4 Methodology

Based on the problem definition introduced in Section 3.3, in this section we first define the notation and the parameters used and then present the mathematical formulation developed for the problem.

3.4.1 Notation

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

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

Shipment–Service ways. Let Σ denote the set of all scheduled services. A shipment–service way is an ordered sequence of one or more services from Σ that connects an origin-destination (OD) terminal pair, under a given schedule. Let $\mathcal{J}$ be the set of all such sequences in the network. The services within any sequence must be temporally and spatially matched (the arrival of each service precedes the departure of the next one, and consecutive terminals are consistent).

For each $j \in \mathcal{J}$, let $\mathrm{Serv}(j) = \big(\sigma_1(j), \sigma_2(j), \ldots, \sigma_{|\mathrm{Serv}(j)|}(j)\big)$, with $\sigma_k(j) \in \Sigma\ \forall k$, denote the ordered sequence of scheduled services composing j. The origin and destination terminals of j are given by the first and last services in the sequence, respectively, denoted $o(j)$ and $d(j)$.

Each scheduled service $\sigma \in \mathrm{Serv}(j)$ has a nominal capacity, an availability time, a travel time, and cost components (transportation and handling such as loading/unloading). The aggregated attributes of j are defined as follows: its *physical* capacity W_j equals the minimum capacity across all $\sigma \in \mathrm{Serv}(j)$. Since water level directly affects the draught and thus the effective load capacity of vessels, each W_σ is a function of the water level $l \in \mathcal{L}$, where $\mathcal{L}$ denotes the set of considered water levels (e.g., discretized levels over the planning horizon). Consequently, the capacity of j is also water-level-dependent and is expressed as: $W_j(l) = \min_{\sigma \in \mathrm{Serv}(j)} W_\sigma(l)$, ; the *temporal–spatial* attributes are given by t_j, the availability time of the first service, and α_j, the sum of the travel and terminal times of all scheduled services in the sequence; and its *economic* attribute C_j is the sum of transportation and handling costs of the scheduled services used by j. An additional cost component β_i is considered to capture the cost of using a spot market path by item $i \in \mathcal{I}$. Given an item $i \in \mathcal{I}$, we denote by $\mathcal{J}_i \subseteq \mathcal{J}$ the set of *item-feasible service ways* for i (i.e., service ways that match the origin and destination of item i and whose departure and arrival times fall within the item's availability and due time window): $\mathcal{J}_i = \big\{j \in \mathcal{J} : \mathrm{o}(j) = o(i),\ \mathrm{d}(j) = d(i),\ t_j \geq \underline{t}_i,\ t_j + \alpha_j \leq \bar{t}_i\big\}$.

3.4.2 Mathematical formulation

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

To formulate the three models, we first introduce the following binary variables:

- $y_{ij} = \begin{cases} 1, & \text{if item-feasible service way } j \in \mathcal{J}_i \text{ is selected for item } i \in \mathcal{I}, \\ 0, & \text{otherwise;} \end{cases}$
- $x_{ij} = \begin{cases} 1, & \text{if item } i \in \mathcal{I} \text{ is assigned to item-feasible service way } j \in \mathcal{J}_i, \\ 0, & \text{otherwise;} \end{cases}$
- $u_i = \begin{cases} 1, & \text{if item } i \in \mathcal{I} \text{ is assigned to a spot-market path,} \\ 0, & \text{otherwise.} \end{cases}$

3.4.2.1 The MBBP model

The MBBP can be formulated as follows:

$$\max \sum_{i \in \mathcal{I}} \phi_i \left[\sum_{j \in \mathcal{J}_i} x_{ij} + u_i \right] - \sum_{i \in \mathcal{I}} \beta_i \, u_i - \sum_{j \in \mathcal{J}_i} C_j \sum_{i \in \mathcal{I}} x_{ij} \tag{3.1}$$

subject to:

$$\sum_{j \in \mathcal{J}_i} x_{ij} + u_i = 1, \quad \forall i \in \mathcal{I}^C \tag{3.2}$$

$$\sum_{j \in \mathcal{J}_i} x_{ij} + u_i \leq 1, \quad \forall i \in \mathcal{I}^{NC} \tag{3.3}$$

$$\sum_{i \in \mathcal{I}} w_i \, x_{ij} \leq W_j(l) \, y_{ij}, \quad \forall j \in \mathcal{J}_i, \quad \forall i \in I \tag{3.4}$$

$$\sum_{i \in I} \sum_{\substack{j \in J_i: \\ \sigma \in \text{Serv}(j)}} w_i x_{ij} \leq W_\sigma(l), \quad \forall \sigma \in \Sigma \tag{3.5}$$

$$y_{ij}, x_{ij}, u_i \in \{0, 1\}, \quad \forall i \in \mathcal{I}, j \in \mathcal{J}_i \tag{3.6}$$

The objective function (3.1) maximizes the profit by adding the revenues from items assigned to item-feasible service ways and spot-market options, and subtracting the corresponding costs for using those service ways or spot-market paths. Constraint (3.2) ensures that each compulsory item is either assigned to a item-feasible service way or sent to the spot-market, while constraint (3.3) states that non-compulsory

items may or may not be assigned. Constraint (3.4) enforces the capacity limits of each item-feasible service way, explicitly considering the corresponding water level. Constraint (3.5) ensures that the total load assigned to all service ways sharing the same scheduled service does not exceed that service's physical capacity, thereby linking capacity usage across different item-feasible service ways that include the same scheduled services. Finally, constraint (3.6) defines all decision variables as binary.

