The Impact of Autonomous Intra-Terminal Barge Concepts

A Case Study of the Port of Rotterdam

Graduation Assignment

Maartje Janszen

Delft University of Technology

TU Delft

Port of Rotterdam

The Impact of Autonomous Intra-Terminal Barge Concepts

A Case Study of the Port of Rotterdam

by

Maartje Janszen

in partial fulfilment opf the requirements for the degree of
Master of Science Mechanical Engineering

at the department Maritime and Transport Technology of Faculty Mechanical Engineering of Delft University of Technology (Mekelweg 2, 2628 CD Delft, The Netherlands).
To be defended publicly on Thursday November 27, 2025 at 10:00 AM.

Student number:	4869559	
MSc Track:	Multi-Machine Engineering	
Report number:	2025.MME.9123	
Version:	November 20, 2025	
Project duration:	June 1, 2025 – November 1, 2025	
Thesis committee:	Dr. B. Atasoy,	Chair and supervisor, Delft University of Technology
	Dr. J. Jovanova,	Supervisor, Delft University of Technology
	Dr. A. Napoleone	Delft University of Technology
	Dr. N. Pourmohammadzia	Delft University of Technology
Company supervisors:	Dr. ir T. Verduijn,	Port of Rotterdam, MAGPIE
	Prof. Dr. R.A. Zuidwijk,	Rotterdam School of Management, Erasmus University Rotterdam

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.

TU Delft

Port of Rotterdam

Preface

When I started this thesis, I was not very enthusiastic about the idea of spending so much time on one project by myself. Throughout my studies, I was always most motivated by working in teams, solving problems together and bouncing ideas around. A long, mostly individual project felt quite far from that. One thing was clear though: if I was going to do this, I wanted a topic where I could make some impact, where I could dream a little about what might be possible, speculate about what-ifs, and work with people outside TU Delft after spending more than enough hours at the university.

This thesis gave me that opportunity. It allowed me to work on a complex logistical problem at the scale of the whole port system rather than a single component. I am grateful to my graduation supervisor, Dr Bilge Atasoy, for her guidance and support. Our conversations were never just about the project plan or the next set of results. She also asked how I was doing as a person and shared experiences about things that women still encounter in STEM environments. Those honest talks made me feel less alone and more confident that what I noticed was real and worth addressing. As a running joke, the sun always seemed to shine during our meetings, no matter what the weather was doing before or after, which somehow matched the tone of those conversations, as I usually left them feeling a bit lighter and more confident.

I would like to thank the external partners involved in this project, in particular the people from the port and organisations who made time for meetings, interviews and discussions. Their openness and practical insights made the research much more concrete and reminded me that behind every model there is a messy, real system with people trying to make it work.

Doing a thesis can feel like a lonely battle, but my friends made sure it did not stay that way. Our very strict coffee break rhythm kept me sane. The collective sighing and “neerploffen” together on the Leeghwater couch helped more than I can describe. That couch has seen me stressed, tired, proud, and everything in between, and it was rarely empty. Thank you to everyone who joined me there with coffee, snacks, gossip and bad jokes, who listened to me complain about simulations and LaTeX, and who reminded me to go home at some reasonable hour. You turned a solitary project into a shared daily routine.

Finally, I want to thank my family for their endless support. Thank you for listening to my explanations about ports, barges and simulations, even when it was probably not very clear what I was talking about, and for reminding me to rest, sleep and keep things in perspective. Knowing that you were there in the background made it much easier to keep going.

Writing this thesis has been intense, frustrating, fun and rewarding, often all in the same week. It marks the end of my time as a student in Delft and the start of a new chapter that is not fully defined yet. I hope that reading this work is at least a bit more straightforward than writing it was, and that it may offer something useful to anyone interested in ports, logistics or autonomous inland shipping.

Maartje Janszen
Delft, November 2025

Abstract

Global container trade continues to grow while climate policy tightens emission and efficiency requirements in major hub ports. European ports such as Rotterdam must accommodate higher volumes and stricter reliability expectations without proportional expansions of quay length or yard space. Within this context, intra-port inter-terminal container exchanges remain a bottleneck. Inland barges provide a low-emission alternative to trucks, but fragmented planning and long waiting times undermine reliability and the business case for modal shift. This thesis examines whether autonomy-enabled modular splitting of barge calls can improve operational performance at the Maasvlakte container terminals.

A port-wide discrete event simulation is developed with deep sea, feeder, conventional barge, and autonomous module services calling five terminals that share a dedicated module-crane pool. Current multi-stop barge operations are compared with scenarios in which barges detach short-calling modules under different barge–module mixes, sea-freight demand levels, module-crane inventories, and module capacities. Performance is assessed using throughput, turnaround and waiting times, berth occupancy, crane utilisation, and anchorage behaviour.

In the recommended 50-50 barge-module mix, modular splitting increases port-wide throughput by about 12% and reduces barge turnaround by more than half compared with current operations, while deep sea and feeder vessels remain largely unaffected. These gains arise from parallel module calls that use residual quay pockets more effectively. Hybrid fleets with around half of inland work carried by modules therefore provide the best compromise between higher throughput and manageable growth in vessel calls. Configuration experiments indicate that two dedicated module cranes per large terminal are sufficient under the tested loads and that medium-sized modules around 24 to 36 TEU perform robustly. Within the limits of the stylised simulation, the results indicate that modular intra-port services can improve inter-terminal operational performance, support port sustainability and modal-shift objectives, and provide guidance for the design of MAGPIE Demo 6.

Keywords: Autonomous e-barges; modular splitting; intra-port transport; discrete event simulation; Port of Rotterdam.

List of Figures

List of Tables

Contents

Nomenclature

Abbreviations

Abbreviation	Definition
ABM	Agent based modelling
AGV	Automated Guided Vehicle
AIS	Automatic Identification System
API	Application Programming Interface
ASC	Automated stacking crane
ASV	Autonomous Surface Vessel
ATB	Actual Time of Berthing
ATD	Actual Time of Departure
BPMN	Business Process Model and Notation
CATWOE	Customers, Actors, Transformation, World view, Owners, Environmental constraints
CCNR	Central Commission for the Navigation of the Rhine
CER	Container Exchange Route
COLREGs	Convention on the International Regulations for Preventing Collisions at Sea
CT	Container terminal
DEA	Data Envelopment Analysis
DES	Discrete Event Simulation
DS	Deepsea
EDD	Earliest due date
ETA	Estimated time of arrival
ETB	Estimated Time of Berthing
ETD	Estimated Time of Departure
FCFS	First come, first served
GHG	Greenhouse gas
IMO	International Maritime Organization
ITT	Inter terminal transport
IWT	Inland Waterway Transport
LOA	Length Overall, the maximum length of a ship or boat measured in a straight line from the very front of the bow to the very back of the stern
KPI	Key Performance Indicator
MAGPIE	sMArt Green Ports as Integrated Efficient multimodal hubs (Horizon 2020)
MASS	Maritime Autonomous Surface Ship
MC	Modular crane
MCA	Melding Container Achterland
MHE	Material Handling Equipment
MILP	Mixed integer linear programming
OECD	Organisation for Economic Co operation and Development
P50	50th percentile (median)
P90	90th percentile
PoR	Port of Rotterdam
PTA	Planned time of arrival
PTD	Planned time of departure

Abbreviation	Definition
QC	Quay crane
RMG	Rail mounted gantry crane
ROI	Return on investment
RTA	Requested Time of Arrival
RTG	Rubber tyred gantry crane
SDG	Sustainable Development Goal
SLA	Service level agreement
SPT	Shortest processing time
SQ	Subquestion
STO	Strategic, Tactical, Operational (decision hierarchy)
STS	Ship to shore crane
TEU	Twenty foot equivalent unit
TEU-km	Twenty foot equivalent unit kilometre
TOS	Terminal Operating System
ULCV	Ultra large container vessel
USV	Unmanned Surface Vessel
VTS	Vessel Traffic Service
WSPT	Weighted shortest processing time

Definitions and terminology

Terminology	Definition
Container transhipment	Includes intra-port movements between terminals. This is commonly referred to as inter-terminal transport
Inland waterway transport	Freight transport using inland vessels within the port area or to the hinterland
Autonomous e-barge	An electrically propelled barge with automation features that reduce crew-related coordination constraints; control-level autonomy is not modelled

1

Introduction

Ports are critical nodes in global logistics networks, responsible for efficiently handling and transhipping vast volumes of containerised goods. As trade volumes grow and sustainability targets become more stringent, port performance, and in particular the coordination of container transhipment, becomes increasingly important. This chapter introduces the context, motivation, and scope of the research. It begins by outlining the operational and strategic relevance of barge-based container handling and transhipment, and then defines the research problem, objectives, and methodology.

1.1. Research Context

The continuous growth in global trade and the increasing demand for sustainability pose growing pressure on container ports to improve operational efficiency, while reducing environmental impact. Container throughput in global ports is projected to increase significantly, by three to six times by 2050, according to long-term forecasts of the Organisation for Economic Co-operation and Development (OECD, 2012). Ports must therefore adapt to this growth by implementing innovative technologies and strategies to enhance capacity and resilience, without proportional physical expansion.

One area of particular importance is the internal movement of containers between terminals, known as container transhipment. Within large multi-terminal ports, this process is crucial for enabling transhipment between modes, such as deep-sea to inland, supporting hub-and-spoke vessel strategies, and reducing dependency on truck-based transfers. Inland Waterway Transport (IWT) provides a low-emission alternative for such movements, particularly for high-volume or medium-distance transport within or between port clusters. However, despite these advantages, inland barge operations remain affected by systemic inefficiencies.

In ports such as the Port of Rotterdam, barge-based transhipment frequently suffers from fragmented planning processes, unpredictable berthing, underutilised quay space, and long vessel waiting times. According to Shobayo and Van Hassel (2019), barges in the Port of Rotterdam spend up to 60% of their total time in port waiting or idle, resulting in substantial delays and reduced reliability. These issues are not merely operational, they undermine the competitiveness of IWT, contribute to excess emissions, and complicate efforts to shift traffic away from congested road corridors such as the A15 Motorway (Froeling et al., 2008).

In response to these challenges, several initiatives have emerged. The Dutch digital coordination platforms Nextlogic and Portbase aim to synchronise barge and terminal planning through improved data sharing and optimisation algorithms. However, these platforms operate within the existing logic of conventional barge operations, and their effectiveness depends on the willingness and capacity of stakeholders to participate in centralised coordination.

Meanwhile, the MAGPIE project (Smart Green Ports, Horizon 2020)[1] has launched a broader effort to investigate transformative innovations for future-ready ports. This research contributes to Work Package

[1] https://www.magpie-ports.eu/work-package-presentation/

5 on Maritime and Inland Waterway Transport, specifically to Demo 6, which addresses Autonomous E-barge and Transhipment. This demonstration explores how automated barge transport and modular vessel concepts might improve intra-port logistics by enabling more scalable and efficient transhipment solutions. This thesis supports that initiative by evaluating the operational effects of such coordination strategies through simulation.

1.2. Research Field

This research sits at the intersection of port logistics, vessel coordination, and operations research. It engages with work on container terminal planning, intermodal synchronisation, and the design of digitally enabled transport systems.

Within this field, studies recognise the environmental and spatial advantages of inland waterway transport over road, including lower external costs for emissions, congestion, and noise (European Commission Mobility and Transport and Essen, 2018). At the same time, authors point to service reliability issues for small and port-internal operations, including waiting times and crew-related costs (Interreg, 2021; Alzate et al., 2024). Methodologically, discrete-event simulation is widely used to analyse time-dependent terminal processes such as berth allocation, vessel handling, and crane scheduling (Negenborn et al., 2023; A. Douma, 2008).

What is less developed is a system-level assessment of port-internal container transhipment that compares alternative vessel coordination strategies, including autonomy-enabled concepts, using a common performance framework. This thesis adresses to that gap. The Port of Rotterdam is used as a representative case, while the analysis and metrics are designed to be generalisable.

1.3. Research Problem

Barge-based transhipment is a spatially efficient and low-emission way to move containers within large ports. Yet, practice shows long and uncertain waiting times, fragmented berth planning, and underused quay capacity. These frictions reduce reliability for operators, create peak clustering for terminals, and hinder modal-shift objectives for authorities.

Existing work often addresses the technical feasibility of autonomous or modular vessels, or it models isolated terminal processes. What remains limited is quantitative evidence at system level that compares autonomy-enabled coordination strategies against current practice for port-internal transhipment. In particular, it is not well understood under which conditions alternative strategies reduce turnaround time, raise berth utilisation, or lower emissions.

This thesis responds to that gap with a simulation-based evaluation framework that compares vessel coordination strategies, including concepts that consider autonomous e-barges, on a shared set of key performance indicators. The model captures vessel scheduling, quay availability, handling time variation, emissions metrics, and scenario-based demand patterns. The Port of Rotterdam provides a representative case, while the design of the analysis and metrics aims to deliver generalisable insight.

1.4. Research Objective and Scope

The objective is to evaluate how autonomous e-barge coordination strategies affect the operational efficiency of intra-port container transhipment. The study uses discrete-event simulation to compare alternatives on common performance indicators and to explore the conditions under which each strategy performs well.

The analysis contrasts present-day multi-stop barge calls with two innovation-oriented concepts. Scenarios are used to represent distinct coordination logics, which are then tested under varying demand and resource settings. The strategies considered are:

1. Baseline: current multi-stop barge calls with limited cross-operator coordination.
2. Roundtrip strategy: fixed cyclic routes with coordinated schedules and terminal time windows.
3. Modular e-barge strategies, two variants:

(a) Modular splitting: a convoy arrives as coupled modules, which detach to berth at different terminals and are reassembled afterwards.

(b) Modular interchange: modules can be swapped between convoys or staged for later pickup to reduce berth conflicts and idle time.

The Port of Rotterdam is used as the case due to its scale and operational complexity. The modelling approach, the scenario logic, and the indicators are defined to support general insights for other large ports that consider automation and coordination improvements.

The study is bounded as follows. Geographically, it covers the Maasvlakte area only. In terms of modal scope, it focuses on intra-port container exchanges by inland barge; very small or time-critical exchanges may be routed by truck outside the simulation, while rail and deep-sea operations are treated as exogenous. Technologically, autonomy is represented as reduced coordination constraints, rather than control-level automation. For performance, the analysis considers turnaround time, waiting, berth utilisation and crane productivity, while broader economic and regulatory analyses lie outside the scope.

1.5. Research Questions

This research is guided by the following overarching question:

> *What is the impact of autonomous e-barge coordination strategies on the operational efficiency of intra-port container transhipment?*

To answer this main question, the study addresses the following subquestions:

1. **SQ1:** What are the current operational characteristics and bottlenecks of barge-based intra-port transhipment?
2. **SQ2:** Which autonomous or coordination-enhanced vessel strategies are applicable and how do they differ from current practice?
3. **SQ3:** How can discrete-event simulation be used to model port-intra-port transhipment processes under varying demand and scheduling conditions?
4. **SQ4:** What are the effects of different coordination strategies on key performance indicators such as turnaround time, quay occupancy, and emissions?
5. **SQ5:** How robust are the different strategies when demand levels, vessel types, or terminal capacity vary?
6. **SQ6:** What coordination strategy offers the best balance between operational performance and implementation feasibility?

These questions guide the structure of the research and its contribution to both academic inquiry and practical innovation in port logistics.

1.6. Research Methodology and Structure

Building on the objective in Section 1.4, this study adopts a theory-oriented case approach and uses discrete-event simulation (DES) as the primary method. DES is selected, since it represents time-dependent processes, stochastic arrivals, queuing, and resource constraints typical of port operations, which are essential to evaluate coordination strategies under varying conditions.

The research process comprises six steps:

1. System analysis: Characterise the current transhipment system, including barge flows, quay layouts, and operational bottlenecks.
2. Scenario design: Define the baseline, roundtrip, and modular concepts and motivate the choice of scenario.

3. Model development: Implement a model with arrival, berthing, and handling logic, including berth-admission and queuing policies.
4. Verification and validation: Verify internal logic and validate outputs against plausible ranges and expert judgement.
5. Verification and validation: Establish internal correctness and external plausibility, with detailed checks and traces.
6. Performance analysis: Comparing baseline and autonomous E-Barge on shared indicators relevant to terminals, barge operators, and the port authority.
7. Sensitivity testing and synthesis: Test robustness under alternative demand, capacity, and handling settings, then synthesise results into design insights and recommendations.

Figure 1.1 summarises the workflow from system definition to scenario evaluation and synthesis.

Figure 1.1: Research methodology

The document structure follows this workflow. Chapter 2 synthesises the literature on inland waterways, terminal operations, intra-terminal transport and performance measurement, and distils the operational characteristics and bottlenecks that, together with the case description, answer SQ1. The conceptualisation chapter defines the coordination strategies, clarifies their intended mechanisms, and motivates the experimental focus on modular splitting; by contrasting the alternatives with current practice, it answers SQ2. Chapter 4 sets the system boundary for the Maasvlakte case and provides data and assumptions for the model; together with Chapter 5, which details the discrete-event simulation design, implementation and experimental set-up, and the dedicated Verification and Validation chapter, these chapters answer SQ3. Chapter 7 reports the quantitative comparison of the modelled strategies on the agreed indicators (turnaround time, waiting, berth utilisation, crane productivity), which provides the evidence for SQ4, and presents the sensitivity analyses for demand, capacity and module settings that answer SQ5. Chapter 8 interprets trade-offs and feasibility and synthesises the recommended coordination logic, thereby answering SQ6. Chapter 9 consolidates the findings and states the overall answer to the research question.