3.4.2.2 The LBBP model

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

The LBBP can be formulated as follows:

$$\begin{aligned}\max \sum_{i \in \mathcal{I}^C \cup \mathcal{I}^{NC}} \phi_i \left[\sum_{j \in \mathcal{J}_i} x_{ij} + u_i \right] + \sum_{i \in \mathcal{I}^{PNC}} \theta_i \, \phi_i \left[\sum_{j \in \mathcal{J}_i} x_{ij} + u_i \right] \\ - \sum_{i \in \mathcal{I}^C \cup \mathcal{I}^{NC} \cup \mathcal{I}^{PNC}} \beta_i \, u_i - \sum_{j \in \mathcal{J}_i} C_j \sum_{i \in \mathcal{I}^C \cup \mathcal{I}^{NC} \cup \mathcal{I}^{PNC}} x_{ij}\end{aligned} \tag{3.7}$$

subject to:

$$\sum_{j \in \mathcal{J}_i} x_{ij} + u_i = 1, \quad \forall i \in \mathcal{I}^C \tag{3.8}$$

$$\sum_{j \in \mathcal{J}_i} x_{ij} + u_i \leq 1, \quad \forall i \in \mathcal{I}^{NC} \cup \mathcal{I}^{PNC} \tag{3.9}$$

$$\sum_{i \in \mathcal{I}^C \cup \mathcal{I}^{NC}} w_i \, x_{ij} \; + \sum_{i \in \mathcal{I}^{PNC}} \hat{w}_i \, x_{ij} \; \leq \; W_j(l) \, y_{ij}, \quad \forall j \in \mathcal{J}_i, \quad \forall i \in I \tag{3.10}$$

$$\sum_{i \in \mathcal{I}^C \cup \mathcal{I}^{NC}} \sum_{\substack{j \in \mathcal{J}_i: \\ \sigma \in \text{Serv}(j)}} w_i \, x_{ij} \; + \sum_{i \in \mathcal{I}^{PNC}} \sum_{\substack{j \in \mathcal{J}_i: \\ \sigma \in \text{Serv}(j)}} \hat{w}_i x_{ij} \; \leq \; W_\sigma(l), \quad \forall \sigma \in \Sigma. \tag{3.11}$$

$$y_{ij}, x_{ij}, u_i \in \{0, 1\}, \quad \forall i \in \mathcal{I}, j \in \mathcal{J}_i \tag{3.12}$$

The objective function (3.7) maximizes profit by adding the revenues from items assigned to item-feasible service ways and spot-market options, with predicted items scaled by their confidence parameter, and subtracting the corresponding spot-market and item-feasible service way assignment costs. Constraint (3.8)

ensures that each compulsory item is either assigned to an item-feasible service way or sent to the spot market. Constraint (3.9) allows non-compulsory and predicted items to be optionally assigned. Constraints (3.10) and (3.11) enforce individual item-feasible service way capacities and shared scheduled-service capacities, respectively. Finally, constraint (3.12) defines all decision variables as binary.

The LBBP-WL model. The LBBP-WL extends the LBBP by replacing the fixed service capacities with point-estimated capacities. Specifically, the parameters W_j and W_σ in constraints (3.10)–(3.11) are replaced by their point-estimated counterparts $\hat{W}_j(l)$ and $\hat{W}_\sigma(l)$, which reflect the variations in capacity induced by the water level l.

3.5 Experimental Results and Analysis

In this section, we present a series of computational experiments designed to assess the performance of the proposed decision-making models and to evaluate the influence of diverse operational and environmental conditions on booking outcomes. The analyses aim to provide both a quantitative comparison of alternative policies and qualitative insights into their managerial implications. The organization of this section is as follows: in Section 3.5.1 , we present the characteristics of the instances generated for the computational experiments—based on realistic cases—and the rolling-horizon booking simulation framework. We examine the behavior of the model in terms of computational time in Section 3.5.2. In Section3.5.3, we define the performance indicators required for analysing the computational results. Five experimental settings are then examined in Subsections 3.5.3. The first compares myopic, lookahead, and lookahead-with-water-level-variation policies to assess the value of anticipating future item and environmental changes. The second explores how booking behavior heterogeneity, including early and late booking tendencies, affects acceptance rates and profit. The third examines the impact of tactical planning choices—such as service selection and fleet composition—on downstream booking efficiency and profitability, highlighting the interdependence between the planning levels. The fourth evaluates the system's sensitivity to the prediction confidence parameter, offering insights into its robustness against forecast uncertainty. Finally, the fifth investigates the impact of smooth versus abrupt water level variations on booking decisions.

3.5.1 Test instances

Shipment-service way generation. We define a set of scheduled services inspired by the structure introduced in Payami *et al.* (2025a), following a one-week horizon divided into 14 half-day periods, consistent

with observed practices of morning and afternoon departures. We apply the procedure detailed in Payami *et al.* (2025a), of which we just give the main lines. Each scheduled service corresponds to a tactical planning decision and is characterized by a fixed route connecting a predefined origin–destination terminal pair, an availability time selected from the discrete set of 14 half-day periods $\{0, 1, \ldots, 13\}$, and a travel duration selected from $\{2, 3, 4\}$, corresponding to a minimum of one day and a maximum of two days. Each service carries a fixed operational cost of 5 units and is assigned to a specific vessel type, either small or large, with small vessels having capacities in the range $[25, 35]$ and large vessels in $[35, 50]$. We consider three distinct scheduled service networks, each established at the tactical level and reflecting different assumptions about navigability at the time of planning. In the VC-U (Uniform) setting, all services have capacities drawn uniformly from $[25, 50]$, representing a balanced fleet with no structural bias. The VC-SF (Split Fleet) setting introduces a heterogeneous composition in which half of the services are operated by small vessels and the other half by large vessels, allowing the model to benefit from service flexibility across vessel types. Finally, the VC-R (Restricted) setting simulates reduced navigability by assigning 75% of the services to small vessels and only 25% to large vessels, resulting in an asymmetric and capacity-limited network. Based on each of these service sets, we construct shipment–service ways by combining temporally and spatially feasible sequences of scheduled services. As a result, both VC-U and VC-R yield 36 distinct service ways, while VC-SF produces 72.

Item generation. We generate a heterogeneous shipment request set composed of compulsory and non-compulsory items. Compulsory items ($\mathcal{I}^C$) represent contractual shipments that must be served, while non-compulsory items ($\mathcal{I}^{NC}$) may be accepted or rejected depending on their contribution to profitability. Each item i is defined by its availability time, due time, weight, and fare. Availability times are selected from the discrete set of 14 half-day periods, and due times are defined by adding a small offset $\Delta \in 2, 3, 4$, corresponding to a delivery window of 1 to 2 days after the item's availability time, thereby introducing variation in delivery time tightness. Item weights are set according to a uniform distribution $w_i \sim U\{5, 15\}$, and fares are calculated as $\phi_i = m \cdot b$, where $b \in \{1, \ldots, 9\}$ is a base fare randomly selected from this set, and $m \in \{1.0, 1.5, 2.0, 2.5\}$ is a multiplier capturing the service/booking combination (standard-early, express-early, standard-late, express-late). Both demand sets include the same 60 compulsory items. The difference lies in the composition of non-compulsory items. The *LF-480* set contains 480 non-compulsory requests, of which approximately 34% belong to the high-fare category (express-late), with the remainder distributed among the other service/booking combinations. The *HF-528* set contains 528 non-compulsory requests, with a substantially higher proportion-about 65%-in the high-fare category, the remainder again

covering the other combinations. Predicted items follow the same fare generation rules as observed items, while their weights are drawn from a uniform distribution $w_i \sim U(0, 15)$. This structure ensures that both compulsory and non-compulsory item patterns are consistent across experiments, while allowing variation in total volume, fare composition, and the share of high-value shipments.

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

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.

3.5.2 Model performance

Tables 3.1 to 3.3 summarize the computational characteristics of the three models under study—MBBP, LBBP, and LBBP-WL—across different network topologies (VC-U, VC-SF, and VC-R) and two shipment-request configurations (LF-480 and HF-528). For each model and instance, the tables report the average and standard deviation of (i) the number of decision variables, (ii) the number of constraints, and (iii) the total solution time in seconds. These metrics serve to assess and compare the scalability and computational efficiency of each model formulation under varying problem sizes and network complexities.

Tables 3.1–3.3 reveal several consistent patterns across the three models—MBBP (myopic), LBBP (lookahead), and LBBP-WL (lookahead with water-level variation)—that reflect both the structural characteristics of the formulations and the computational effects of the rolling-horizon approach. Increasing the shipment-request size from LF-480 to HF-528 leads to a clear expansion in problem size, with higher numbers of decision variables and constraints, and, in most cases, longer average solve times. This is expected because larger request sets create more item–service-way combinations, thereby increasing the number of binary selection decisions. The effect is particularly pronounced in the VC-SF topology, which consistently yields the largest problem sizes and often the longest solution times; this topology offers more shipment–service-way options, resulting in a denser assignment structure. Comparing models, LBBP consistently produces the largest formulations because it accounts for both revealed and predicted requests at each decision stage. Incorporating predictions increases the number of possible states for assigning requests to services and adds more capacity-related constraints. MBBP, by contrast, works only with revealed requests and thus has fewer variables and constraints, resulting in smaller problem sizes and generally shorter solve times. LBBP-WL lies between these two in size, as the introduction of water-level-dependent capacity constraints removes some assignment possibilities when low-water scenarios occur, effectively reducing available service capacity. Interestingly, larger problem size does not always translate into proportionally higher solution times. A key reason lies in the nature of the bin packing problems that are solved iteratively within the rolling-horizon procedure. The computational difficulty of bin packing models often stems not from the number of items or bins alone, but from the degree of symmetry in the solution space—namely, the number of equivalent solutions (in terms of objective value) that differ only in their decision variables (i.e., the

specific item-to-bin assignments). When there is a large number of small items relative to bin capacity (as in the lookahead models, which incorporate predicted items), and when bins are highly homogeneous—such as in VC-SF or VC-R configurations where bins share identical or nearly identical capacities—the number of equivalent item-to-bin assignments increases substantially. This symmetry can make some instances significantly harder to solve, even when the nominal problem size is smaller. This behaviour aligns with what we observe in Tables 3.2 and 3.3, where LF-480 requires longer solution times than HF-528 despite involving fewer items. Overall, the computational results align with expectations: LBBP is the largest and often the slowest due to its predictive nature, MBBP is the smallest and generally the fastest, and LBBP-WL occupies a middle ground—except in cases where its additional constraints significantly accelerate convergence by tightening the feasible region.

Table 3.1: MBBP performance with shipment-request/Item structure

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

Table 3.2: LBBP performance with shipment-request/Item structure

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

3.5.3 Rolling Horizon Experiment

3.5.3.1 Performance metrics

Before presenting the performance metrics, we first recall that the booking process is dynamic: demand arrives over a finite booking length, and booking decisions must be made at specific points in time. To

Table 3.3: LBBP-WL performance with shipment-request/Item structure

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

handle this sequential structure, the proposed bin packing models are solved iteratively within a rolling-horizon procedure. In our setting, the booking length is divided into five decision stages to reflect the sequential nature of booking arrivals and to capture the operational reality that barge services typically experience a small number of well-defined booking waves over a weekly cycle, ranging from early to late requests. Using five stages therefore provides a realistic level of temporal granularity—sufficiently fine to capture booking dynamics while remaining consistent with operational practice—and offers a consistent basis for comparing the performance of alternative booking policies. The following metrics are therefore used to evaluate system performance across the five decision stages:

- **Total profit:** The overall profit obtained by the carrier calculated as the total revenue minus the total costs. This metric reflects the system's ability to maximize economic efficiency while limiting the use of costly alternatives (i.e., spot-market paths).
- **Acceptance rate:** The ratio of all accepted non-compulsory items to the total number of non-compulsory items offered throughout the horizon.
- **Total cost:** The cumulative cost incurred over the horizon, comprising:
  - *way cost:* The total cost associated with assigning items to shipment-service ways.
  - *path cost:* The total cost associated with routing items via spot-market paths.

3.5.3.2 Impact of Anticipative Information on Booking System Performance

In this subsection, we assess the use of the three proposed models within the context of the rolling-horizon procedure: a myopic model (MBBP), a lookahead model (LBBP), and a lookahead model with water-level

variations (LBBP-WL). The first model relies solely on realized information available at each decision point, whereas the second incorporates anticipative contextual information, such as expected future requests. The third model extends this anticipative framework by additionally accounting for random water-level variations that reduce available capacities, thereby capturing a more constrained and operationally realistic environment.