2

Literature Review

This chapter reviews literature on efficient barge planning for container transport in large, multi-terminal ports. It first outlines how container flows are organised and how governance assigns responsibilities, then examines barge and terminal operations, planning structures and rules, inter-terminal transport, and a stakeholder-based KPI framework. Suitable simulation approaches are discussed, followed by research gaps that motivate the experiments developed later in this thesis.

2.1. Container Transport in Ports and Inland Waterways

The movement of containers through seaports involves a sequence of tightly connected processes spanning maritime arrival, terminal handling, and hinterland dispatch. Containers arriving on deep-sea vessels are discharged, sorted, and transferred to the appropriate mode; truck, rail, or barge. In high-throughput hub ports such as Rotterdam, transhipment between terminals and onward inland distribution represent a substantial share of total volume.

Container handling begins with vessel arrival and berthing, followed by quay crane operations. Terminal vehicles, such as Automated Guided Vehicles (AGVs), straddle carriers, and terminal trucks, transfer containers to the yard for stacking and retrieval by RTG or RMG cranes. The final leg is dispatch to hinterland by truck, rail, or barge. Figure 2.1 summarises the process and modal transitions.

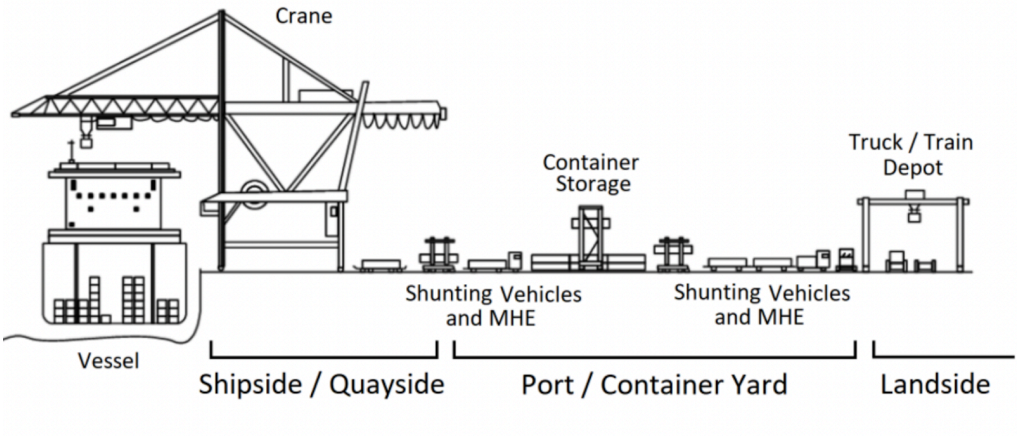

Figure 2.1: Overview of a container terminal (A. Y. Cil, Abdurahman, and I. Cil, 2022).

Hinterland transport by barges, IWT, provides high-capacity and sustainable connectivity between terminals and inland destinations (Konings et al., 2010). Compared with road, barges reduce congestion and emissions and suit bulk container flows over predictable routes, while requiring less fixed infrastruc-

ture than rail. Constraints include lower speeds, lock and bridge restrictions, and coordination across multiple independent terminals. Studies in Rotterdam indicate that up to 60% of barge port time is spent waiting at terminals or sailing between them (Port of Rotterdam, 2025a). Digitalisation and data-sharing initiatives such as *Nextlogic* aim to improve reliability and reduce idle time, while pilots with autonomous and electrified barges target further efficiency and environmental gains (Interreg, 2021; AUTOBarge, 2025).

When container discharge and pick-up locations differ, inter-terminal transport (ITT) is used. Mode choice reflects flow characteristics and infrastructure: small, heterogeneous, or time-critical flows are typically moved by truck, which allows quick consolidation without berth access; barges are preferred when water-connected terminals exchange larger volumes that justify a berth call; and rail is used where on-dock access and timetables support block movements between rail-served terminals. Empty containers add a separate layer: they are stored at dedicated depots and repositioned to terminals as needed, often by truck at short notice, and by barge when volumes build.

Synchronisation challenges include limited quay availability at peaks, uncertainty in container readiness, heterogeneous terminal rules and windows, and fixed nautical resources. Barges often serve multiple terminals per rotation, which makes sequence choice and coordination pivotal; according to the Barge Performance Monitor, waiting and inter-terminal sailing can dominate total port time (Konings, 2007; Shobayo and Van Hassel, 2019). These choices interact with the planning structures and rules outlined in this chapter and with the KPI framework in Section 2.4.1. With the physical chain outlined, the next section examines how governance assigns responsibilities and performance objectives.

2.2. Port Governance and Management

The container-handling chain involves multiple actors, each with distinct roles and objectives. The port authority ensures nautical access and provides infrastructure; terminal operators manage the quay, yard, and landside interfaces; barge and other hinterland carriers move containers to and from the port; shipping agents and forwarders arrange documentation and schedules; and customs oversee regulatory compliance. These actors coordinate across strategic, tactical, and operational levels. Figure 2.2 shows typical contractual and coordination relationships between stakeholders involved in container logistics (A. Douma, 2008).

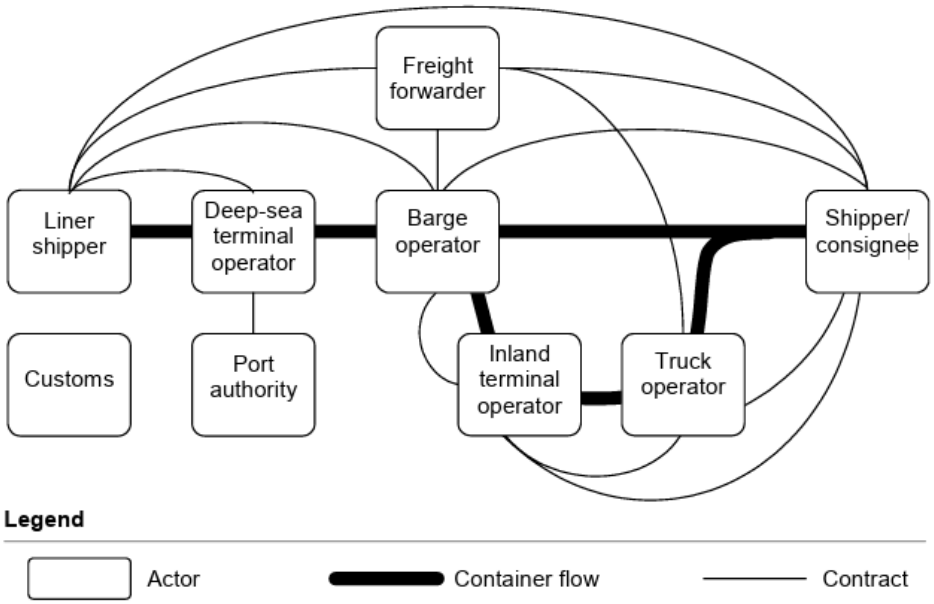

Figure 2.2: Actors involved in hinterland container transport and their contractual relationships, adapted from Douma (2008) (A. Douma, 2008).

These relationships are shaped by the port's governance model. Port governance is commonly grouped into public service, tool, landlord, and private service models. In Europe, the landlord model dominates; infrastructure is publicly owned while private operators manage terminals and most operational decisions. This division of responsibilities and decision rights determines who is accountable for which KPIs. Alignment mechanisms such as shared data platforms and service windows are therefore important precursors to improving planning. In landlord settings the port authority steers system objectives and enables digital coordination, terminal operators optimise local execution, and barge operators optimise fleet cycles across multiple terminals (World Bank, 2007; Port of Rotterdam, 2025a).

2.2.1. Decision-Making Levels in Port Logistics

The governance model is enacted through three decision-making levels, strategic, tactical, and operational (STO). These levels differ by time horizon, actors, decision rights, information required, and outputs, as summarised in Table 2.1.

Strategic - Long-term planning and policy, including infrastructure development, investment strategies, governance design, and data-sharing agreements. In landlord ports this is led by the port authority and public bodies, with input from terminal operators and carriers. Typical horizons are multi-year. Outputs include infrastructure plans, concession terms, fairway and berth capacity targets, and policies on digital platforms and sustainability.

Tactical - Medium-term planning that translates strategy into capacity and service plans. Decisions include berth and window policies for barges, crane and labour rosters, planned maintenance, target service levels, and route patterns for barge rounds. Ownership is shared; terminals set yard and quay plans, barge operators design rotations and buffers, and the port authority facilitates alignment, for example via platform rules. Horizons range from weeks to one season. Outputs include berth plans, window allocations, resource schedules, and performance targets.

Operational - Day-of-operations execution and short-horizon replanning. Decisions cover berth assignment and sequencing, crane dispatch, start and completion of moves, re-routing during disruptions, and choices to split or return calls when cargo is not ready. Actors are terminal planners, barge skippers or remote operators, Vessel Traffic Services (VTS) and nautical services. Horizons are minutes to days. Outputs include actual berth sequences, crane work orders, timestamps, and exception logs.

Table 2.1: STO levels mapped to actors, horizons, decisions, information, and example KPIs

Level	Main actors	Typical decisions	Information required	Example KPIs
Strategic	Port authority, public agencies, terminal owners	Infrastructure capacity, concession terms, platform governance, sustainability targets	Demand outlook, nautical capacity, investment options, stakeholder commitments	Modal split, total throughput, emissions inventory, digital adoption
Tactical	Terminal operations management, barge operators, platform operator	Berth and window policies, crane and labour rosters, maintenance, barge rotations and buffers	Forecast arrivals, handling times, yard capacity, service targets	Berth occupancy, crane utilisation (mean, P90), service level within window
Operational	Terminal planners, barge skippers or remote operators, VTS, pilotage and towage	Berth assignment and sequence, crane dispatch, start times, return or split calls, disruption handling	Real-time arrivals, container readiness, resource status, nautical constraints	Waiting time, service time, turnaround, on-time calls, moves per hour

With roles and horizons fixed, the next subsection sets out how planning and control are structured.

2.3. Barge Planning and Terminal Operations

A port call is the organised set of nautical and cargo activities that allow a vessel to enter the port, be serviced, and depart. The port call framework standardises the milestones and timestamps used to plan and monitor progress, providing a common reference for coordination between actors, particularly at the tactical and operational levels summarised in Table 2.1. Phases include pre-arrival, arrival, alongside and readiness checks, cargo operations, departure, and post-departure reporting. Sub-processes may

run in parallel, for example pilotage concurrent with towage. Common timestamps include estimated time of arrival (ETA), preferred time of arrival (PTA), required time of arrival (RTA), estimated/actual time of berthing (ETB/ATB), cargo start and complete, and estimated/actual time of departure (ETD/ATD). These provide a shared language for plans and updates across vessel traffic services, terminals, carriers, and platforms. A Rotterdam-specific process illustration is provided in Appendix B.

2.3.1. Coordination Structures and Rules

Efficient communication and coordination are essential to align independent actors in the port logistics system. In practice, interactions span multiple decision-making levels and are executed through specific planning structures. Three distinct planning structures are commonly recognised in port environments, as illustrated in Figure 2.3:

- Centralised planning: A single coordinating authority makes system-wide decisions on vessel and container movements across terminals, with full information and direct control.
- Decentralised planning: Individual terminals and transport operators make independent operational decisions based on local information and objectives, without overarching coordination.
- Distributed planning: Actors retain decision-making autonomy but exchange information and coordinate through structured digital platforms, improving alignment while preserving local flexibility (Truong, 2016; A. Douma, 2008).

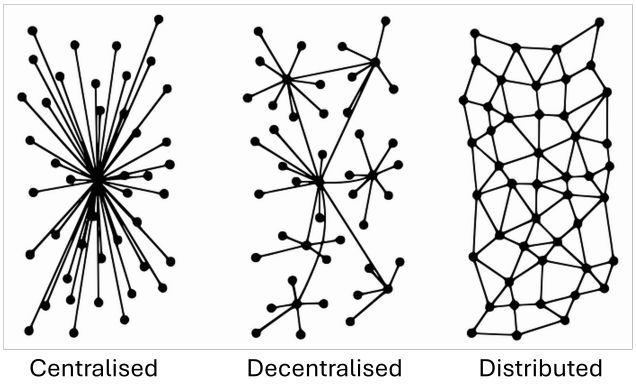

Figure 2.3: Centralised, decentralised, and distributed planning networks (adapted from Truong, 2016).

From a control perspective, decentralised planning lacks communication between agents, whereas distributed planning introduces structured information sharing, often supported by platforms such as *Nextlogic* and *Portbase*. These systems offer real-time visibility of quay availability and handling schedules so that barge operators and terminals can align plans without surrendering control to a central authority (Nextlogic, 2024; Portbase, 2025). Fully decentralised decision-making can lead to misalignments or conflicts, while distributed systems require reliable data sharing and mutual trust to function effectively.

As vessels arrive, terminals apply access and sequencing rules to balance transparency, perceived fairness, contractual commitments, and system throughput. Liner deep-sea services typically operate with fixed windows and often priority arrangements, feeders run shorter cycles, and barges and inter-terminal moves rely more on day-of-operations alignment and container readiness. In its simplest form, first come, first served admits vessels strictly by arrival time, which is transparent but can be inefficient under congestion. Priority-based policies order calls by class or contract, for example deep-sea ahead of barge, minimum call size, or cargo type, thereby protecting service levels. Shortest processing time favours the quickest calls to raise berth turnover, whereas earliest due date protects time-critical chains; a weighted SPT variant adds penalties or values to reflect differentiated service. Where digital coordination is mature, integrated or distributed optimisation is used, with terminals or platforms solving a scheduling problem, often via mixed-integer models or heuristics, to minimise waiting, tardiness, or other system KPIs. Table 2.2 summarises these rules and typical use cases (Imai, Nishimura, and Papadimitriou, 2001; Steenken, Voß, and Stahlbock, 2004; Bierwirth and Meisel, 2010).

Table 2.2: Overview of common vessel and barge planning rules

Rule	Description	Typical use case
FCFS	Serve in strict arrival order; simple but can increase delay under congestion.	Default in decentralised or low-automation settings.
Priority-based	Order by class, e.g., deep-sea over barge, size, contract, cargo type.	High-throughput terminal, service-level commitment.
SPT	Shortest handling time first to maximise berth turnover.	Congested periods, delay minimisation.
EDD	Earliest cargo due date first to protect time-critical chains.	Tight delivery windows, synchronisation needs.
WSPT	Processing time weighted by value or penalty.	Mixed-priority customers, differentiated service.
Integrated / distributed optimisation	MILP, heuristics, or platform-aided scheduling to minimise system KPIs.	Digitally mature ports with TOS or platform links.

The choice of coordination structure and access/sequencing rules conditions where delays arise; the next subsection summarises observed bottlenecks in large multi-terminal hubs.

2.3.2. Bottlenecks in Barge Handling at Large Multi-Terminal Hubs

Terminals coordinate berth allocation, crane assignment, and yard flows under objectives for moves per hour, resource use, and service reliability for deep-sea, feeder, and barge calls (Steenken, Voß, and Stahlbock, 2004). For barges, asynchronous container readiness due to yard priorities, cut-off times, and hinterland connections creates additional coordination needs, which motivates explicit planning rules and timely data exchange (A. Douma, 2008).

In this context, efficient barge handling in large multi-terminal hubs is constrained by the interaction of multi-stop rotations, heterogeneous terminal rules, and partial information sharing. Evidence from Rotterdam and comparable ports shows that these frictions manifest as long and variable waits, missed windows for small calls spread across terminals, and sensitivity to quay-space fragmentation, effects that are operational rather than purely technological. The following observations summarise the bottlenecks most relevant in practice.

First, long and variable waiting times persist at deep-sea terminals. The Central Commission for the Navigation of the Rhine (2024) reports average barge waits of 20 to 30 hours in Rotterdam (2019), with peaks up to 60 hours during congestion episodes. Recent platform-enabled planning via Nextlogic reports around 20% shorter port stay for participating vessels, underlining that access, sequencing, and information sharing are binding constraints beyond pure handling speed.

Second, many calls with small call sizes are spread across terminals (Van der Horst et al., 2019). This pattern increases exposure to missed windows and lengthens total port stay; institutional analyses for Rotterdam link weak alignment between barge rotations and terminal schedules to delay and unreliability (A. M. Douma, Schuur, and Schutten, 2011).

Third, minimum-move expectations and fixed windows at major terminals favour fewer, larger calls. Public procedures at ECT state minimum call sizes, for example 175 moves at Delta DDE and 150 at Euromax with limited deviation, which improves crane productivity but raises barriers for small or late flows.

Finally, quay-space fragmentation arises from irregular arrivals and vessel lengths. Berth allocation studies formalise quay planning as a rectangle packing problem (continuous BAP), which implies that flexibility in effective vessel length helps occupy residual pockets and raise utilisation (Dai et al., 2008; Frojan et al., 2015).

The coordination choices and bottlenecks identified above determine where time, capacity, and reliability are lost or preserved. The next section sets out the performance lenses and indicators that quantify these effects at call, terminal, and port levels.

2.4. Port Performance Measurement

Port performance describes how effectively a port system converts resources and capacity into outcomes that users value. A widely used framing distinguishes efficiency, productivity, effectiveness, and resilience (Pallis and Rodrigue, 2024). Efficiency concerns the ratio of inputs to outputs, productivity the level of output achieved, effectiveness the degree to which service promises are met, and resilience the ability to maintain or quickly recover performance under disruption. These dimensions interact; coordination that reduces waiting lifts productivity and effectiveness and may indirectly improve environmental efficiency through smoother flows.