To facilitate interpretation, the experiment is organized into two distinct settings. The first setting corresponds to cases in which water-level variability is not considered; in this environment, MBBP and LBBP operate under identical capacity assumptions and can thus be directly compared in terms of decision quality and efficiency. The second setting explicitly incorporates water-level–induced capacity reductions. In this more restrictive environment, LBBP-WL is not evaluated as a competing policy against MBBP or LBBP, but is instead used to quantify how water-level variability alters booking decisions and overall system performance relative to a setting in which such variations are ignored.

Table 3.4 summarizes the key performance indicators (KPIs)—including total profit, way cost, path cost, and acceptance rate—highlighting how different levels of anticipative information influence booking performance across these distinct informational and environmental contexts.

Table 3.4 shows that the average profit in LBBP is notably higher than in MBBP (1580.66 vs. 1369.83). This improvement results from the anticipative nature of LBBP: by incorporating predictions of future requests, the model strategically accepts fewer current non-compulsory shipments (as reflected in its lower acceptance rate, 0.49 vs. 0.51 in MBBP) to preserve capacity for potentially more profitable future demand. Comparing MBBP with LBBP-WL reveals an even larger drop in acceptance rate (0.51 vs. 0.44), which reflects a combination of anticipative demand forecasting and more restrictive operational conditions. Similar to LBBP, the LBBP-WL model incorporates forecasts of future shipments, encouraging conservative acceptance decisions to maintain future flexibility. In addition, the updated water-level information available to LBBP-WL allows the model to anticipate reduced future capacity under low-water conditions. Consequently, it rejects more current requests to avoid future infeasibility, which directly contributes to its lower average profit relative to MBBP (1301.55 vs. 1369.83).

The comparison between LBBP and LBBP-WL highlights the impact of ignoring water-level variability and the operational consequences that emerge when capacity reductions are taken into account. In LBBP, service-

Table 3.4: KPI comparison for MBBP, LBBP, and LBBP-WL

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

way capacities remain fixed and optimistic, allowing the model to consolidate more shipments and rely more heavily on regular services. From a cost perspective, this results in the highest way cost (968.66), as the model can select more regular services to handle both current and predicted demand. In contrast, LBBP-WL updates service-way capacities based on anticipated low-water conditions, leading the model to expect tighter capacity in upcoming stages. As regular-service capacity becomes limited, the system must rely more frequently on spot-market services to serve already-accepted and contractual shipments. This shift produces the lowest way cost (458.66), due to reduced feasible capacity for regular services, but simultaneously results in a higher path cost because of the increased dependence on spot-market alternatives.

Overall, all three models remain profitable, but their outcomes reflect the informational and environmental conditions under which they operate. LBBP yields the highest profit because anticipative demand information enables it to make more selective acceptance decisions and reserve capacity for potentially more profitable future requests. MBBP, which relies solely on realized information, accepts a larger share of current requests but at the expense of lower overall profit. In contrast, the reduced profit observed in LBBP-WL is not due to weaker decision-making, but results from the more restrictive capacity conditions imposed by anticipated low-water levels, which limit feasible acceptance opportunities and increase reliance on expensive spot-market services.

3.5.3.3 Item impact

To isolate the effect of demand composition on model performance, we analyze how the fare structure of the shipment-request set influences booking outcomes. Specifically, we distinguish between two input categories: LF-480 (low-fare–dominant) and HF-528 (high-fare–dominant). Figure 3.2a and 3.2b provide a visual comparison of the total number of accepted items and the resulting profits across the three models for each input type. This comparison helps assess the sensitivity of each booking policy to variations in revenue potential, and highlights how predictive or capacity-constrained models adjust their decisions based

on expected profitability.

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

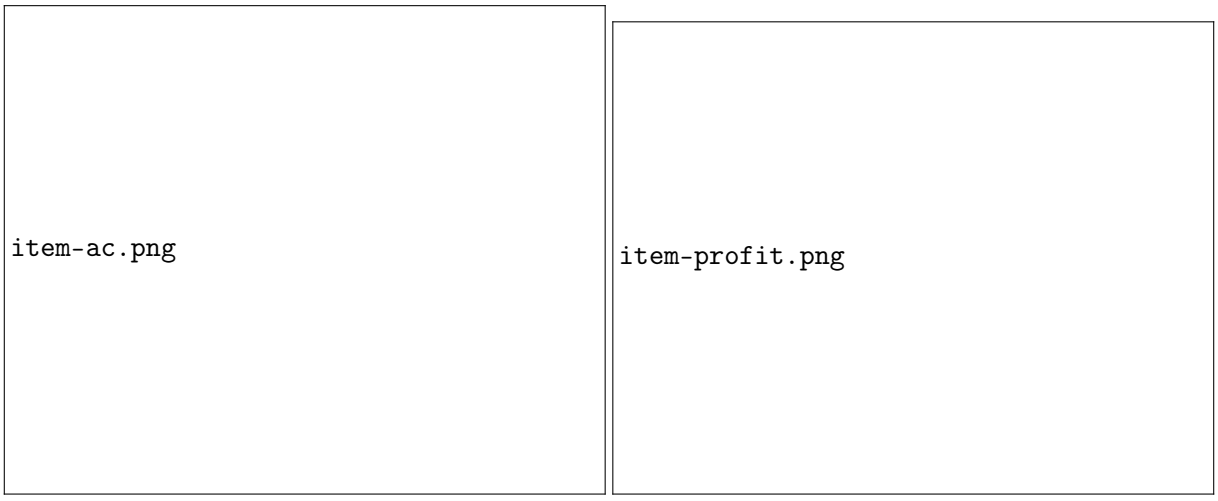

(a) Accepted item by model and item (b) Profit by model and item

Figure 3.2: Item impact

Overall, the results reveal a clear and systematic pattern driven by the revenue composition of the request set. When the input is dominated by low-fare items (LF-480), all models accept fewer shipments and generate lower profits, as each accepted item contributes little revenue. Under these conditions, the system becomes far more sensitive to operational restrictions: once capacity is reduced due to low water levels

(LBBP-WL), acceptance drops sharply. In contrast, when the request set consists primarily of high-fare items (HF-528), acceptance levels and total profits increase substantially. The high revenue per item makes each accepted booking worthwhile, enabling the lookahead model (LBBP) to better exploit profitable opportunities. Moreover, the negative impact of water-level–induced capacity reductions is noticeably weaker in this setting, because even a smaller number of accepted high-fare shipments is sufficient to maintain strong overall profitability. In other words, how strongly capacity reductions affect the system depends entirely on the revenue potential of the incoming items: they amplify losses when fares are low and dampen them when fares are high.

3.5.3.4 Shipment-request ways impact

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

Figure 3 illustrates the effect of two shipment-service way configurations on booking system performance across the MBBP, LBBP, and LBBP-WL models. The VC-R configuration represents a network with lower aggregate capacity due to a higher share of small vessel supported-services, whereas VC-SF provides substantially greater total network capacity.

In terms of accepted shipments (Figure 3.3a), increasing total capacity from VC-R to VC-SF leads to a rise in median acceptance across all models, but this increase is modest relative to the scale of capacity growth. For MBBP, the median rises from roughly 45 to 55, with the overall range expanding from about 20–70 in VC-R to 25–85 in VC-SF. In LBBP, the median increases from about 50 to 65, with the upper range exceeding 80, representing the largest change among the models. In LBBP-WL, the median grows from around 40 to 55 but remains below LBBP due to the additional consideration of water-level–induced capacity reductions. These patterns confirm that acceptance decisions are primarily driven by the economic value of shipments rather than the nominal capacity of the network, as even in the higher-capacity case, the system accepts only as many shipments as remain profitable.

For profit (Figure 3.3b), the increase from VC-R to VC-SF is more pronounced. In MBBP, the median profit rises from approximately 1200 to 1500–1800. In LBBP, the growth is more substantial—from around 1600 to over 2000, with a maximum near 2600. In LBBP-WL, the median improves from about 1100 to 1400–1500, with the maximum reaching around 2000. The primary driver of this improvement is the reduced reliance on costly outsourced paths in the VC-SF configuration. Higher in-network capacity allows a greater proportion of shipments to be routed internally at lower cost, thereby increasing total profit even if accepted items do not grow substantially.

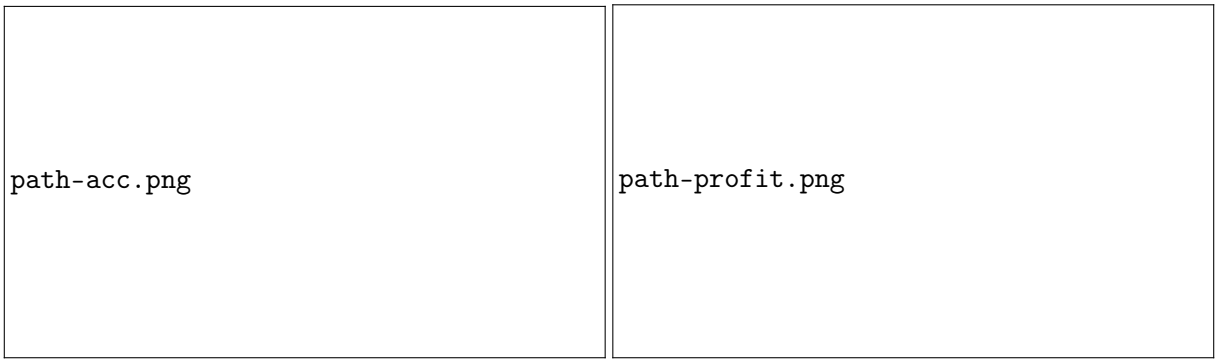

(a) Accepted item by model and ways (b) Profit by model and ways

Figure 3.3: Shipment-request ways impact

Overall, higher capacity (VC-SF) yields limited gains in accepted volumes but significant improvements in profit across all models. This outcome indicates that the booking system's acceptance policy is profitability-oriented rather than purely capacity-driven, with capacity expansion primarily enhancing economic efficiency by reducing dependence on expensive outsourced services.

3.5.3.5 Prediction confidence impact

This section explores how the level of confidence associated with predicted demand (θ) influences booking decisions and overall system performance. Table 3.5 presents the results across three confidence levels—$\theta = 0.25, 0.5$, and 0.85—for each combination of way configuration and item mix. The parameter θ reflects the model's trust in the accuracy of predicted item attributes and plays a crucial role in shaping acceptance behavior, particularly under uncertainty.

Table 3.5 shows that increasing prediction confidence (θ) from 0.25 (low confidence)to 0.85 (high confidence) systematically increases profit while slightly reducing acceptance rates across all path configurations

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

Table 3.5: Prediction confidence impact

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

3.5.3.6 Water level impact

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

Table 3.6: Impact of smooth ($\gamma = 0.75$) and abrupt ($\gamma = 0.35$) water-level fluctuations on system performance

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

Table 3.6 reports the evolution of profit, path cost, and way cost over five discrete decision steps under smooth and abrupt water-level fluctuations. In both scenarios, profit increases from step 1 to step 4 and then exhibits a slight decline at step 5. However, profit levels differ substantially across scenarios: under smooth fluctuations, profit remains consistently higher and reaches a maximum of approximately 1,350 at step 4, whereas under abrupt fluctuations the peak profit is lower, around 1,200. This gap directly reflects the impact of tighter capacity restrictions under low-water conditions, which constrain revenue generation even when demand remains available.

Path costs also rise with successive decision steps in both scenarios, but their magnitude is markedly higher under abrupt fluctuations. In the smooth case, path costs increase broadly in line with profit and reach approximately 1,400 at step 4. Under abrupt fluctuations, however, path costs escalate much more sharply, exceeding 1,700 at step 4 and clearly surpassing profit levels. This behavior indicates that severe capacity reductions force the system to rely on more expensive routing and allocation options, thereby eroding margins and overall profitability.