Port performance is also viewed through managerial lenses that reflect stakeholder perspectives: economic (cost/TEU, productivity), investment (asset utilisation, ROI), technological (digital integration, platform adoption), and ecological (energy/emissions per TEU) dimensions (Pallis and Rodrigue, 2024). Emission-related accounting follows the GHG Protocol Corporate Standard with Scope 1 (direct), Scope 2 (purchased energy), and Scope 3 (value chain) boundaries, enabling consistent attribution across actors (GHG Protocol Team, 2011). In a landlord setting, operators' reporting focuses on Scopes 1-2, for example barge propulsion and terminal electricity, whereas Scope 3 is often addressed in authority-led programmes such as modal shift or shore-power policy. Figure 2.4 summarise the scope definitions and their distribution across the value chain. In port environments, these lenses connect most directly to Sustainable Development Goal 9 (industry, innovation and infrastructure), via digital coordination and asset utilisation, and to Goal 13 (climate action), via energy and emissions per TEU (UN Office for Sustainable Development, 2015). Broader SDG linkages, including Goal 14 (life below water), are acknowledged but fall outside the measurement scope used here.

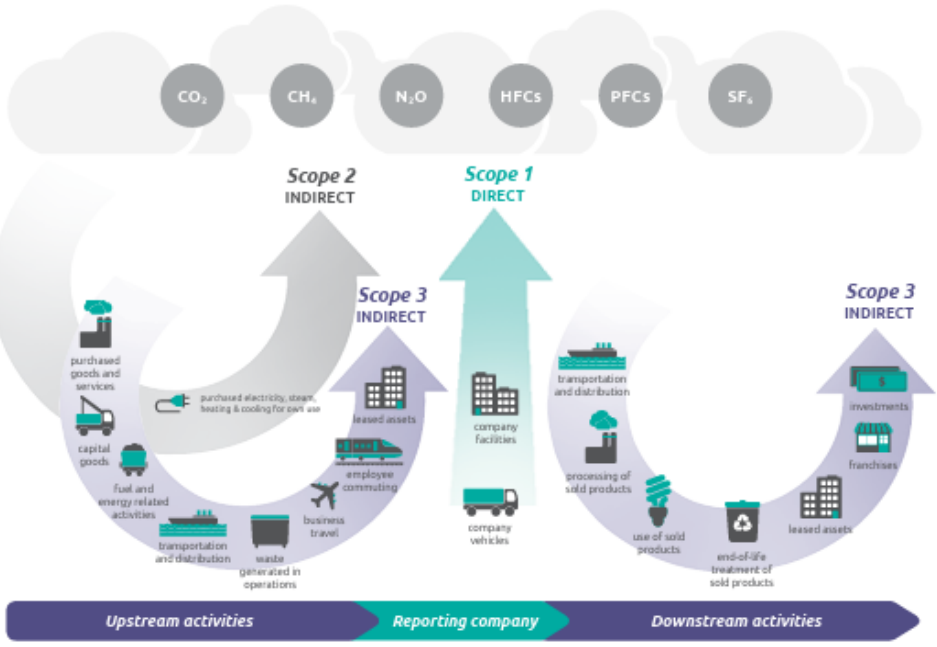

Figure 2.4: Overview of GHG protocol scopes and emissions across the value chain (GHG Protocol Team, 2011)

Focussing on efficiency at the operational level, different types of efficiency affect the performance in container ports (Pallis and Rodrigue, 2024):

- Operational/technical efficiency concerns how well resources are used in day-to-day operations, for example crane moves per hour and berth or yard utilisation.
- Financial/allocative efficiency concerns the choice and use of resource mixes, for example cost per TEU and revenue per quay-metre.
- Systemic/coordination efficiency reflects how well interactions between actors are aligned, for example predictability, adherence to agreed windows, and split or return rates.

- Environmental efficiency describes the output achieved per unit of energy use or emissions, for example energy or emissions per TEU.

These constructs complement the conceptual dimensions and managerial lenses. Operational gains that reduce waiting can lift productivity; better coordination improves effectiveness (on-time service) and, indirectly, environmental efficiency through smoother flows. The next section defines the attribution framework by stakeholder and decision level and introduces the KPI families used in the results.

2.4.1. Operational Definitions

Operational performance is read and reported at three levels: per call, per terminal or route over a period, and for the whole port (Steenken, Voß, and Stahlbock, 2004; Pallis and Rodrigue, 2024). Typical horizons are a shift or day for operations, a week or month for planning, and a quarter or year for strategy. The port-call timeline (Section 2.3) defines the key events and timestamps: arrival (t_{arrival}), all fast/berth start ($t_{\text{berth start}}$), cargo start/complete, and departure ($t_{\text{departure}}$). Predictability compares actual times to a reference; either the submitted plan or the agreed requirement. Each report should state which reference is used and which summary number is shown (for example, median P50 or 90th percentile P90)(Nextlogic, 2024; Portbase, 2025).

Capacity-based KPIs use available time rather than raw calendar time so that denominators reflect capacity that could actually be used. Berth-hours exclude closures and maintenance, and crane-hours exclude breakdowns and planned outages. This accounting choice is standard in terminal studies and reviews to avoid diluting utilisation and productivity with off-line capacity (Steenken, Voß, and Stahlbock, 2004).

Across these objects and horizons, four KPI families recur in reviews and practice. Utilisation captures how intensively capacity is used, for example berth and yard occupancy. Productivity reflects output per constrained resource, for example gross berth productivity and crane moves per hour. Service and predictability assess whether operations meet plans, for example waiting time, turnaround time, window adherence, and plan adherence. Environmental impact reports energy or emissions per TEU, typically attributed under the GHG Protocol's Scope 1–3 boundaries to separate direct fuel use, purchased electricity, and value-chain emissions (Pallis and Rodrigue, 2024; GHG Protocol Team, 2011). A fifth theme, reliability under stochastic conditions (e.g., variability of arrivals or handling times), is frequently discussed in reviews but is only assessed when uncertainty is explicitly modelled (Dragovic, Tzannatos, and Park, 2017).

Stakeholders emphasise different KPI families at different decision levels in landlord ports. Table 2.3 summarises the typical emphases.

2.4.2. KPIs Alongside Coordination and STO

Changes in coordination structures and access or sequencing rules (Section 2.3.1) register in indicators on different timescales and at different decision levels (Section 2.2.1). At the operational level, moving from decentralised to distributed planning typically improves predictability (smaller P90 start and finish deviations) and reduces idle time (shorter waiting, fewer idle crane-hours) over shift or day horizons. At the *tactical* level, once windows, buffers, and rotations are updated, changes appear in window adherence, split and return rates, and berth-plan adherence. At the strategic level, throughput, modal split, and emissions per TEU adjust once operational and tactical changes stabilise, typically over month or quarter horizons.

Attribution should follow the level that controls the relevant decision right. For example, priority classes and minimum call sizes are tactical policies implemented through operational sequencing rules (Section 2.3.1). If the split and return rate falls, the improvement originates tactically even if observed in operational statistics. When interpreting results, it is useful to distinguish direct indicators driven by vessel–terminal interaction (turnaround, waiting time, moves per hour) from indirect indicators that shift through second-order effects as arrivals smooth and conflicts fall (berth occupancy, yard dwell, modal split). Accordingly, the evaluation reports three groups: direct operational indicators for immediate effects, tactical indicators for coordination choices, and system-level indicators for policy relevance. Table 2.3 summarises how responsibilities and example KPIs are organised by stakeholder and decision level.

Table 2.3: KPIs by stakeholder and decision level

Stakeholder	Level	Purpose and example KPIs, units
Port Authority	Strategic	Modal split [%], total port throughput [TEU], predictability of port calls (P90 deviation) [min], Scope 1–3 inventory, digital platform adoption [%].
	Tactical	Average berth occupancy across terminals [%], ship/barge delay rate [% calls delayed], planned vs. actual nautical service times [min].
	Operational	Simultaneous movements [#], conflict events resolved [#], VTS response times [min].
Terminal Operators	Strategic	Annual TEU per hectare, energy per TEU, service level for barge calls [% within window].
	Tactical	Mean and P90 crane utilisation [%], berth plan adherence [min], yard occupancy [%], container dwell time [h].
	Operational	Crane moves per hour [moves/h], barge turnaround time [h], start delay after all fast [min].
Barge Operators	Strategic	Trips per barge per week [#], on-time performance [%], energy per TEU-km.
	Tactical	Calls per round trip [#], planned buffer per leg [min], split/return rate [% of calls].
	Operational	Waiting time [h], berth time [h], waypoint deviation [m].

Taken together, these responsibilities determine how coordination choices propagate through operational, tactical, and strategic indicators.

2.5. Autonomous Barges and Automated Transhipment

Autonomy in the maritime sector is framed by a small set of widely used taxonomies. The International Maritime Organization (IMO) distinguishes four degrees of autonomy for Maritime Autonomous Surface Ships (MASS); ships with automated processes and decision support, remotely controlled ships with crew on board, remotely controlled ships without crew, and fully autonomous ships that make and execute decisions on their own (IMO, 2018; IMO, 2021).

Inland and port applications are motivated by a combination of reliability, safety, labour availability and decarbonisation. Autonomy can raise schedule adherence and allow finer service frequencies, reduce exposure to crew shortages, and pair naturally with electrification and digital coordination in terminals (Schiaretti, Chen, and Negenborn, 2017; Negenborn et al., 2023; Puleri, 2023). Most current prototypes operate with remote supervision and partial autonomy across subsystems (decision-making, action execution, exception handling, cooperation) (Schiaretti, Chen, and Negenborn, 2017; Negenborn et al., 2023).

Classification societies provide complementary engineering views. Lloyd's Register's procedures describe how autonomy is accepted and assured in class, from decision support through remote operation to higher levels of automation, specifying system assessments, integration, and validation activities that accompany each step (Register, 2017; Register, 2017). The commonly reproduced Lloyd's "levels" are best read as an assurance scaffold rather than a single ordinal ladder of capability (Figure 2.5).

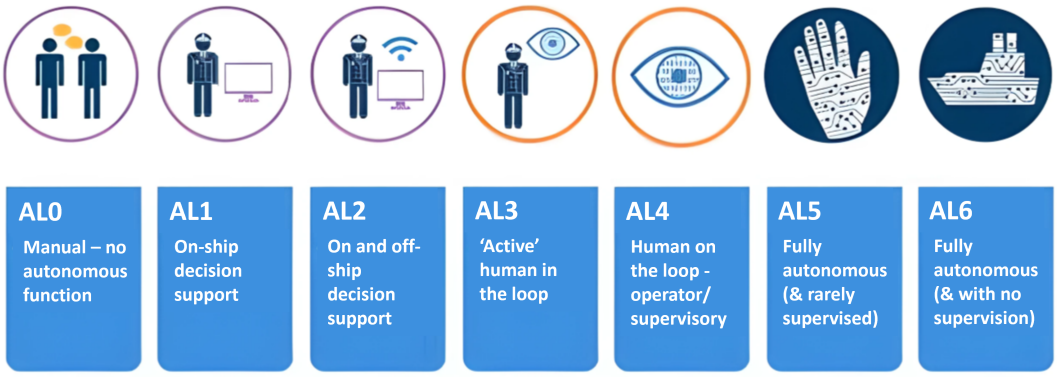

Figure 2.5: Levels of autonomy according to Lloyd's Register(Vidan et al., 2019)

To operationalise these notions for vessels calling at multiple terminals per rotation, a subsystem view is analytically useful. Schiaretti, Chen, and Negenborn 2017 decompose autonomy into decision, action, exception, and cooperation, arguing that an overall level emerges from the maturity and coupling of these components rather than from a single headline label. For port-facing operations, the Decision component covers arrival timing and sequencing against service windows, the Action component covers low-speed approach, berthing and mooring, the Exception component covers recovery when cargo or quay is not ready, and the Cooperation component covers data exchange with terminals and platforms. Recent refinements recommend distinguishing automation capability from where operators reside, onboard versus shore control centres, clarifying how IMO degrees map to concrete control architectures (Ø. Rødseth, L. A. Wennersberg, and Nordahl, 2022; Ø. J. Rødseth and L. A. L. Wennersberg, 2023). For inland shipping specifically, structured descriptions link autonomous barge functions with supporting shore procedures and infrastructure, which is directly relevant in landlord settings with distributed decision rights (Hagaseth, Meland, and Wille, 2023).

Technically, port-facing autonomy depends on three coupled capabilities. First, perception and localisation fuse AIS, VTS and cooperative messages with cameras and LiDAR to maintain safe clearances in dynamic, degraded conditions. Second, low-speed guidance and control support approach, berthing and departure. Third, resilient ship-shore communications support remote supervision and data exchange (Prasad et al., 2017; Ø. J. Rødseth, L. A. L. Wennersberg, and Nordahl, 2023; Neurohr et al., 2025). Competence frameworks and recommended practices for shore control centres have recently emerged from classification societies, reflecting the shift of certain bridge functions to shore (Neurohr et al., 2025).

Regulatory integration rests on two pillars. First, MASS trials follow IMO interim guidelines that set expectations for risk management, human–machine interfaces, and engagement with coastal, flag, and port states (Maritime Safety Committee, 2019). Second, collision-avoidance obligations continue to derive from the COLREGs; autonomy affects how compliance is achieved, not whether it is required, which anchors exception handling and right-of-way logic for intra-port movements irrespective of automation level (International Maritime Organization, 1972).

Benefits, reduced crew dependence, more predictable operations, and potential energy savings, materialise only when autonomy is coupled to operational integration; digital berth planning and TOS connectivity, shore control centres with reliable communications, and charging where fleets are electrified (Puleri, 2023). In Rotterdam, platform participation via Nextlogic and Portbase provides information discipline for appointment adherence. Early operational evidence reports roughly a 20% reduction in port stay for participating barges, indicating that access, sequencing and data sharing are binding constraints even before fully unmanned operation (Nextlogic, 2024; Portbase, 2025).

2.6. Modelling Approaches in the Literature

Research on port and terminal planning typically employs three methodological families. Discrete-event simulation (DES) captures event-driven flows with resource constraints and queues and is widely applied across berth, yard and inter-terminal contexts. Surveys and handbooks in container-terminal operations position DES as the dominant tool for analysing operational dynamics and policy levers (Steenken, Voß, and Stahlbock, 2004). Agent-based modelling (ABM) is used when heterogeneous actors and adaptive or strategic behaviour matter (e.g., platform participation, learning, distributed decision-making), with barge-handling studies in Rotterdam using serious-gaming/ABM to introduce multi-agent coordination concepts and assess organisational feasibility rather than system KPIs. Optimisation methods, classical MILP/heuristics for berth, quay, crane and integrated scheduling, generate static or rolling plans that are often embedded into hybrid DES+optimisation frameworks. Representative contributions include berth-allocation models and solution methods as well as dynamic berth-allocation formulations that inform DES scenario design (Imai, Nishimura, and Papadimitriou, 2001). Field-level reviews note that existing studies rarely cover system-level intra-terminal coordination, split/return calls and stakeholder-owned KPIs in landlord ports, motivating a port-wide simulation perspective that treats such mechanisms as first-class objects.

Port-Wide and Barge-Related Simulation Studies: Scope and Limits

Early port-wide simulation for Rotterdam formalised a generic multiterminal model by composing transport, transfer and stacking functions, and applied it to the complete set of existing and planned terminals on the Maasvlakte. The experiments report requirements for deep-sea quay length, storage capacity and inter-terminal transport capability, but appointment reliability and barge-specific dispatch rules are not isolated as experimental factors (Ottjes et al., 2007).

A complementary line compares inter-terminal transport systems at the Maasvlakte, multi-trailer systems versus automated guided/lifting vehicles, showing how networked flows create joint bottlenecks between vehicles and handling cranes and require complex planning and control. The analysis remains equipment-centric and land-sided, rather than evaluating water-borne barge coordination and window adherence (Duinkerken et al., 2007).

For barge planning and coordination itself, work based on a simulation game introduces and evaluates a multi-agent, distributed-planning concept for Rotterdam with an emphasis on organisational acceptance and change management. This stream addresses the viability of coordination but does not quantify system-wide delay impacts of specific appointment or dispatch rules under realistic stochasticity (A. Douma, 2008).

Beyond Rotterdam, traffic-generation methodology for port simulators provides the calibration discipline needed to construct arrival processes and functional relationships (e.g., between ship dimensions, capacity and crane deployment). These methods are formulated for sea-going traffic and have not been adapted to multi-call intra-port barge rotations and their cascading disruptions (Pachakis and Kiremidjian, 2003).

Field-wide reviews corroborate the picture: most simulation studies concentrate on terminal-level operational problems, while integrated, system-level analyses are comparatively scarce, motivating a port-wide perspective that recognises inland-waterway operations as non-peripheral (Dragovic, Tzannatos, and Park, 2017).

These modelling strands motivate a barge-centred, port-system simulation that adapts traffic-generation discipline to multi-call barge rotations and treats appointment and dispatch rules as policy levers. This setup interacts directly with the autonomy and automated-transhipment concepts reviewed above and is evaluated through the KPI framework in Section 2.4.

2.7. Synthesis and Research Gaps

The literature reviewed in Sections 2.1-2.6 shows that barge-based transhipment in large, multi-terminal ports is shaped by landlord governance, distributed decision-making, heterogeneous local planning rules, and inter-terminal coupling of resources. Studies on barge planning and coordination document long and uncertain waiting times, fragmented berth planning, and emerging platform-based initiatives that aim to improve reliability (Section 2.3). Work on port performance measurement provides indicator frameworks that cover turnaround times, berth occupancy, crane productivity, and emissions at different decision levels (Section 2.4). Research on autonomous and modular barge concepts concentrates on technical feasibility, navigation, and terminal-level interactions rather than system-wide impacts (Section 2.5). Simulation studies of terminals and port areas offer robust methods for traffic generation and time-dependent operations (Section 2.6), yet they largely treat inland barges in an aggregate fashion and seldom focus on intra-port rotations.

Relative to the research problem in Section 1.3 and the main question in Section 1.5 (how autonomy-enabled coordination strategies for inland barges affect the operational efficiency of intra-port container transhipment), this body of work leaves four concrete gaps.

- System-level evidence on barge coordination
  Existing studies on barge coordination and serious-game based planning concepts for Rotterdam primarily assess organisational feasibility and qualitative effects of distributed planning, rather than quantitatively comparing appointment and dispatch policies under realistic berth, crane, and schedule constraints (Section 2.3). As a result, there is limited quantitative evidence on how specific coordination rules shape waiting-time distributions, service times, window adherence, and berth occupancy across multiple terminals. This thesis addresses this gap by embedding alternative appointment

and dispatch policies in a port-wide discrete-event simulation and evaluating their impact on these indicators relative to current practice.

- Rotation-aware traffic generation
  Port-wide traffic-generation methods developed for sea-going vessels provide calibrated arrival processes and functional links between vessel characteristics and handling resources (Section 2.6). However, they generally treat inland barges as single, independent calls and do not capture multi-call intra-port rotations in which delay at one terminal propagates to subsequent calls and triggers re-sequencing decisions. The simulation model developed in this thesis adapts traffic-generation discipline to represent barge rotations explicitly, including call order, slack, and disruption propagation, so that coordination strategies can be tested under realistic rotation effects.

- Comparative evaluation of planning rules
  Studies of port and terminal operations employ a variety of berth-access and sequencing rules, from first-come-first-served to priority classes and due-date based policies, yet these are rarely compared side-by-side using a common case, a shared set of performance indicators, and a landlord-port setting (Sections 2.3.1 and 2.4). This thesis takes a first step toward that comparison by embedding a realistic class-priority access rule with FCFS tie-breaking in a port-wide landlord model and quantifying how its performance, in terms of turnaround times, berth occupancy, crane utilisation, and anchorage behaviour, changes under different coordination concepts (modular splitting) and demand and resource regimes. A full side-by-side comparison of FCFS, due-date based, and platform-optimised scheduling policies remains an avenue for future work.

- Impact of autonomous and modular barge concepts at port scale
  Research on autonomous and modular barges focuses mainly on technology, navigation, and conceptual designs or on terminal-level process changes (Section 2.5). System-level implications of such concepts for berth occupancy, quay productivity, and service reliability, under realistic deep-sea traffic and resource constraints, remain largely unquantified. This thesis fills this gap by representing modular e-barge strategies, such as splitting and interchange, explicitly in a port-wide discrete-event simulation of the Maasvlakte area and comparing their performance with current multi-stop barge calls under common demand scenarios and key performance indicators.

Together, these gaps link the literature review to the research problem and frame the conceptual scenarios (Chapter 3), model design (Chapter 5), and experimental analysis (Chapter 7) that form the core of this thesis.

3

Conceptualisation and scenario selection

This chapter introduces the coordination concepts considered in the study and defines the scope for scenario comparison. It documents the design space explored during ideation and links the concepts to operational problems observed in large hub ports. This is the first chapter where the concepts are defined at a high level. The detailed model rules and parameters follow in Chapter 5.

3.1. Autonomous E-Barge and Coordination Strategies

The autonomous e-barge is conceived as a small, electrically propelled barge with automation features that ease crew-related coordination. The emphasis is on operational autonomy rather than navigational autonomy, enabling more flexible dispatch, shorter workable windows, and operation in modular configurations. The unit is intended to shorten berth visits, make use of short planning windows, and provide planners with a flexible asset that can be deployed more frequently. Electrification is assumed to lower local emissions, while the analysis focuses on the coordination levers that this configuration enables. Two application patterns guide the design space:

- Intra-port shuttle: Short port-internal moves that collect or deliver containers between deep-sea terminals, including leftover and late-registered units.
- Scalable call size: Service that adapts effective vessel length and call size to fit residual quay pockets and available crane time.

Ports differ in access rules and window regimes. The concept therefore aims to work with, not against, prevailing practices such as minimum-move expectations, dedicated barge spots, and container registration cut-offs.

The application patterns above translate into two primary coordination concepts, plus one interactive variant. These will be compared to the baseline, which represents present-day multi-stop barge operations in Rotterdam and serves as the reference for comparison. Figures 3.1- 3.3 provide concise sketches of the concepts introduced below.

1. Roundtrip service: A fixed loop visits a predefined set of terminals on a repeating circuit. The rotation is predictable, which can help synchronise with terminal windows and reduce ad hoc repositioning.
2. Modular splitting: A mother arrives with detachable modules that berth in parallel at different terminals, then recombine for departure.
3. Modular splitting with interchange (variant): Completed modules may interchange between mother ships to reduce rendezvous waiting, subject to additional coordination and governance.

Roundtrip Service - A roundtrip strategy uses a fixed cyclic loop that visits a predefined set of terminals in sequence (for example A $\rightarrow$ B $\rightarrow$ C $\rightarrow$ A). This raises predictability at each terminal and, when window regimes are stable, can reduce arrival variance and the need for ad hoc repositioning. The approach presupposes timetable discipline and platform support that may not always hold. Thin flows can lead to sub-threshold calls, delay at one stop can cascade to later stops in the loop, and a single larger hull cannot exploit residual quay pockets efficiently. The concept is visualised in Figure 3.1.

Figure 3.1: Coordination concept: Roundtrip service with multiple routes serving terminals and depots in the Port of Rotterdam (Flikkema et al., 2022).

Modular Splitting - Modular splitting deploys a mother with detachable modules that berth in parallel across terminals where work and crane time are available, then rendezvous to depart together. Parallel handling can compress total port time where cranes are available. Variable module length improves fit to residual quay pockets, and concentrating a terminal's work in its own module can support minimum-move expectations without proliferating many small sequential calls. The main exposure is the departure rendezvous, where a late-finishing module can delay the convoy. Additional module sailing introduces overhead, and clear rendezvous rules, buffers, and traffic coordination are required to protect punctuality and safety. The concept is visualised in Figure 3.2.

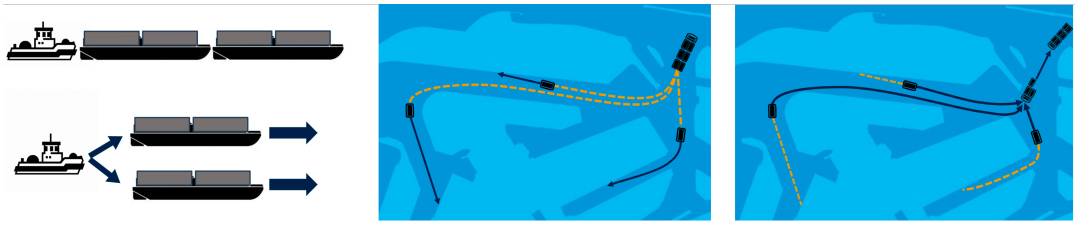

Figure 3.2: Coordination concept: Modular splitting

Interchange Variant - An interchange variant (Figure 3.3) extends modular operations by allowing completed modules to transfer between mothers or be staged for later pick-up, decoupling module completion from convoy schedules. This can reduce idle time at the rendezvous, lower the number of small craft concurrently under way by aggregating modules into larger composite movements, and increase the likelihood that brief windows translate into completed work. These benefits come with coordination and governance requirements, including permissions, traceability, liability, and alignment with port traffic services. Additional shunting or staging could offset some gains if poorly designed.

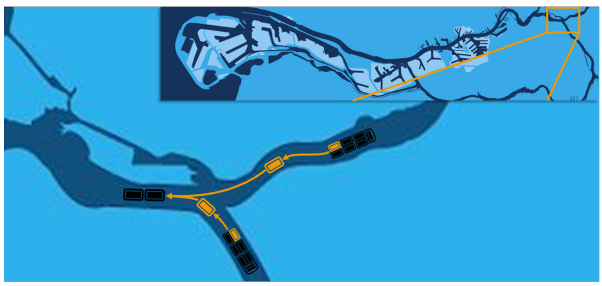

Figure 3.3: Coordination concept: Modular splitting with interchange variant.

3.2. Targeted Bottlenecks and Mechanisms

The concepts respond to recurrent operational bottlenecks in large hub ports: (i) small and time-sensitive calls, (ii) short berth opportunities, (iii) minimum-move expectations that favour fewer, larger calls, and (iv) quay-space fragmentation that leaves residual pockets difficult to use. A fifth concern arises within modular operations, namely (v) rendezvous delay when modules complete at different times. The literature points to operational design features that are likely to alleviate these bottlenecks:

1. Concentrate work per terminal to meet minimum-move expectations with fewer under-sized calls, improving eligibility for fixed windows.
2. Convert short berth opportunities into completed work by keeping visits short and predictable and by using integrated cross-actor planning (e.g., Nextlogic) to raise the value of responsiveness.
3. Use flexible effective vessel length so calls fit residual quay pockets and berth utilisation rises, consistent with continuous berth allocation and rectangle-packing insights (Dai et al., 2008; Frojan et al., 2015).
4. Reduce sequential in-port sailing by handling work in parallel where feasible and by limiting repositioning that cascades delays over multi-stop rounds.

These options provide a neutral foundation for the concepts introduced in Section 3.1 and for the scenario selection that follows. Table 3.1 summarises how the mechanisms implied by the concepts are expected to affect the headline indicators used later in the study.

Table 3.1: Mechanisms and expected KPI directions relative to the baseline

Mechanism	Expected effect on KPIs
Regular loop with consolidated stops (Roundtrip)	Predictability ↑, repositioning slack ↓; waiting remains sensitive to missed windows; no direct effect on quay-pocket fit
Parallel berthing with flexible effective length (Modular splitting)	Waiting at anchorage ↓, turnaround time ↓, berth-occupancy utilisation ↑, crane productivity ↑
Synchronisation at recombination (Modular splitting)	Departure-delay risk ↑ if one module lags; service reliability ↓ driven by the slowest module
Concentration of moves per terminal (Modular splitting)	Alignment with minimum-move expectations ↑; share of sub-threshold calls ↓
Interchange between mothers (Modular variant)	Idle time at rendezvous ↓, departure punctuality ↑; additional coordination and shunting requirements

3.3. Rationale for Focusing on Modular Splitting

Modular splitting embeds levers that directly target the observed constraints in Rotterdam-like hubs. Concentrating work per terminal aligns with minimum-move expectations and fixed-window regimes. Flexible effective length allows modules to occupy residual quay pockets that a single larger hull cannot use efficiently. Parallel handling converts brief windows into completed work when crane time is available. The concept also matches the scope of this study, since it introduces day-of-operations flexibility without requiring network-wide timetable design.

3.4. Comparison Scope and Future Work

The scenario experiments compare the baseline and modular splitting under shared environmental conditions. Roundtrip and the interchange variant are retained for conceptual completeness but are not modelled. For these two concepts, detailed process-flow BPMN diagrams are provided in Appendix C. Two extensions are identified for future work: (i) recombination policies and buffers, including their effect on punctuality and reliability, and (ii) richer module interactions such as interchange between convoys or staged pick-up, which require additional coordination protocols, permissions, and traceability as outlined in Appendix C. Chapter 5 presents the modelling framework, parameter settings, and scenario design used to evaluate the selected concept.

4

Case Study – Port of Rotterdam

The Port of Rotterdam is the largest seaport in Europe and one of the most significant logistics hubs worldwide. In 2024, it handled approximately 13.8 million TEU, representing a critical share of Europe's container throughput (Port of Rotterdam, n.d.(b)). Its strategic location, deep-water access, and extensive hinterland connections by road, rail, and inland waterways make it a central node in global supply chains. Within this extensive transport network, inland waterway transport plays a particularly important role, accounting for around 34% of the hinterland container volumes (Port of Rotterdam, n.d.(b)).

As illustrated in Figure 4.1, terminals and depots are spatially dispersed across the Maasvlakte and Waalhaven areas. This geography reinforces the need for efficient short-distance movements and careful coordination across multiple, independently operated facilities. The Port of Rotterdam therefore serves as the empirical setting for this study, providing a large, multi-terminal hub where inland barge coordination is operationally significant. The coordination concepts examined, including autonomous and modular barges, and the associated performance indicators are formulated such that the findings are informative for other hub ports with similar structures and hinterland roles.

Figure 4.1: Terminals & depots in the Port of Rotterdam (Port of Rotterdam, n.d.(a))

4.1. Container Transport in Rotterdam

Container flows in Rotterdam can be divided into four categories: deep-sea, feeder/shortsea, inland, and inter-terminal transport. Deep-sea vessels connect the port to global trade networks, with multiple ultra-large container vessels (ULCVs) calling at the Maasvlakte terminals each week. Feeder and short-sea vessels redistribute containers regionally. Inland barges serve the hinterland, connecting Rotterdam with destinations in the Netherlands, Belgium, Germany and beyond. They increasingly

support short-haul intra-port exchanges when volumes justify a berth call. Together, these segments shape berth utilisation, crane demand, and quay access. While deep-sea calls dominate quay occupation, feeders and barges account for a significant share of total vessel movements and are material to day-to-day capacity management at the Maasvlakte terminals.

4.2. Inter-Terminal Transport and Intra-Port Exchanges

Inter-terminal transport (ITT) moves containers between terminals and depots within the port area. These flows arise when deep-sea discharge and hinterland pick-up occur at different facilities, when containers require inspection, or when empty equipment must be repositioned. In a multi-terminal hub such as Rotterdam, the volumes are significant and are handled by several modes. The spatial distribution in Figure 4.1 explains both the frequency of these moves and the need for careful coordination.

Truck-based ITT remains important for small or time-critical exchanges because it offers speed and flexibility. On the Maasvlakte, the Container Exchange Route (CER) is a dedicated, closed corridor of about 17 km that links RWG, ECT Delta, QTerminals Kramer locations (KDD, RCT, DCS), and the State Inspection Terminal for Customs. The CER keeps exchanges off public roads to reduce delays and emissions, and to improve safety and integrity. Permitted carriers include Multi Trailer Systems, terminal tractor units with chassis, and conventional trucks (Port of Rotterdam Authority, 2023; Port of Rotterdam Authority, 2025a).

Barge-based exchanges are effective for consolidated, water-connected flows and they reduce road pressure with favourable emissions per TEU. Their performance depends on berth access, crane time, and aligned handling windows at both ends. Coordination in Rotterdam has strengthened for inland services through integrated barge planning, while some intra-port barge exchanges between seagoing terminals and depots are still planned bilaterally.

Rail complements truck and barge through PortShuttle Rotterdam, a neutral intra-port service that organises daily shuttles between RSC Rotterdam (Waalhaven) and CCT Pernis and the Maasvlakte deep-sea terminals, with exchange between Maasvlakte terminals also possible (Port of Rotterdam Authority, 2025c). Capacity on the rail side is supported by the Maasvlakte–Zuid (EMZ) yard, which provides up to four bundles of six tracks for 740 m trains (Port of Rotterdam Authority, 2025b).

Empty containers are stored at dedicated depots and are repositioned according to demand. Short-notice empty moves typically use truck, often via the CER, while barge or rail are preferred when volumes accumulate or when exchanges align with scheduled rotations. The placement of depots relative to terminals, shown in Figure 4.1, therefore has a direct effect on routing and resource use.

Mode choice reflects volume, time-sensitivity, distance, and access to berths or tracks. Trucks, including the CER, minimise lead time and avoid quay conflicts for small or heterogeneous flows. Barges are suitable for pooled, water-connected volumes when windows are aligned. Rail supports block moves and predictable consolidation. These characteristics shape the key performance indicators used later in this thesis, including waiting and turnaround times, berth and track occupancy, and schedule predictability, and they motivate the scenario design in Chapter 5.

The reliance of ITT on compatible windows and resources across independent terminals connects directly to the digital coordination layer and automation discussed next. This provides the context for assessing autonomous and modular barge concepts and their interaction with terminal operations in Chapter 5.

4.3. Automation and Digital Coordination

Rotterdam's leading terminals apply high levels of automation, including automated guided vehicles and automated stacking cranes, supported by advanced planning and data exchange. Port-wide platforms enable information sharing among barge operators, terminals, and authorities, which improves window alignment and reduces idle time at the quay. This digital layer provides the foundation to integrate new vessel concepts such as autonomous or modular barges that can shuttle container modules between terminals and depots with minimal dwell and predictable service profiles. In such configurations, short and frequent moves can be coordinated with automated equipment to create a low-emission, high-reliability intra-port link.

4.4. Data and Case Relevance for the Model

Data from the Rotterdam case provide the main inputs to the simulation model in Chapter 5. Public terminal documentation from the Port of Rotterdam Authority is used to parameterise the physical layout, including quay lengths per terminal, approximate terminal area and the number and type of quay cranes (Port of Rotterdam et al., n.d.). Terminal locations and inter-terminal distances are derived from the same dataset and port charts. Numerical values for the main inputs are summarised in Appendix D.

Operational parameters are based on a combination of port statistics and practitioner input. Average call sizes, barge shares in hinterland volumes and typical inter-arrival patterns for deep-sea, feeder and inland services are derived from the Barge Performance Monitor and the Port Performance dashboards of the Port of Rotterdam (Port of Rotterdam, 2025a; Port of Rotterdam, 2025b). Practitioners provided estimates for crane productivity in moves per hour, the typical number of cranes assigned to each vessel type and minimum move expectations per barge call. These values are used as exogenous model inputs rather than simulation outputs; where detailed data were not available, stylised assumptions were adopted within the ranges indicated by experts and port statistics.

The model captures four vessel categories, namely deep-sea, feeder, inland barge and autonomous short-distance barge, and represents their interaction with quay availability and crane capacity. Barge planning practices are reflected through arrival processes and handling windows rather than explicit modelling of truck or rail systems. Dedicated road or rail infrastructure is therefore treated as contextual background and is not simulated. Performance outputs include waiting time at berth, vessel turnaround time, quay occupancy and crane utilisation, which are used to compare coordination concepts in Chapter 5.