Way costs remain comparatively stable across decision steps, varying within a limited range in both scenarios. They are systematically lower under abrupt fluctuations, as stricter capacity constraints reduce the number of feasible service-way allocations, leading to lower overall way usage while shifting a larger share of demand toward costlier path-level decisions.

Overall, the results indicate that smooth water-level fluctuations allow the system to preserve a closer balance between revenues and costs, whereas abrupt fluctuations disrupt this balance by disproportionately increasing path costs and limiting achievable profits. The slight decline in profit observed at the final decision step in both scenarios further suggests diminishing returns once accumulated commitments and tight capacity limits reduce the flexibility of feasible allocations.

3.6 Conclusions

In this paper, we investigated the booking problem for consolidation-based intermodal freight transportation carriers. The study explicitly incorporated acceptance and rejection decisions for non-contractual shipment requests submitted throughout the season, while ensuring that the carrier fulfills all contractual shipment commitments. To the best of our knowledge, this is the first work to address booking-level control in such a setting using a bin-packing framework, thereby linking tactical planning with booking planning.

The proposed approach proceeds in two steps: first, identifying feasible service ways for each shipment item based on time–space compatibility; and second, applying bin-packing models to evaluate profitability and make accept/reject decisions for incoming requests. We proposed three models to capture different booking policies: a myopic bin-packing booking problem (MBBP), a lookahead bin-packing booking problem (LBBP), and an extended lookahead version with water-level constraints (LBBP-WL). These formulations provide a unified framework for analyzing how different informational contexts affect booking decisions and, ultimately, how acceptance and rejection choices influence overall system performance.

Computational experiments on realistic-size instances across multiple service topologies and request settings revealed consistent trade-offs between profit, acceptance rates, and computational complexity across the three models. While MBBP proved computationally efficient due to its limited scope, LBBP achieved higher profit by exploiting request predictions, albeit with increased problem size and solution times. Overall, the results emphasized the importance of anticipating request and accounting for capacity variability in order to make profitable accept/reject decisions. We further analyzed the impact of different settings incorporated into the problem, including request heterogeneity in terms of fare values (high-fare versus low-fare), variation in service size and fleet composition, the role of prediction confidence levels, and the effect of water-level changes under both abrupt and smooth scenarios. The results showed that these factors significantly influence the structure of optimal solutions, the utilization of available services, and ultimately the overall system profit.

Finally, several directions remain open for future research. Extending the framework to consider order-based booking decisions, rather than batch-based decisions, would enable a more detailed representation of how shipments arrive and are processed, and allow comparisons between the two approaches to evaluate how the timing of decision-making influences system profitability and booking structures. Another promising extension is to consider flexible delivery times, where late deliveries are permitted with explicit penalty costs. The modelling structure developed in this paper naturally supports such an extension: since each shipment is assigned to a service way with well-defined timing attributes, the model can be augmented to represent cases where the assigned service departs or arrives after the shipment's due time. This would allow late deliveries to be explicitly captured and penalty costs to be computed as a function of the degree of lateness (e.g., linearly scaled with delay duration). Implementing this feature would provide additional flexibility in planning and enable a systematic evaluation of how different penalty structures influence the profitability and feasibility of booking decisions. Furthermore, while contractual shipment requests were

treated as deterministic commitments in this study, relaxing this assumption and considering stochastic patterns even for contractual flows would provide a more realistic decision environment. Integrating such uncertainty into the accept/reject process for non-contractual requests, combined with efficient heuristic or decomposition methods, would improve scalability and extend the applicability of the models to larger and more complex problem instances.

CONCLUSION

Inland waterway freight transportation systems face structural and operational challenges that distinguish them from other transport modes. Among the most critical is the spatial and temporal variability of water levels, which imposes segment-specific navigational constraints on vessel operations. These variations arise from both natural phenomena—such as rainfall, snowmelt, and sedimentation—and anthropogenic factors including dam operations, irrigation withdrawals, and lock schedules. As a result, the available navigable depth differs across network segments, complicating the design of reliable and efficient service networks.

A central physical limitation introduced by this variability is the risk of grounding due to inadequate submerged clearance. As a vessel's load increases, its draught deepens, which may exceed the local water depth in shallow segments and render services infeasible. In this thesis, we focused exclusively on this grounding constraint, given its prominence and restrictive impact on tactical planning for consolidation-based carriers. In the first chapter, we extended the Scheduled Service Network Design with integrated Resource and Revenue Management (SSND-RRM) framework by embedding segment-level draught feasibility conditions. This enhancement ensures that service selection, demand routing, and vessel utilization decisions remain operationally viable under real navigational conditions.

Computational analyses on various network topologies—including linear and star topologies showed the limitations of tactical planning models that overlook water-level-dependent capacity constraints. Results showed that such omissions lead to overestimated profit margins, inflated shipper satisfaction, and unrealistic capacity utilization. These discrepancies were especially pronounced in linear topologies with limited rerouting flexibility. In contrast, star-shaped topologies exhibited more resilience due to centralized hub structures. Overall, the findings validate the necessity of integrating segment-specific water level considerations in service network planning for consolidation-based carriers.

In the second chapter, the Water-Level-Constrained Scheduled Service Network Design with Resource and Revenue Management (WL-SSND-RRM) problem was introduced to explicitly address uncertainty in water levels. Recognizing that fluctuating water depths influence available vessel capacity and route feasibility, we developed two stochastic formulations. The first penalizes unmet demand due to unforeseen water-level-induced infeasibility, while the second allows for dynamic reallocation of demand itineraries following the realization of water levels.

To solve these formulations efficiently, we proposed a decision-based scenario clustering method that reduces problem size while preserving decision quality. The results revealed the value of incorporating environmental uncertainty in tactical planning. Both stochastic models outperformed deterministic model. Among them, the dynamic itinerary adjustment model exhibited higher adaptability, better vessel utilization, and improved shipper satisfaction, particularly under adverse conditions. Scenario-specific evaluations confirmed that low water levels triggered a shift toward smaller vessel deployment, more flexible routing, and higher service costs, all of which impact profitability and network performance.