Rotterdam's landlord governance structure and decentralised terminal ownership mirror the coordination constraints present in other large multi-terminal hubs on the North Sea range, where intra-port exchanges and barge congestion have been documented (Shobayo and Van Hassel, 2019; Central Commission for the Navigation of the Rhine, 2024).

4.5. Conclusion

Rotterdam exemplifies a multi-terminal, digitally coordinated port where deep-sea, feeder and inland barge operations coexist with spatially separated facilities. The chapter has shown how container flows are divided across deep-sea, feeder and inland segments, and how inter-terminal transport is organised through truck-based services, including the Container Exchange Route, as well as barge and rail shuttles. Empty depots and the spatial distribution of terminals were identified as key drivers of routing choices, resource use and the frequency of intra-port exchanges.

The chapter also outlined the existing automation levels at leading terminals and the port-wide digital coordination layer that links barge operators, terminals and authorities. This context explains both the persistence of bottlenecks for short-distance barge moves, such as quay congestion, limited crane time and misaligned windows, and the feasibility of testing more integrated planning and autonomous barge concepts. These conditions directly connect to the performance dimensions and stakeholder perspectives outlined in Chapter 2 and motivate the scenario design and simulation-based evaluation in Chapter 5.

5

Simulation Methodology and Model Design

Simulation provides a controlled environment to analyse the performance of complex logistics systems and to explore alternative coordination concepts without disrupting real-world operations. For inter-terminal barge transport, simulation makes it possible to evaluate how planning rules, resource constraints and scheduling practices influence vessel turnaround times, quay utilisation and port-wide efficiency.

The purpose of this chapter is to describe the methodological approach taken to model inter-terminal barge transport in the Port of Rotterdam. Section 5.1 frames the problem using CATWOE and a black box view of the system. Sections 5.2 and 5.3 then describes the technical implementation of the simulation, including the choice of simulation approach, the process logic and the main variables and assumptions. Section 5.4 presents the experiments designed for evaluation, together with the model outputs and key performance indicators.

5.1. Problem Framing

This section introduces the conceptual boundary of the system and clarifies who is affected, who can change it, and what the system is supposed to achieve. This is done in two complementary ways. First, a CATWOE description articulates stakeholders and worldview. Second, a black box view separates inputs, transformations, and outputs, keeping the model concise, while still capturing the performance measures that matter for evaluation.

5.1.1. CATWOE Description

This subsection summarises the barge–terminal planning system using the CATWOE lens. The aim is to make explicit who is affected, who acts, what is transformed and which constraints shape feasible changes. The elements are listed in Table 5.1.

The CATWOE framing emphasises that environmental constraints are treated as given, while coordination rules and information sharing are the main levers available to port and terminal actors. It also underlines that performance needs to be interpreted from multiple stakeholder perspectives, which is reflected in the model design and KPI selection in the remainder of this chapter.

Table 5.1: CATWOE description for the barge–terminal planning system.

Element	Description
Customers (C)	Shippers/cargo owners and freight forwarders who experience reliability and dwell time; barge operators whose cycle times and utilisation depend on berth access; hinterland carriers (road/rail) whose pick-up/drop-off reliability depends on terminal synchronisation; ultimately consignees in the hinterland.
Actors (A)	Terminal planners (berth/crane/yard), barge planners/dispatchers, port authority traffic control (nautical services, fairway capacity), platform operators (e.g. Nextlogic/Portbase), customs/inspection services.
Transformation (T)	Convert unsynchronised arrivals and container flows, constrained by quay and crane capacity, into executed service plans (berth allocation, routing, crane assignment) that achieve timely handovers and reduced waiting.
World view (W)	A distributed, data-driven coordination approach (rather than purely bilateral FCFS) can reduce idle time and variability, enabling higher quay productivity and more competitive, low-emission inland waterway transport.
Owners (O)	Port authority (infrastructure, access rules, digital interoperability mandates), terminal operators (resource policies, yard/berth rules), platform owners (algorithm and data-sharing policies). These parties can enable or constrain change.
Environmental constraints (E)	Physical: quay length, draught, lock/bridge windows, fairway capacity. Operational: crane fleets, yard capacity, shift patterns, priority rules. Institutional: contracts, service windows, customs hold/release. Technological: TOS maturity, API availability, autonomy level of barges, charging/shore power.

5.1.2. Black Box System Boundary

The system is represented as a black box, meaning that only the relationship between inputs and outputs is modelled, while internal micro-level details are abstracted away. In this formulation, inputs such as vessel arrivals, resource capacities, and coordination rules enter the system. Within the black box these are processed through allocation and sequencing rules, producing executed service plans. The outputs are twofold; the extent to which operational requirements are met, for example whether barges are served within their requested window, and the performance indicators that allow evaluation of efficiency and reliability.

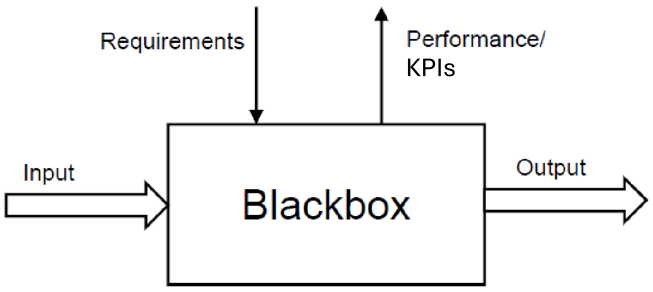

Figure 5.1: Black box view (Veeke, Lodewijks, and Ottjes, 2008).

The key inputs include arrival processes, terminal resources and planning rules, and port constraints. The transformation consists of berth allocation, crane assignment, routing of barges across multiple terminals, including split and return calls, and stochastic handling times. The outputs are measured in terms of requirements and KPIs, including barge turnaround and waiting times and berth and crane

utilisation. This separation highlights that the model is designed as an evaluation tool that captures how different coordination rules transform input conditions into measurable stakeholder outcomes. The detailed event structure and equations that implement this black box view are introduced in the subsequent sections on simulation implementation and process logic.

5.2. Operational Context and Strategies Under Study

The conceptual framing above sets the stage for the technical design of the model. This section explains the methodological choices that shape the simulation, the process logic on which it is built, and the assumptions that define its level of detail. It then shows how the generic model is instantiated with data from the Port of Rotterdam.

5.2.1. Process Logic

The vessel handling process in the simulation is formalised using a Business Process Model and Notation (BPMN) representation. BPMN provides a structured and standardised way of capturing operational processes, highlighting the sequence of activities, the actors responsible, and the decision points that influence vessel–terminal interactions. Compared to a simple flowchart, BPMN also allows for the representation of concurrency (e.g. parallel crane operations), decision gateways (e.g. berth availability), and coordination activities across multiple terminals. To aid interpretation, Figure 5.2 illustrates the basic BPMN symbols used.

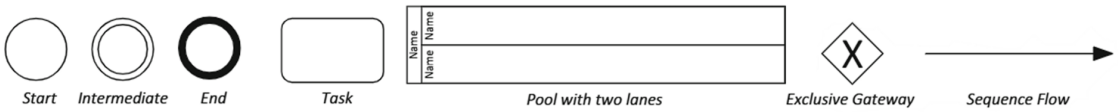

Figure 5.2: BPMN basic elements

The general flow of a vessel through the port can be summarised as follows. Before departure, a vessel communicates its estimated time of arrival (ETA), cargo manifest, and destination terminal(s) to a coordination and data-sharing platform. Based on this information, a preliminary berth and turnaround plan is circulated, ETA, planned time of arrival (PTA), and planned time of departure (PTD), and refreshed at regular intervals as the vessel progresses and terminal conditions change.

On approach, nautical services organised by the Port Authority/harbour master (e.g. VTS, pilotage, towage) coordinate access and guide the vessel to berth. Deep-sea vessels may adjust speed or wait at anchorage to synchronise with berth availability, whereas inland barges, often visiting multiple terminals, typically wait at designated areas inside the port (e.g. buoys/dolphins).

Once berthed, cranes are assigned and container handling begins. Deep-sea vessels are normally worked by three or four ship-to-shore cranes in parallel, whereas inland barges are typically served by a single crane. If cargo is not immediately available for loading, barges may (i) proceed to the next terminal and return later, (ii) accept partial handling, or (iii) rely on re-stowage moves to create space. For small inter-terminal flows, containers are often moved by truck or train rather than barge, as these modes offer greater flexibility with shorter advance notice requirements.

After all planned operations have been completed, the vessel departs its berth. Inland barges continue to their next terminal until all calls are completed, while deep-sea vessels leave the port directly. Departure is again assisted by nautical services. Throughout the process, deep-sea vessels have priority in berth and crane allocation, while inland barges are served on a first-come, first-served basis within their class.

This generic flow is summarised in Figure 5.3, which presents the baseline BPMN representation of the vessel-terminal process at a multi-terminal port. This BPMN diagram forms the conceptual backbone of the simulation.

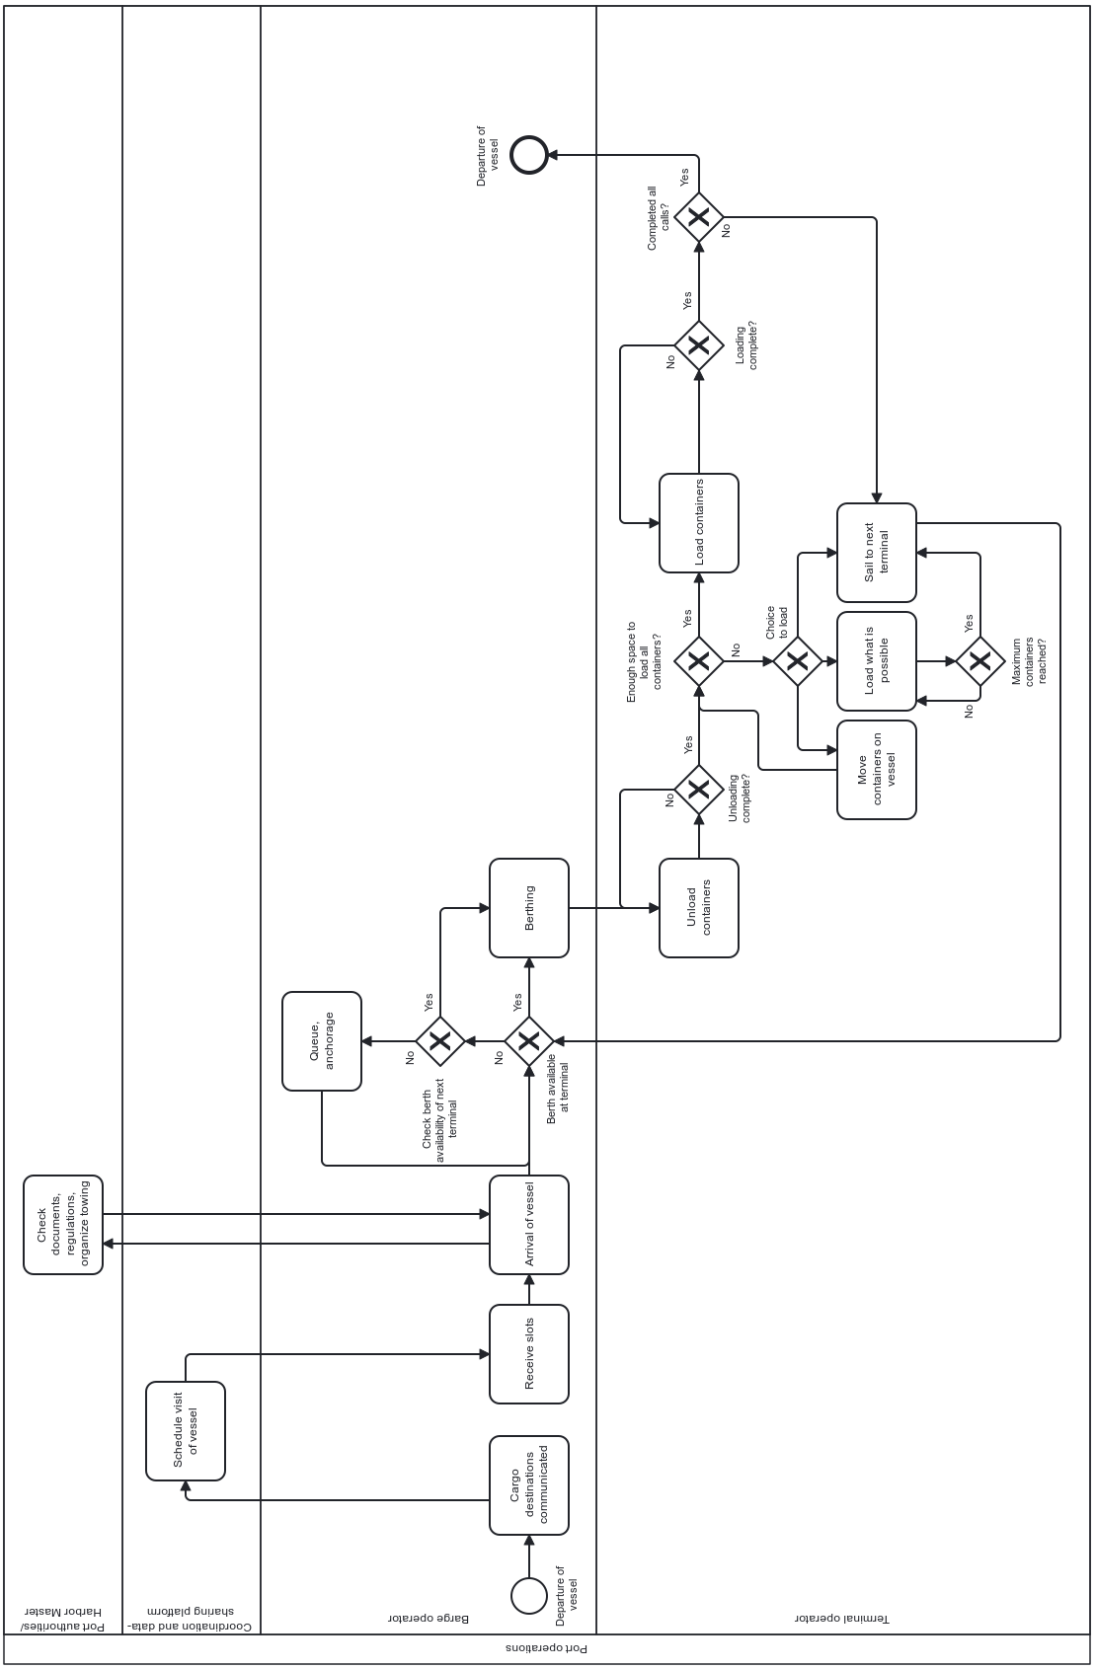

Figure 5.3: General BPMN representation of vessel–terminal interactions in a multi-terminal port.

The process is divided into four lanes, each corresponding to a main actor:

- The coordination and data-sharing platform integrates information from vessels and terminals and supports coordinated slot allocation and plan updates (ETA/PTA/PTD).
- The Port Authority/harbor master provides nautical access and safety services, issues/records arrival clearances and related formalities, and coordinates pilotage and towage for inbound/outbound movements.
- The vessel operator conducts sailing, waiting, berthing, requests/receives handling, and proceeds along the planned route.
- The Terminal operator manages quay and crane resources, confirms berth availability, assigns cranes, and executes loading/unloading, including any necessary re-stowage.

Where case-specific assumptions treat nautical services and clearances as non-delay-generating, interactions in the Port Authority lane are represented as instantaneous events and not explicitly queued (see Section 5.2.3).

This BPMN process logic provides the event structure for the discrete event simulation model. In the implementation, each BPMN element type is mapped to a corresponding simulation construct, summarised in Table 5.2.

Table 5.2: Mapping of BPMN elements to simulation constructs

Element type	Role in the simulation model
Events	Arrival at port boundary, arrival at terminal, berth allocation, handling start, handling finish, departure from terminal, port departure.
Gateways	Berth availability, crane availability, cargo readiness,including the split or return decision, and choice of planning rule, e.g. FCFS with class priority.
Activities	Plan and slot communication and updates, sailing or relocation, waiting at anchorage or buoys, mooring and unmooring, and crane handling with class- and equipment-dependent rates.

5.2.2. Port Setting and Terminals

The simulation is instantiated using a stylised representation of the Maasvlakte area of the Port of Rotterdam. Five container terminals are included, each characterised by its quay length, handling productivity, and position in the distance matrix used for inter-terminal sailing times. This abstraction reduces the full set of Rotterdam terminals to a manageable system size while preserving the spatial separation and resource constraints that shape barge coordination. Data on quay lengths, crane rates, and inter-terminal distances are drawn from publicly available sources provided by the Port of Rotterdam Authority and terminal operators. In the scenarios described later, these five terminals act as the operational environment in which barges perform inter-terminal transport and are therefore the concrete setting for the model experiments.

5.2.3. Rotterdam-Specific Adaptations

To align the generic process with observed Rotterdam practice and the available data, the following adaptations are applied. Where relevant, the coordination and data-sharing platform[1] (referred to as platform) is assumed to provide integrated planning services.

- Integrated coordination (plan, not slots): All vessel classes (deep-sea, feeder, inland barge, autonomous module) are generated by stochastic arrival processes (Section D.6); no liner timetables or standing appointments are modelled. The coordination platform constructs a single conflict-aware berth sequence from these arrivals. Operators follow the issued plan; opportunistic local re-routing

[1] In Rotterdam this role is implemented via Portbase/Nextlogic. The generic term is used here to retain generalisability.

(e.g., skipping a terminal if a berth is busy) is not applied. Residual short waiting may occur due to arrival noise and handling overruns.

- Priority and service discipline: Priority for berth and crane allocation is by vessel class in the order: deep-sea, feeder (shortsea), inland barge/autonomous module. Within each class, first-come, first-served applies. Equivalently, the rule is implemented via a priority score with class weights α_{class} such that $\alpha_{\text{deepsea}} > \alpha_{\text{feeder}} > \alpha_{\text{barge}} \geq \alpha_{\text{module}}$.
- Pre-checks on manifest and vessel capacity: The platform validates cargo manifests and planned stowage at the planning stage so that load and discharge orders are feasible and vessel capacity is not a binding constraint. The intra-call space check is therefore omitted from the BPMN used for simulation.
- Crane allocation norms: Deep-sea vessels are worked by multiple STS cranes in parallel, feeders typically by one or two STS cranes, inland barges by one STS crane, and modules by one dedicated module crane. Exact crane counts and handling rates are parameterised (see Section D.3); the scheduler may allocate differently at run-time.
- Yard blockage: Yard re-stows that free blocked containers are represented as additional handling moves, without an explicit yard process. Yard storage does not bind within the horizons studied, inter-terminal transport is absorbed into handling times.
- Environmental effects: Weather, tide and hydrodynamic effects are omitted.
- Anchorage and reassessment triggers: If a berth is not assigned, the vessel joins anchorage. Anchorage has no capacity limit in the model.
- Nautical services as non-queuing with deterministic time: Interactions with the harbour master and service providers (access clearances, VTS, pilotage and towage dispatch) do not create queues in the model. A deterministic time allowance is included for these activities, with no stochastic delay, as nautical service performance is outside scope.

To reflect these adaptations visually, a simplified Rotterdam BPMN is used in the results that follow (Figure 5.4), in which the manifest/space checks are consolidated at planning time, the cargo-readiness loop is retained, and nautical services are shown as non-queuing activities.

5.2.4. Strategies Under Study

The simulation evaluates two coordination strategies for intra-terminal waterborne transport in the Port of Rotterdam: a baseline reflecting current practice without modular splitting, and a modular splitting concept that allows parallel handling of smaller units. Both strategies share identical case inputs, resource limits, and the same priority policy unless stated otherwise.

Baseline Scenario, Conventional Barge Handling - In the reference case, barges arrive and visit multiple terminals sequentially according to the plan issued by the coordination platform. Deep-sea and feeder calls have priority over inland barges for berth and crane allocation, within each class a first-come, first-served rule applies. Truck and rail flows are not modelled, interactions with those modes are outside the scope of the experiments. No additional rules are added beyond the event flow and queuing logic in Section 5.3.3.

Modular Splitting - A large barge arrives as a single vessel, then splits into autonomous modules within the port area. Each module sails independently to its designated terminal, which enables parallel berthing across multiple terminals. A waiting location is defined for the mother vessel until all modules complete operations, after which units recombine and depart together. Modules are serviced by dedicated module cranes with parameterised handling rates. The process flow is shown in Figure 5.5.

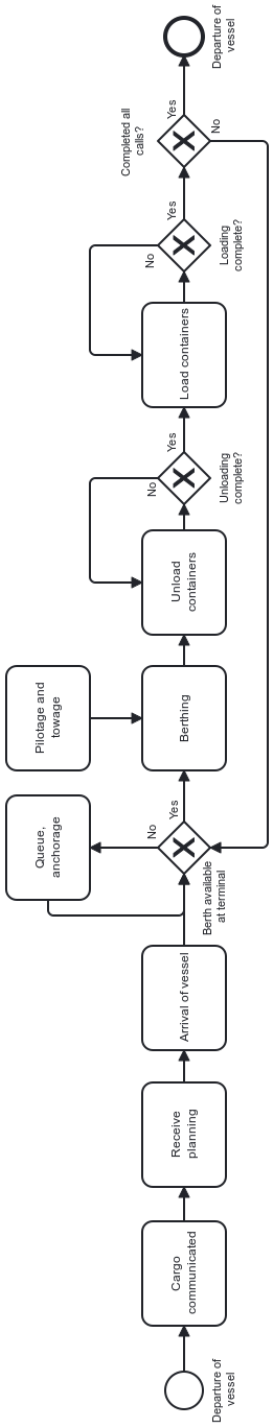

Figure 5.4: BPMN representation of vessel–terminal interactions in the Port of Rotterdam.

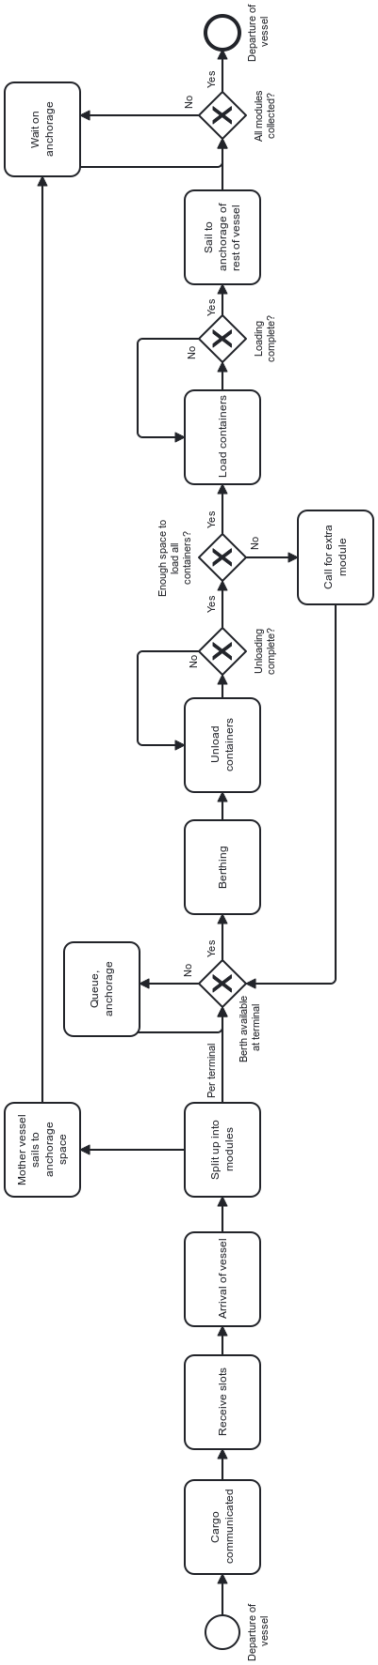

Figure 5.5: BPMN representation of the modular splitting strategy.

5.3. Model Design and Simulation Engine

This section explains how the real-world port process is represented in the simulator and how the engine executes scenarios. It first motivates the choice of modelling approach, then maps the operational process into simulation constructs. The event flow of vessels and terminals is described next, followed by the berth-admission and queuing policy that drives ordering at the quay. The section closes with a concise note on implementation, configuration, and assumptions.

5.3.1. Choice of Simulation Approach

Several simulation paradigms can be applied to study logistics systems, each with its own strengths and limitations. System Dynamics (SD) is well suited to exploring long-term feedback loops and strategic policy effects, but lacks the temporal resolution needed to model event-driven operations such as vessel arrivals and quay allocation. Agent-Based Modelling (ABM) captures heterogeneous decision-making by individual actors, but for this study it would introduce unnecessary complexity at the micro-level (e.g., modelling each crane operator or each barge captain as an autonomous agent). As Abdul Majid et al. (2007) show in their quantitative comparison, DES provides predictable, testable results for sequential processes, whereas ABM offers greater flexibility at the cost of complexity. The research question in this thesis focuses less on behavioural heterogeneity and more on system-level performance under alternative coordination rules, which makes ABM less suitable.

DES is chosen since it provides an appropriate balance between detail and tractability. DES explicitly represents the system as a sequence of discrete events (arrival, berthing, container handling, departure), allowing bottlenecks in quay allocation and resource scheduling to be captured with high temporal accuracy. It also supports the systematic comparison of scenarios under controlled parameter variations, which aligns with the experimental design presented later in this chapter. Moreover, DES has an established track record in port and terminal studies, such as congestion management analyses at marine terminals (Neagoe et al., 2021) and container terminal performance studies (Abourraja, Meijer, and Boukachour, 2021).

For these reasons, DES is the most suitable approach to capture the operational dynamics of inter-terminal barge coordination in the Port of Rotterdam, while keeping the model transparent, interpretable, and aligned with the stakeholder-focused KPIs defined in Section 5.1.

5.3.2. Configuration and Implementation

The behavioural assumptions and policies that govern the model are stated earlier in Section 5.2.3. The present section records how the system is instantiated for the case study and how the simulator executes scenarios. Numerical parameters, distribution fits, and spreadsheet schemas are listed once in Appendix D and referenced here as needed.

The simulated system is a stylised representation of the Maasvlakte container hub with four vessel classes (deep-sea, feeder, inland barge, autonomous module), quays subject to effective-length constraints (vessel length plus fore and aft buffers), two crane pools (a general quay-crane pool and a dedicated module-crane pool), nautical legs between entrance, anchorage, and terminals, and an anchorage that buffers arrivals when berths are unavailable. Five terminals are instantiated with their quay lengths, crane pools, and inter-terminal distances as summarised in Section D.2 (full definitions in Appendix D).

Traffic and work content enter through the demand-generation procedure in Section D.3. For each vessel class, arrivals are realised from the specified stochastic processes and assigned to terminals according to class–terminal shares derived from the case inputs. Per-call load and discharge work is sampled from the class-specific distributions reported in Section D.3. Routing rules determine barge itineraries across multiple terminals. In modular scenarios, a mother vessel anchors in the port area and spawns child modules that sail independently to their designated terminals before rejoining for departure. Sailing and manoeuvring times are applied using the distance matrix and type-specific speeds. Admission and service follow the event flow in Section 5.3.3.

Handling productivity is parameterised per equipment type as a nominal per-crane rate with a bounded multiplicative deviation to reflect operational variability; outages and maintenance are not simulated. Cranes are attached up to the class cap and pool availability, subject to spacing along the quay, and

may be rebalanced as other vessels finish. In the baseline configuration two dedicated module cranes provide flexible handling capacity for module operations; the experiments vary the module-crane pool size and explores substitution with quay cranes to separate policy effects from pure capacity effects (scenario settings are listed in Appendix D).

Runs are configured for a fixed warm-up and analysis horizon, replication count, and precision target as specified in Section D.1. Random seeds are controlled per scenario (and per stream) to enable reproducibility and common-random-number comparisons. Scenario inputs are parameterised in a single settings file. Arrivals and work content are created by the internal generator described in Section D.1.

5.3.3. Simulation Event Flow

The simulator first instantiates the port environment in Figure 5.6. Entrance and departure points are defined, an anchorage is created for vessels that cannot berth, and each terminal is configured with quay length, berth windows, and two crane pools, a general quay-crane pool and a dedicated module-crane pool. Scenario inputs are read, random seeds are set for reproducibility, and the future-event list is prepared.

Traffic is then injected over the horizon. For each vessel type, arrivals are either read from an external schedule or generated from counts that are placed across the horizon with hour-of-day weighting and a small earliness or lateness draw. Each arrival receives a route and a cargo plan. In the baseline setting, a barge may visit several terminals in sequence. In the modular setting, a mother vessel anchors and child modules are created, each with one terminal to serve using the module-crane pool. Vessel attributes such as length and speed are sampled by type, since these determine effective quay usage and sailing times.

As the clock advances, vessels sail from the entrance to the first terminal or to the anchorage if policy dictates. A vessel is admitted to a berth only when three conditions are satisfied. First, enough *effective* quay length is free, the vessel length plus fore and aft buffers that protect crane working space. Second, the berth lies within the terminal's opening window. Third, the crane rule for the class can be met. Barges require at least one free general quay crane, modules draw on the module-crane pool, sea-going vessels may berth even if cranes are not immediately free, since crane attachment can lag berthing without changing safety or sequencing. If any test fails, the vessel sails or waits at anchorage according to the queuing policy in Section 5.3.4.

When a vessel is admitted, cranes are attached up to three limits, the vessel's class cap, the number of free cranes in the relevant pool, and spacing along the quay. The workload for the call is the sum of unload and load moves. Handling productivity is deterministic and type specific through a per-crane rate for that vessel type, and cranes divide the work in parallel. Service time is therefore governed by the amount of work and the average number of cranes actually attached during the call:

$$T_s \approx \frac{\mathrm{TEU}_{\mathrm{unload}} + \mathrm{TEU}_{\mathrm{load}}}{\overline{c}}\, h_{\mathsf{type}},$$

where $\overline{c}$ is the average concurrent cranes on the vessel and h_{type} is the per-crane handling time for the vessel type. Cranes can be rebalanced during handling as other vessels finish, so $\overline{c}$ may change over the call. After a call the vessel either sails to the next terminal in its itinerary or returns to anchorage and re-enters the queue, depending on policy and the state of the next berth. For modular missions, each child module returns to the mother at anchorage once its terminal is finished, and the mother departs only after all children have rejoined.

Departures are recorded at the port exit. Throughout the run, the simulator collects time stamps for sailing, waiting, berthing, and handling, and aggregates these into daily and steady-state indicators.

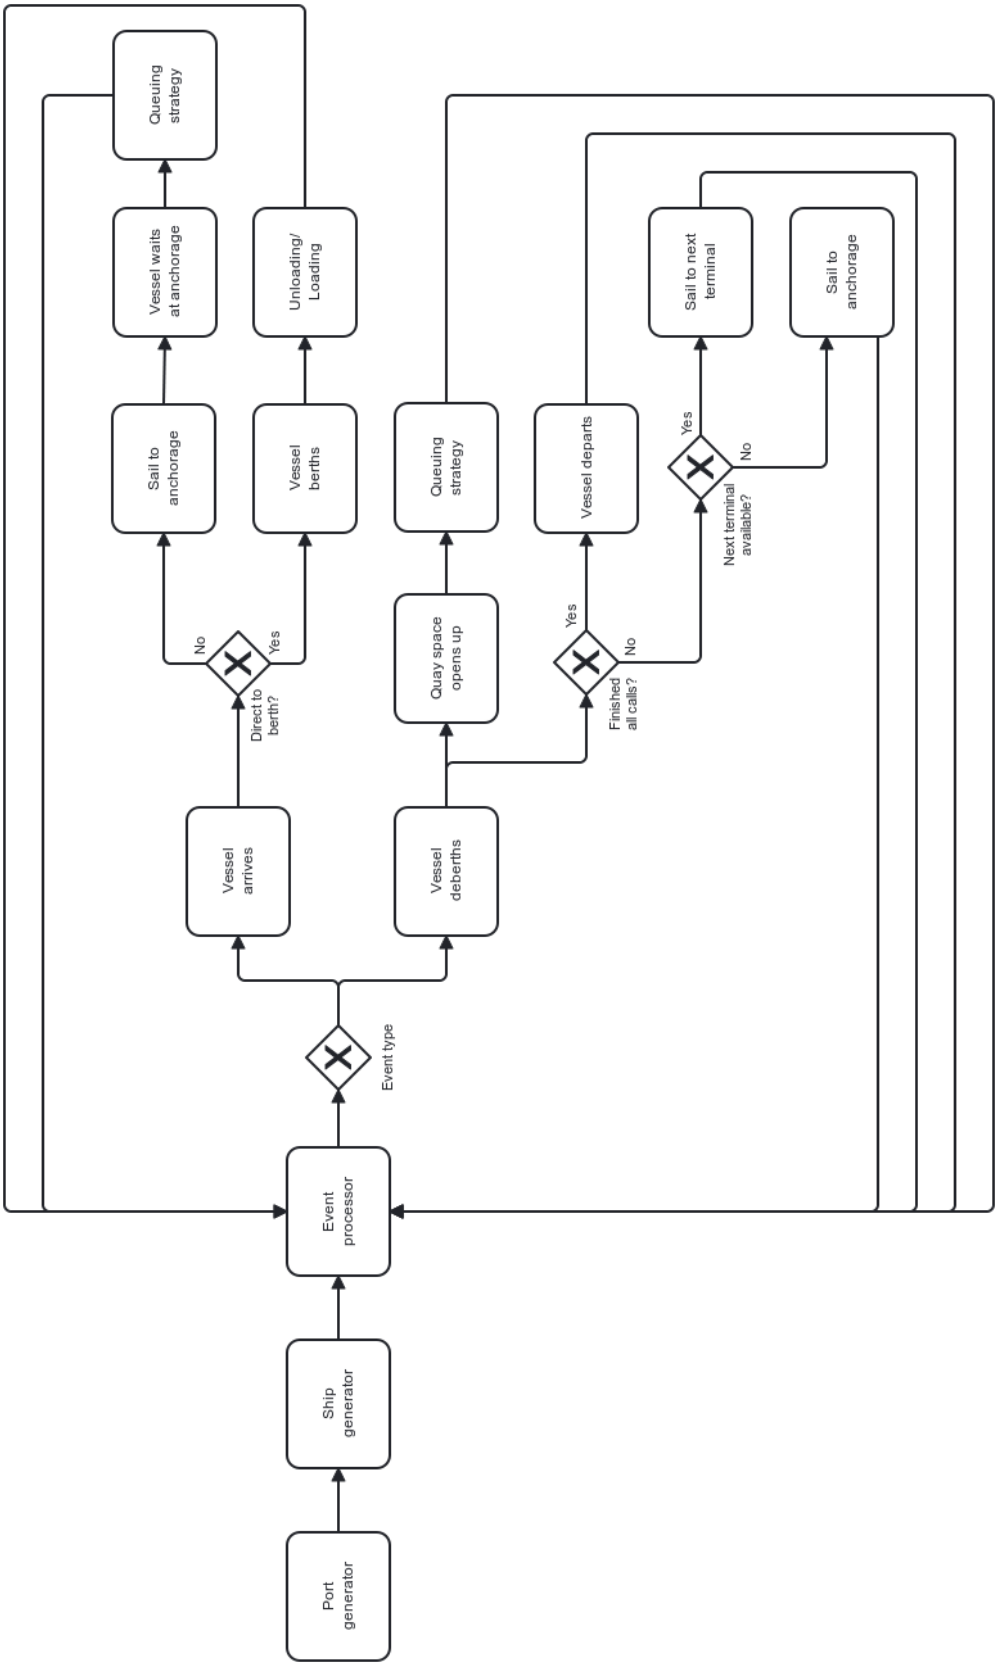

Figure 5.6: High-level event flow of the simulation engine.