The third chapter focuses on the booking-level decision-making system for consolidation-based intermodal freight transportation carriers, where each incoming shipment request must be evaluated under tight time constraints and incomplete information about future demand. In this context, booking decisions must consider both the contractual obligations already in place and the arrival of non-contractual shipments throughout the season. The model developed in this chapter provides a structured and computationally efficient framework to assess the economic viability of each request before its physical entry into the system. At the core of the methodology lies a bin-packing structure used to identify feasible time–space ways for each shipment and determine accept/reject decisions based on profitability . Three distinct booking strategies are introduced, reflecting different levels of foresight: a myopic strategy that relies solely on current information, a lookahead strategy that incorporates forecasted future demand, and an extended lookahead variant that also integrates environmental uncertainties such as fluctuating water levels. These strategies allow carriers to tailor their responses to the degree of uncertainty and operational variability they face.

Results demonstrate that although anticipating future conditions may reduce the acceptance rate in certain cases, it typically leads to higher profitability. In particular, in scenarios where capacity becomes constrained due to environmental factors such as low water levels, lookahead-based strategies offer the flexibility needed to maintain robust system performance. The model also reveals key behavioral patterns across different demand mixes, network structures, and forecast accuracy levels, shedding light on how the system adapts to uncertainty. Overall, this chapter contributes a fast, scalable, and adaptive framework that bridges tactical planning with dynamic booking operations in intermodal container freight systems. It lays a solid foundation for the development of intelligent decision-making tools that are capable of responding to real-world operational constraints and variability.

BIBLIOGRAPHY

Agamez-Arias, A.-d.-M. et Moyano-Fuentes, J. (2017). Intermodal transport in freight distribution: a literature review. *Transport Reviews*, *37*(6), 782–807.

Akyüz, M. H., Dekker, R. et Sharif Azadeh, S. (2023). Partial and complete replanning of an intermodal logistic system under disruptions. *Transportation Research Part E: Logistics and Transportation Review*, *169*.

Aminzadegan, S., Shahriari, M., Mehranfar, F. et Abramović, B. (2022). Factors affecting the emission of pollutants in different types of transportation: A literature review. *Energy Reports*, *8*, 2508–2529.

Bai, R., Wallace, S. W., Li, J. et Chong, A. Y.-L. (2014). Stochastic service network design with rerouting. *Transportation Research Part B: Methodological*, *60*, 50–65.

Bakir, I., Erera, A. et Savelsbergh, M. (2021). Motor carrier service network design. In T. Crainic, M. Gendreau, et B. Gendron (dir.), *Network Design with Applications in Transportation and Logistics* chapitre 16, 427–467. Springer, Boston, MA.

Bektaş, T. et Crainic, T.G. (2008). A Brief Overview of Intermodal Transportation. In Taylor, G.D. (dir.), *Logistics Engineering Handbook* chapitre 28, 1–16. Taylor and Francis Group, Boca Raton, FL, USA.

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

Bilegan, I. C., Crainic, T. G. et Wang, Y. (2022). Scheduled service network design with revenue management considerations and an intermodal barge transportation illustration. *European Journal of Operational Research*, *300*(1), 164–177.

Braekers, K., Caris, A. et Janssens, G. K. (2013). Optimal shipping routes and vessel size for intermodal barge transport with empty container repositioning. *Computers in industry*, *64*(2), 155–164.

Caris, A., Macharis, C. et Janssens, G. K. (2012). Corridor network design in hinterland transportation systems. *Flexible Services and Manufacturing Journal*, *24*, 294–319.

CEIC Data (2023). United states inland waterways freight transport, total. Accessed from CEIC Data.

Chen, L. et Miller-Hooks, E. (2012). Resilience: An indicator of recovery capability in intermodal freight transport. *Transportation Science*, *46*(1), 109–123.

Christiansen, M., Fagerholt, K., Nygreen, B. et Ronen, D. (2007). Maritime transportation. In C. Barnhart et G. Laporte (dir.), *Transportation*, volume 14 de *Handbooks in Operations Research and Management Science* 189–284. North-Holland, Amsterdam.

Christodoulou, A., Christidis, P. et Bisselink, B. (2020). Forecasting the impacts of climate change on inland waterways. *Transportation Research Part D: Transport and Environment*, *82*, 102159.

Cordeau, J.-F., Toth, P. et Vigo, D. (1998). A survey of optimization models for train routing and scheduling. *Transportation science*, *32*(4), 380–404.

Crainic, T. (2003). Long-haul freight transportation. In R. Hall (dir.), *Handbook of Transportation Science* 451–516. Kluwer Academic Publishers, Norwell, MA, (second éd.).

Crainic, T. et Hewitt, M. (2021). Service network design. In T. Crainic, M. Gendreau, et B. Gendron (dir.), *Network Design with Applications in Transportation and Logistics* chapitre 12, 347–382. Springer, Boston, MA.

Crainic, T. et Kim, K. (2007). Intermodal transportation. In C. Barnhart et G. Laporte (dir.), *Transportation*, volume 14 de *Handbooks in Operations Research and Management Science* chapitre 8, 467–537. North-Holland, Amsterdam.

Crainic, T. et Rei, W. (2024). *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. Forthcoming in EURO Journal on Transportation and Logistics.

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

Crainic, T. G., Fu, X., Gendreau, M., Rei, W. et Wallace, S. W. (2011). Progressive hedging-based metaheuristics for stochastic network design. *Networks*, *58*(2), 114–124.

Crainic, T. G., Hewitt, M., Maggioni, F. et Rei, W. (2021b). Partial benders decomposition: general methodology and application to stochastic network design. *Transportation Science*, *55*(2), 414–435.

Crainic, T. G., Hewitt, M. et Rei, W. (2014). Scenario grouping in a progressive hedging-based meta-heuristic for stochastic network design. *Computers & Operations Research*, *43*, 90–99.