5.3.4. Berth Admission and Queuing Policy

Berth admission is priority-based. The order is deep-sea $\rightarrow$ feeder $\rightarrow$ barge/module. Within a class, the queue is first-come, first-served. When quay space opens or handling completes, and at each time step while a berth is idle, the terminal scans the waiting list in that order. The decision flow is shown in Figure 5.7.

A candidate is admitted only if all three gate checks all pass. If a candidate fails any gate, the controller proceeds to the next vessel in the priority order (see the left branch in Figure 5.7). Lower-priority vessels also face a finish sooner test. Let $T_{\text{low}}^{\text{finish}}$ be the predicted completion time of the lower-priority vessel if admitted now, and $T_{\text{high}}^{\text{start}}$ the earliest feasible start of the next higher-priority vessel, including sailing, quay, and crane constraints, with a safety buffer Δ. The lower-priority vessel is admitted only if

$$T_{\text{low}}^{\text{finish}} + \Delta \leq T_{\text{high}}^{\text{start}}.$$

If the test fails, the berth is deliberately kept empty to protect the imminent higher-priority start, and the queue is re-evaluated at the next time step. If no vessel in the ordered list passes all checks, the berth remains idle and the same evaluation repeats every step until a candidate fits. Because deep-sea has the highest priority and no immediate crane requirement at admission, any deep-sea vessel that satisfies the window and effective-quay checks will occupy the berth ahead of lower-priority traffic. Once admitted, the vessel keeps the berth until its call completes and cranes are attached up to the class cap and the free cranes in the appropriate pool, subject to spacing along the quay. Cranes may be rebalanced as other vessels finish.

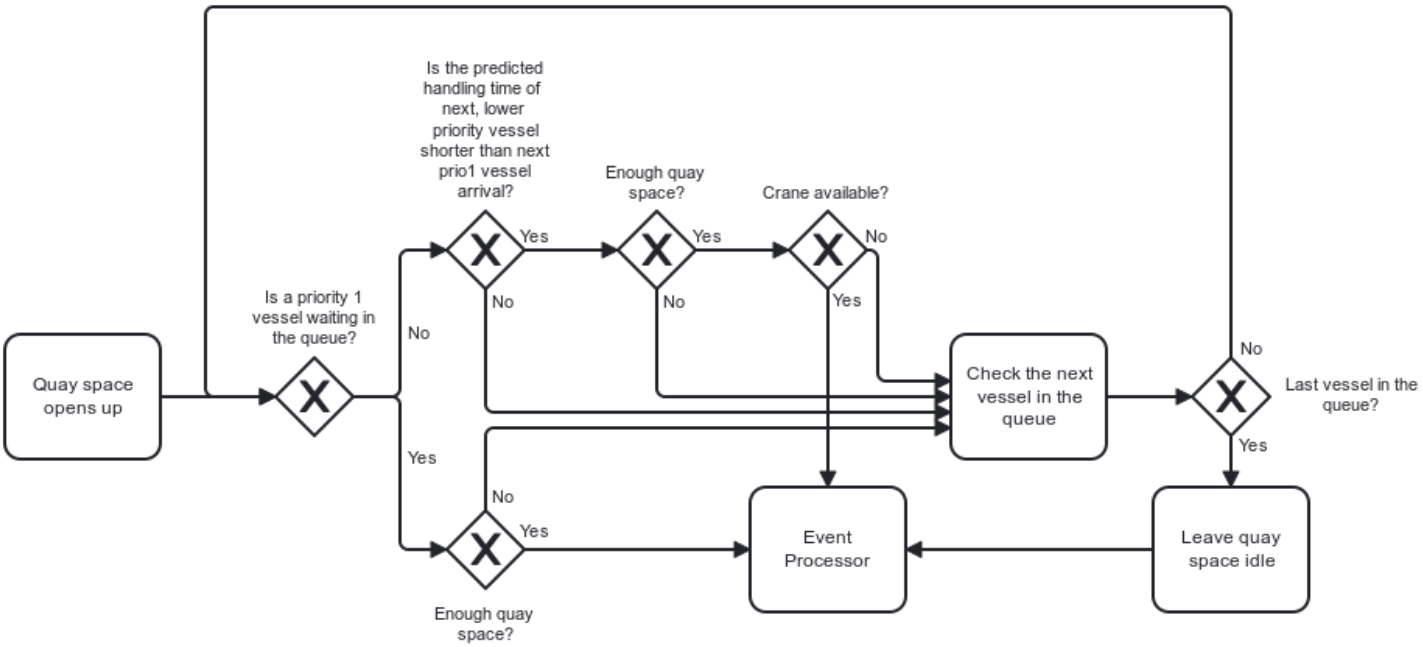

Figure 5.7: Berth-admission decision flow at a terminal.

5.4. Experiments and Analysis Setup

This section defines the experiments used to evaluate the impact of the modular concept on port performance. Each experiment builds on the baseline configuration introduced in Chapter 5, with assumptions, priorities, and simplifications following Section 5.2.3. Collectively, the experiments quantify the performance of the modular concept under varying demand regimes, isolate the effects of equipment availability and module capacity, and assess a near-term adoption scenario that shifts demand from road to barge.

1. Inland vessel ratio design: The mix of inland vessels is varied by adjusting the share of modular vessels relative to conventional barges to find an operating range that balances waiting, service time, and quay utilisation. Levels for the module share are listed in Table D.8. Results from E1 identify a short-list of ratio settings for subsequent tests.

2. Demand robustness under sea-freight variation: The inland vessel ratios are tested under lower and higher exogenous sea-freight demand to examine robustness. Sea demand is scaled by a factor $\alpha \in \{0.8, 1.0, 1.2\}$. The low case relaxes congestion and increases effective quay availability for inland shipping. The high case intensifies competition for terminal resources.

3. Module and equipment configuration: Module design and equipment are varied to identify bottlenecks and performance trade-offs:
   (i) module capacity $c \in \{12, 24, 36\}\,\mathrm{TEU}$, which determines the module design and length overall (LOA),
   (ii) size of the module-crane pool $m \in \{1, 2, 3\}$,
   (iii) module-crane handling rate h as a percentage change around the reference rate h_0.
   Two substitution variants are included where one or two quay cranes are replaced like-for-like by module cranes so that the total number of cranes remains constant. This family evaluates effects on operational flexibility, quay utilisation, and service times.

4. To reflect an early implementation, a scenario shifts $\sigma = 0.10$ of relevant road cargo to barge, with the additional barge work handled by modules. The inland ratio and equipment reflect settings close to current availability.

All scenarios follow the run design in Section D.1: 32 days per run (2 days warm-up, 28 days analysis, 2 days cool-down) with independent replications until the 95% CI half-width of primary KPIs is $\leq 5\%$ of the mean. The warm-up period allows the system to fill and reach a representative operating state before statistics are collected, so that early vessels that experience unrealistically low waiting times are excluded. The cool-down period allows vessels that arrive during the analysis window to complete their port call, so that their turnaround and waiting times are fully observed. Seed details are listed in Appendix D. Across all scenarios, the admission rule, within-class sequencing, the finish-sooner test based on Δ, terminal opening windows and event tick size are held constant. Unless the scenario specifies a substitution, the number of quay cranes is unchanged.

5.4.1. Model Outputs and KPIs

The simulation records time-stamped events for every vessel, berth and crane. The raw traces are aggregated into performance indicators at three levels: vessel, terminal and port. In the landlord-port setting, the public and private actors are primarily concerned with four questions; whether the port can process the offered demand without building up queues, whether sea-going vessels remain well served while inland shipping grows, whether scarce quay and crane resources are used efficiently, and whether access conditions at anchorage remain acceptable. The KPI families introduced in the literature review (Chapter 2) are instantiated here to address these questions. The KPIs focus on five primary groups:

- throughput,
- turnaround and waiting times by vessel class,
- berth occupancy and effective berth utilisation,
- crane utilisation and productivity,

- anchorage and access behaviour.

Together, these indicators test several of the benefits attributed in the literature to autonomous and modular inland concepts, such as improved use of inland capacity, better use of quay length, and maintaining service levels for sea-going traffic. Autonomy or modularity related claims about safety and fine-grained schedule reliability are not evaluated directly, because the present model uses stylised demand patterns and limited uncertainty. Reliability appears only indirectly, through the spread of waiting and service times.

Vessel KPIs - At vessel level, the analysis tracks waiting, service and turnaround times per vessel class. For modular configurations, mothers and modules are reported separately so that the division of work and the associated berth use become visible.

For each completed call i in class c the following event times are recorded: arrival to the port area a_i, berth start b_i, service completion c_i, and departure from the port d_i. Service begins at the first lift and ends at the last lift. The time components per call are

$$\text{Waiting time } W_i = b_i - a_i, \qquad \text{Service time } S_i = c_i - b_i, \qquad \text{Turnaround time } T_i = d_i - a_i.$$

These show how access rules and vessel mixes affect service quality for each vessel class. They are complemented, where relevant, by empirical distributions and percentiles to indicate the spread in time spent in port. Given the simplified and mostly deterministic demand processes, these distributions are interpreted as indicative variation within the model rather than as full schedule-reliability measures.

Terminal KPIs - At terminal level, the model reports how each terminal converts berth time and crane effort into throughput. The primary terminal KPIs used in the experiments are (daily) throughput in TEU, crane productivity, mean service time per call, and berth occupancy and effective utilisation. These quantify whether the modular concept allows terminals to handle more inland work without eroding sea-side performance or requiring disproportionate resource use.

Let terminal k have analysis horizon H hours, D analysis days and B_k berths. Let TEU_i denote the handled volume of call i (discharged plus loaded TEU).

Throughput per day in TEU is

$$\mathsf{TH}_k = \frac{1}{D} \sum_{i \in k} \mathrm{TEU}_i,$$

where the sum runs over all calls handled at terminal k in the analysis window.

Crane productivity measures how efficiently active crane time is converted into moves. For terminal k,

$$\mathsf{CP}_k = \frac{\text{total moves at } k}{\sum_{\ell \in \text{cranes at } k} \text{active hours}_\ell},$$

where moves are counted as TEU-equivalent lifts and active hours are the hours during which crane ℓ is assigned to a vessel and working. In Chapter 7, this is reported as moves per active crane-hour.

Berth occupancy and effective utilisation distinguish between time and space use. The time-based berth-occupancy ratio is

$$\mathsf{BOR}_k^{\text{time}} = \frac{\sum_{j=1}^{B_k} \text{busy time of berth } j}{B_k \, H},$$

that is, the share of open berth-hours spent handling vessels.

In addition, the model tracks the effective occupied quay length $L_k^{\mathsf{occ}}(t)$ over time, which accounts for vessel length and safety buffers to neighbouring vessels and berth limits. With total quay length L_k, the length-based berth utilisation over the horizon is summarised by

$$\mathsf{BOR}_k^{\mathsf{eff,avg}} = \frac{\overline{L_k^{\mathsf{occ}}}}{L_k}, \qquad \mathsf{BOR}_k^{\mathsf{eff,max}} = \max_t \frac{L_k^{\mathsf{occ}}(t)}{L_k},$$

where $\overline{L_k^{\mathsf{occ}}}$ is the time average of the occupied effective length from the simulation time series. The average ratio indicates how heavily the quay is used on a typical basis, while the maximum ratio shows how close the quay comes to being fully packed at peak times, which is important for understanding blocking and waiting.

In the results, references to effective berth utilisation use this length-based measure, while berth occupancy refers to the time-based ratio. This definition is consistent with the continuous berth-allocation rectangle-packing representation of quay planning discussed in Section 2.3.2 and in berth-allocation studies such as Dai et al. (2008) and Frojan et al. (2015). It provides the link to quay and berth productivity indicators used later to assess whether modular inland traffic makes better use of available quay length.

Port-wide KPIs - At the port scale, terminal results are aggregated to describe the overall behaviour of the system and the distribution of benefits between vessel classes. The focus is on whether the modular concept supports a higher inland share while maintaining acceptable sea-going performance and manageable resource use.

Total port throughput is the sum of terminal throughputs,

$$\mathsf{TH}_{\mathsf{port}} = \sum_k \mathsf{TH}_k.$$

Composition by vessel class is summarised mainly through the TEU based share of each class in total throughput. For a given vessel class c,

$$\mathsf{ShareTEU}_c = \frac{\sum_{i \in c} \mathrm{TEU}_i}{\sum_i \mathrm{TEU}_i} \times 100\%.$$

In the results, this is reported in particular as the inland share of total port throughput, obtained by aggregating conventional barge and module calls into a single inland class $\mathcal{I}$,

$$\mathsf{InlandShare} = \frac{\sum_{i \in \mathcal{I}} \mathrm{TEU}_i}{\sum_i \mathrm{TEU}_i} \times 100\%.$$

These composition indicators are used in the results and discussion chapters to show how introducing modular inland vessels increases the opportunities for inland traffic to be handled at terminals, without displacing sea-going work, by capturing a higher inland share of the processed throughput. This links to the literature on modal split and on how the benefits from improved coordination are distributed between actors.

5.5. Conclusion

This chapter has defined the simulation framework used to evaluate barge coordination strategies in a multi-terminal port. The CATWOE and black box framing established the system boundary, highlighted the main stakeholders and decision levers, and clarified that environmental constraints are treated as given while coordination rules and information sharing are the primary variables of interest. Building on this, the BPMN process logic and the Rotterdam-specific adaptations translated the generic vessel–terminal interaction into a stylised Maasvlakte setting with five terminals, integrated planning through a coordination platform, and priority-based service disciplines.

The model design and simulation engine then mapped this operational context onto a discrete event simulation with four vessel classes, effective quay-length constraints, shared and dedicated crane pools, and a priority-based berth-admission and queuing policy. Demand generation procedures, configuration choices and implementation details were specified to ensure that traffic, work content and routing reflect realistic ranges while remaining tractable for scenario analysis. Finally, the experimental design and KPI definitions set out how alternative mixes of modular and conventional barges, sea-freight demand levels and module configurations will be compared in terms of throughput, waiting and turnaround times, berth and crane utilisation, and anchorage behaviour.

Together these elements establish a transparent and reproducible simulation environment that captures the key interactions between vessels, terminals and coordination rules, and is sufficiently flexible to represent both current practice and modular splitting concepts. The next chapter presents verification and validation to assess the correctness of the implementation and the plausibility of outcomes against reference data and stylised behaviours. The following chapter reports the experimental results and interprets performance trade-offs across scenarios and sensitivity variants.

6

Verification and Validation

This chapter explains how confidence in the simulation is established. Verification considers whether the implementation reflects the intended event logic and rules, validation considers whether the model is sufficiently accurate for analysing coordination strategies for barge and modular operations in the Port of Rotterdam. The main text reports only the essential checks that meaningfully de-risk the implementation, while the detailed catalogue, plots, and traces are placed in Appendix E.

6.1. Verification

Verification is carried out to ensure that the simulation model correctly implements the intended logic and that no programming or conceptual errors distort the results. The process focused on confirming that the translation from the conceptual framework into the discrete event simulation was consistent and reliable.

The first step was to examine the event flow of the model in detail. By running simplified test cases with a single vessel or a small number of calls, the sequence of arrivals, berth admission, and crane allocation could be traced step by step. These traces were compared with the expected outcomes derived from the BPMN process logic. In particular, the berth admission policy was tested under controlled conditions to confirm that vessels were only admitted when sufficient quay length and crane capacity were available, and that the queuing discipline followed the intended first come, first served logic.

Special attention was given to the modular splitting strategy, as this introduced additional complexity in the form of synchronisation between modules. To verify this mechanism, deterministic scenarios were created in which modules arrived simultaneously and were allocated to different terminals. The recombination rules were then checked by reproducing simplified Gantt charts by hand, ensuring that departure only occurred once all modules had completed their operations. These checks confirmed that the model respected the intended synchronisation logic and did not allow premature departures.

Resource rules and conservation principles were also verified. Edge-case scenarios were used to test crane allocation bounds, berth length constraints, and quay pocket availability. For example, scenarios with maximum berth occupancy were run to confirm that no vessel was admitted beyond the physical quay length. Similarly, crane allocation was checked to ensure that the minimum crane rule was always satisfied and that crane productivity was conserved across operations. Container balance and time accounting were monitored to confirm that no containers were lost or created within the system and that turnaround times matched the sum of handling and waiting components.

Finally, reproducibility was verified by running the model with fixed random seeds. This ensured that identical input conditions produced identical outputs, which is essential for sensitivity analysis and scenario comparison. In addition, single-vessel hand calculations were performed to validate turnaround times and berth occupancy under simplified conditions. These deterministic checks provided confidence that the model's outputs were consistent with analytical expectations.

Taken together, these verification activities demonstrate that the simulation engine faithfully represents

the intended operational logic. The model can therefore be considered sufficiently robust to support the comparative analysis of coordination strategies presented in the subsequent chapters.

6.2. Validation

Validation was conducted to assess whether the simulation model provides a credible representation of intra-terminal barge operations in the Port of Rotterdam. Whereas verification focused on the internal correctness of the model, validation addressed its external realism by comparing model behaviour with empirical data and expert judgement.