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

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

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

European Commission (2011). *White Paper 2011: Roadmap to a Single European Transport Area – Towards a Competitive and Resource-Efficient Transport System*. Rapport technique, Directorate for Mobility and Transport, Brussels.

European Environment Agency (EEA) (2024). *Greenhouse gas emissions from transport in Europe*. Rapport technique, EEA, Copenhagen.

Eurostat (2015). *Energy, Transport and Environment Indicators – 2015 Edition*. Rapport technique, Publications Office of the European Union, Luxembourg.

Eurostat (2023). *Inland Waterways Freight Transport Statistics – 2022*. Rapport technique, European Commission, Luxembourg.

Federal Statistical Office of Germany (Destatis) (2019). Freight transport on german inland waterways decreased by 11.1% in 2018.

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

He, Z., Navneet, K., van Dam, W. et Van Mieghem, P. (2021). Robustness assessment of multimodal freight transport networks. *Reliability Engineering and System Safety*, *207*.

Hewitt, M., Crainic, T. G., Nowak, M. et Rei, W. (2019). Scheduled service network design with resource acquisition and management under uncertainty. *Transportation Research Part B: Methodological*, *128*, 324–343.

Hewitt, M., Ortmann, J. et Rei, W. (2022). Decision-based scenario clustering for decision-making under uncertainty. *Annals of Operations Research*, *315*(2), 747–771.

Hoff, A., Lium, A.-G., Løkketangen, A. et Crainic, T. G. (2010). A metaheuristic for stochastic service network design. *Journal of Heuristics*, *16*, 653–679.

International Energy Agency (IEA) (2023). *Greenhouse Gas Emissions from Energy*. Rapport technique, IEA, Paris.

International Transport Forum (ITF) (2021). *Freight Transport: Bold Action Can Decarbonise Movement of Goods*. Rapport technique, OECD, Paris.

Ishfaq, R. (2017). Intermodal shipments as recourse in logistics disruptions. *Journal of the Operational Research Society*, *64*(2), 229–240.

Jiang, X., Bai, R., Wallace, S. W., Kendall, G. et Landa-Silva, D. (2021). Soft clustering-based scenario bundling for a progressive hedging heuristic in stochastic service network design. *Computers & Operations Research*, *128*, 105182.

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

Karam, A., Jensen, A. J. K. et Hussein, M. (2023). Analysis of the barriers to multimodal freight transport and their mitigation strategies. *European Transport Research Review*, *15*(1).

Kaut, M. (2014). A copula-based heuristic for scenario generation. *Computational Management Science*, *11*, 503–516.

Kaut, M., Vladimirou, H., Wallace, S. W. et Zenios, S. A. (2007). Stability analysis of portfolio management with conditional value-at-risk. *Quantitative Finance*, *7*(4), 397–409.

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

Lanza, G., Crainic, T., Passacantando, M. et Scutellà, M. (2024). *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.

Lanza, G., Crainic, T. G., Rei, W. et Ricciardi, N. (2021). Scheduled service network design with quality targets and stochastic travel times. *European Journal of Operational Research*, *288*(1), 30–46.

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

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

Lium, A.-G., Crainic, T. G. et Wallace, S. W. (2009). A study of demand stochasticity in service network design. *Transportation Science*, *43*(2), 144–157.

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

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

Müllerklein, D. et Fontaine, P. (2025). Resilient transportation network design with disruption uncertainty and lead times. *European Journal of Operational Research*, *322*(3), 827–840.

Payami, B., Bilegan, I. C., Crainic, T. G. et Rei, W. (2025a). *Service Network Design with Uncertainty on Water Levels for Intermodal River Transport*. Research Report CIRRELT-2025-16, Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique et le transport, Université de Montréal, Montréal.

Payami, B., Bilegan, I. C., Crainic, T. G. et Rei, W. (2025b). *Tactical Network Planning for Intermodal Barge Transportation Considering Varying Water Levels*. Research Report CIRRELT-2025-17, Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique et le transport, Université de Montréal, Université de Montréal, Montréal.

Prandtstetter, M., Widhalm, P., Leitner, H. et Czege, I. (2023). A novel visualisation tool for reliable inland navigation. *Transportation Research Procedia*, *72*, 3537–3544.

Riessen, B. V., Negenborn, R. R., Dekker, R. et Lodewijks, G. (2015). 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.

Rodrigue, J.-P. et Notteboom, T. (2024). *The Geography of Transport Systems* (6 éd.). London: Routledge.

Rousseeuw, P. J. (1987). Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, *20*, 53–65.

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

Sun, Y., Zhang, G., Hong, Z. et Dong, K. (2018). How uncertain information on service capacity influences the intermodal routing decision: A fuzzy programming perspective. *Information*, *9*(1), 24.

TRIP (2023). *America's Freight Network: Vital for Economic Growth and Quality of Life*. Rapport technique, TRIP Report, Washington, D.C.

Uddin, M. et Huynh, N. (2016). Routing model for multicommodity freight in an intermodal network under disruptions. *Transportation Research Record: Journal of the Transportation Research Board*, *2548*(1), 71–80.

Uddin, M. et Huynh, N. (2019). Reliable routing of road-rail intermodal freight under uncertainty. *Networks and Spatial Economics*, *19*(3), 929–952.

U.S. Bureau of Transportation Statistics (BTS) (2024). *Freight Facts and Figures – Long Range Forecast*. Rapport technique, U.S. Department of Transportation, Washington, D.C.

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

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

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

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

Wang, X., Crainic, T. G. et Wallace, S. W. (2019). Stochastic network design for planning scheduled transportation services: The value of deterministic solutions. *INFORMS Journal on Computing*, *31*(1), 153–170.

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

Wang, Z. et Qi, M. (2020). Robust service network design under demand uncertainty. *Transportation Science*, *54*(3), 676–689.

Ypsilantis, P. et Zuidwijk, R. (2019). Collaborative fleet deployment and routing for sustainable transport. *Sustainability*, *11*(20), 5666.

Zhang, R., Huang, C. et Feng, X. (2020). 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.

Zheng, Y. et Kim, A. M. (2017). 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.