The first step involved comparing model outputs with publicly available port statistics. Key indicators such as the proportion of vessels proceeding directly to berth versus waiting at anchorage, average turnaround times, and berth occupancy levels were contrasted with reported figures for the Maasvlakte terminals. While the model does not aim to reproduce exact historical values, the alignment of trends and magnitudes provided confidence that the simulation captures the essential dynamics of port operations.

In addition to data comparison, expert interviews were conducted with practitioners familiar with barge handling and terminal planning in Rotterdam. These interviews served two purposes: to confirm that the process logic and assumptions embedded in the model reflect operational practice, and to evaluate whether the simulated outcomes were plausible from a practitioner's perspective. Experts confirmed that the representation of berth admission, crane allocation, and the challenges of synchronising modular operations were consistent with their experience. They also highlighted that the model realistically reproduced the variability and clustering of barge calls, which are known bottlenecks in practice.

Together, the data comparison and expert validation suggest that the model provides a sufficiently accurate representation of the system for the purpose of comparative scenario analysis. It should be emphasised that the model is not intended as a predictive tool for exact throughput figures, but rather as an experimental environment to explore the relative performance of coordination strategies under varying conditions. Within this scope, the validation results support the credibility of the findings presented in the results chapter.

6.3. Conclusion

The verification activities confirmed that the simulation model is internally consistent and that its logic correctly implements the intended coordination strategies. By tracing simplified scenarios, checking resource rules, and validating conservation principles, confidence was gained that the model behaves as designed.

Validation then demonstrated that the model provides a credible representation of intra-terminal barge operations in the Port of Rotterdam. The comparison with publicly available port statistics showed that the model reproduces the main operational patterns, while expert interviews confirmed that the assumptions and outcomes are plausible from a practitioner's perspective.

Together, these steps provide sufficient assurance that the simulation results presented in the next chapter are both reliable and meaningful for analysing alternative coordination strategies.

7

Results

This chapter presents the outcomes of the simulation experiments defined in Chapter 5 and implemented as described in Appendix D. The experiments explore three main factors: the composition of inland traffic between conventional barges and autonomous modules, the intensity of sea-freight demand, and implementation choices for the autonomous-module concept, including crane layouts, handling speeds, and module capacity. For each factor, results are reported in terms of port-wide throughput and inland share, turnaround and waiting times by vessel class, and the use of berth length and crane capacity at terminals. Unless stated otherwise, values are averages over multiple independent replications of a four-week horizon.

Section 7.1 compares different barge–module mixes under the baseline demand pattern and identifies how modular uptake affects throughput, turnaround, and resource use. Section 7.2 examines how changes in sea-freight demand (80, 100, and 120 % of the baseline) influence these indicators in selected barge–module mixes. Section 7.3 analyses configuration choices for the modular system, including the number of dedicated module cranes, their handling speed, and the capacity of individual modules. Section 7.4 studies an early modular adoption scenario in which a modest increase in inland throughput is accommodated by additional module services, illustrating how future growth in inland demand can be absorbed by scaling up the modular fleet rather than expanding conventional barge or road capacity.

7.1. Results for Different Barge and Modules Ratios

This section compares four mixes of conventional inland barge services and autonomous module services; 75:25, 50:50, 25:75, and 0:100, relative to the baseline (100:0). The baseline represents today's organisation without modules. Increasing the share of modules raises the frequency of short-haul exchanges, changes berth access dynamics, and shifts crane loading patterns. Outcomes are reported per KPI below, with percentage changes summarised in Table 7.7.

7.1.1. Vessel KPIs

Vessel-level performance is summarised in terms of average turnaround time in the port area and the decomposition of turnaround time into terminal and waiting components for barges and modules. Table 7.1 reports average turnaround time by vessel class and barge–module mix, measured from arrival at the port area to departure. Table 7.2 then shows how much of this time is spent at terminals and in waiting for barges and module children.

Average turnaround time decreases monotonically for all vessel classes as the module share increases, with the strongest effects for inland services (Table 7.1). Barge turnaround time falls from 33.7 h in the 100:0 baseline to 10.5 h in the 25:75 mix, after which barges disappear in the 0:100 case. Deep-sea calls shorten more modestly, from 28.5 h to 25.6 h between 100:0 and 0:100, and feeders improve from 15.4 h to 12.9 h over the same range. These gradients show that modular uptake primarily benefits inland services yet still delivers a small but consistent reduction in seagoing turnaround. Module-mother

turnaround times remain very short in the mixed cases, around 1.3–2.1 h, and only increase sharply in the fully modular 0:100 scenario, where they reach 33.1 h.

Table 7.1: Average turnaround time by vessel class and barge-module mix (hours). Turnaround time is measured from arrival at the port area to departure.

Vessel class	100:0	75:25	50:50	25:75	0:100
Deep-sea	28.5	27.8	26.6	26.0	25.6
Feeder	15.4	14.8	13.7	13.1	12.9
Barge	33.7	23.9	15.6	10.5	—
Module (mother)	—	1.3	1.4	2.1	33.1

Note: — denotes a class not present in that mix (no barges at 0:100 and no modules at 100:0).

The breakdown in Table 7.2 (results for all vessel types in Appendix F, Table F.1) shows that the barge improvement is driven almost entirely by reduced waiting rather than faster handling. As detailed in the appendix, average barge time at terminals stays close to 8 h across all mixes where barges are present, whereas average barge waiting time drops sharply from 23.6 h in the 100:0 baseline to 14.5 h at 75:25, 6.7 h at 50:50, and 1.8 h at 25:75. In other words, additional module traffic removes inland vessels from anchorage and terminal queues rather than speeding up crane operations.

For module children, average waiting remains below 1 h in all mixed cases and only becomes substantial in the fully modular 0:100 configuration, where it jumps to 25.3 h. The 0:100 system therefore clears inland work through very short handling cycles but, once modules constitute all inland traffic, does so at the cost of long off-berth queues.

Table 7.2: Average waiting time for barges and module children across barge:module mixes (hours).

Vessel class	100:0	75:25	50:50	25:75	0:100
Barge	23.6	14.5	6.7	1.8	—
Module (child)	—	0.1	0.2	0.6	25.3

Note: — denotes a class not present in that mix (no barges at 0:100 and no modules at 100:0).

Taken together, these results show that mixed barge–module fleets substantially shorten inland turnaround times and slightly improve seagoing turnaround, primarily by reducing waiting rather than changing service times.

7.1.2. Terminal KPIs

Terminal operations are assessed in terms of berth utilisation and crane utilisation. Berth utilisation is measured as effective berth utilisation, that is, the share of quay metre–hours that is occupied by vessels based on effective vessel length (LOA plus buffers). Table 7.3 reports the average and peak effective utilisation at Terminal B for each barge–module mix. Terminal B is chosen as a representative case, because its quay length and crane capacity are close to the system average. The behaviour is more indicative of a typical terminal than of extreme small or large facilities.

Average effective berth utilisation at Terminal B declines systematically as the module share increases, from about 80% in the 100:0 baseline to roughly 39% in the 0:100 case (Table 7.3). Peak effective utilisation remains close to full occupancy in all mixes, between about 97% and 99%. The terminal therefore continues to experience short periods in which almost the entire quay is occupied, but the proportion of time spent in such states falls as more inland work is handled by modules. Combined with the higher port-wide throughput observed in the mixed and modular cases, this pattern shows that Terminal B processes more containers per metre–hour of quay while operating at lower average length–time occupancy. Shorter, more frequent calls and better packing of work into small gaps along the quay lift berth productivity without pushing the terminal into sustained high-occupancy regimes.

Crane utilisation at Terminal B reflects both the change in inland mix and the difference in handling speeds between barges and modules (Table 7.4). The number of installed cranes is fixed at 16 general and 2 module units, but in the 100:0 configuration all inland work is handled by the general pool and

Table 7.3: Average and peak effective berth utilisation at Terminal B by barge:module mix.

Ratio	Avg. effective utilisation (%)	Peak effective utilisation (%)
100:0	80.2	99.2
75:25	72.7	99.3
50:50	64.3	99.2
25:75	52.8	98.3
0:100	38.9	97.4

Effective berth utilisation is defined as the share of quay metre-hours used, based on effective vessel length (LOA plus buffers).

the module cranes remain idle. As the module share increases, average use of the general cranes falls from about 12.0 to 5.3 cranes, while average use of the module cranes rises from 0.0 to almost 2.0 cranes. In the fully modular 0:100 case, average module use reaches 1.96 out of 2 cranes, so average and peak use of the module pool are almost identical, indicating that the dedicated module cranes are close to fully deployed whenever modules dominate inland traffic.

This shift is driven by two effects. First, the number of barge calls decreases and eventually disappears in the 0:100 case, so a large inland workload is removed from the general crane pool; at that point general cranes mainly serve deep sea and feeder calls, which are fewer in number than the barge calls in 100:0 even though several cranes may be attached to a seagoing ship at once. Second, module cranes handle containers roughly twice as fast as the barge operations on the general quay (about 1.5 versus 3 min/TEU), so the same inland volume requires fewer crane hours when it is processed by modules. As a result, total average crane use at Terminal B decreases from about 12.0 to 7.2 cranes and the average to installed ratio drops from roughly 0.67 to 0.40 as the mix moves from 100:0 to 0:100, even though, as discussed in the throughput results in the next subsection, the mixes with modules handle more TEU. Together with the berth utilisation results, this indicates that introducing modular services at Terminal B redistributes crane loading from the general to the dedicated pool and creates additional slack in installed capacity, rather than driving the terminal towards a congestion regime.

Table 7.4: Crane capacity and utilisation at Terminal B.

	General			Module			Total			
Ratio	Inst.	Peak Used	Avg. Used	Inst.	Peak Used	Avg. Used	Inst.	Peak Used	Avg. Used	Avg. vs Inst. (%)
100:0	16.00	16.00	12.01	0.00	0.00	0.00	16.00	16.00	12.01	75.06%
75:25	16.00	16.00	10.57	2.00	2.00	0.52	18.00	17.98	11.10	61.67%
50:50	16.00	15.98	9.08	2.00	2.00	1.06	18.00	17.94	10.14	56.33%
25:75	16.00	15.93	7.25	2.00	2.00	1.58	18.00	17.83	8.83	49.06%
0:100	16.00	15.58	5.27	2.00	2.00	1.96	18.00	16.58	7.23	40.17%

7.1.3. Port-wide KPIs

At port level, the main indicators are total throughput, the inland share of that throughput, and the way vessels are routed between direct berthing and anchorage. Table 7.5 summarises total and inland throughput for the five barge–module mixes.

Total port throughput increases almost monotonically with the module share (Table 7.5). Starting from 1.12×10^6 TEU in the 100:0 baseline, throughput rises to 1.32×10^6 TEU in the 0:100 mix, corresponding to gains of about 7%, 12%, 17%, and 18% relative to the baseline. Inland throughput (barges and modules combined) grows even faster, from 2.20×10^5 to 4.12×10^5 TEU, and the inland share of port throughput increases from 19.7% to 31.2%. Modular uptake therefore raises overall volumes while shifting a larger fraction of port activity to inland services.

Anchorage behaviour differs by vessel class and shifts systematically with the barge:module mix (Table 7.6). Deep-sea and feeder services maintain low anchorage exposure across all scenarios. Conventional barges start from a much heavier anchorage regime in the 100:0 baseline, where roughly

Table 7.5: Port-wide throughput and inland share for the baseline and barge:module mixes. Inland throughput is reported as total TEU of barges and modules.

Ratio	Total TEUs ($\times 10^6$)	Change vs 100:0 (%)	Inland TEUs ($\times 10^5$)	Inland share (%)
100:0	1.12	-	2.20	19.7
75:25	1.20	+7.1	2.91	24.3
50:50	1.26	+12.3	3.49	27.8
25:75	1.31	+16.6	3.96	30.3
0:100	1.32	+17.7	4.12	31.2

32% of arrivals route via anchorage and average anchorage waiting time is 8.4 h. As the module share increases, the barge anchored share drops to about 24% at 75:25, 14% at 50:50, and 6% at 25:75, and average barge anchorage waiting time falls to about 0.5 h. Module children follow the opposite pattern: they are almost always routed direct in the mixed cases, with anchored shares between about 6% and 64% and average waits below 1 h, but in the 0:100 mix nearly all module children are staged off-berth (around 100% anchored share) with an average anchorage wait of 25.3 h.

Table 7.6: Anchorage behaviour by vessel class and barge:module mix. “Anch.” and “Direct” give the percentage of arrivals of that class routed via anchorage or directly to berth.

	Deep-sea			Feeder			Barge			Module (child)		
Ratio	Anch. (%)	Direct (%)	Wait (h)	Anch. (%)	Direct (%)	Wait (h)	Anch. (%)	Direct (%)	Wait (h)	Anch. (%)	Direct (%)	Wait (h)
100:0	15.0	85.0	0.74	6.1	93.9	0.36	31.5	68.5	8.36	—	—	—
75:25	13.8	86.2	0.64	5.8	94.2	0.42	23.9	76.1	4.90	6.3	93.7	0.13
50:50	10.7	89.3	0.36	4.8	95.2	0.19	14.2	85.8	2.14	27.7	72.3	0.20
25:75	8.1	91.9	0.27	4.1	95.9	0.10	6.4	93.6	0.54	63.5	36.5	0.57
0:100	7.0	93.0	0.27	3.1	96.9	0.10	—	—	—	99.9	0.1	25.26

In the mixed configurations, module anchorage is largely a matter of short staging around window alignment and safety separations, since the dedicated module crane pool usually has spare capacity. In the fully modular 0:100 case, by contrast, the high inland volume keeps the module cranes close to fully utilised. Because module crane availability is explicitly checked in the berth-admission logic, modules are routed to anchorage whenever the module pool is saturated or expected to be saturated at the planned start time. The very high anchored share and long average anchorage wait for module children in 0:100 are therefore a direct consequence of relying almost entirely on a heavily loaded dedicated crane pool.

The reported waiting times are per call. Barges that visit several terminals accumulate waiting across multiple calls, so the reduction in per-call anchorage waiting from 8.4 h to well below 1 h represents a substantial reduction in total inland delay per round trip. Taken together with the vessel indicators, terminal indicators, and port-wide KPIs, these results show that mixed barge–module fleets increase total and inland throughput, relieve barge anchorage and waiting, and keep seagoing traffic largely unaffected. Only the fully modular 0:100 configuration shifts significant queuing from terminals to anchorage for the module fleet, in line with the high inland volumes and strong reliance on the dedicated module-crane pool in that scenario.

7.1.4. Overview of Ratio Results

Table 7.7 summarises the main percentage changes relative to the baseline (100:0) across vessel, terminal, and port-wide KPIs. The pattern is consistent; higher module uptake shortens inland turnaround, reduces barge waiting, lowers average berth and crane loading, and simultaneously raises total and inland throughput.

The 75:25 mix already reduces barge time in port by about 30% and barge waiting by almost 40%, while increasing total throughput by just over 7% and inland throughput by roughly one third. Moving to 50:50 more than halves barge turnaround time and cuts barge waiting by over 70%, with a 12% gain

Table 7.7: Summary of relative changes in key KPIs compared with the baseline (100:0). Positive values denote increases, negative values reductions.

KPI (% vs 100:0)	75:25	50:50	25:75	0:100
Vessel KPI				
Deep-sea turnaround time	-2.5%	-6.7%	-8.8%	-10.2%
Feeder turnaround time	-3.9%	-11.0%	-14.9%	-16.2%
Barge turnaround time	-29.1%	-53.7%	-68.8%	N/A
Barge waiting time	-38.5%	-71.5%	-92.2%	N/A
Terminal KPI				
Avg. effective berth utilisation	-9.4%	-19.8%	-34.2%	-51.5%
Avg. cranes used (total)	-7.6%	-15.6%	-26.5%	-39.8%
Port-wide KPI				
Total throughput	+7.1%	+12.3%	+16.7%	+17.7%
Inland throughput (barge + module)	+32.2%	+58.6%	+80.0%	+86.8%
Inland share	+4.6%	+8.1%	+10.6%	+11.6%

in total throughput and almost 60% more inland TEUs. The 25:75 and 0:100 cases continue this trend, but with clear diminishing returns in total throughput: inland throughput and inland share keep rising, while average berth utilisation and general-crane loading fall sharply.

Taken together, these results show that intermediate module shares, particularly 50:50 and 25:75, capture most of the system-wide benefits. They combine substantial reductions in inland turnaround and waiting with higher throughput and lower average resource loading, while avoiding the extreme anchorage and long module waits observed in the fully modular 0:100 case. For this reason, the subsequent experiments focus on 50:50 and 25:75 to examine how these near-optimal mixes behave under varying demand and resource conditions.

7.2. Effects of Different Demands

This section isolates the effect of sea-freight intensity while keeping all other settings as in Chapter 5. Sea-side arrival rates for deep-sea and feeder services are scaled to 80%, 100%, and 120% of the baseline demand vector, with call mix, priorities, berth windows, and crane inventories unchanged. Two module shares that were close on most KPIs in the previous analysis, 50:50 and 25:75, are carried forward to examine whether demand interacts with the balance between general quay cranes and dedicated module cranes. The objective is to quantify the elasticity of KPIs with respect to demand and to identify emerging bottlenecks.

7.2.1. Port-wide KPIs

Table 7.8 reports total port throughput in TEU for sea-freight loads of 80%, 100%, and 120%. As expected, volumes scale almost linearly with the imposed sea-freight intensity: a 20% reduction in deep-sea and feeder arrivals lowers total throughput by roughly 14–15% across all three mixes, while a 20% increase raises throughput by a similar margin. The slopes are very similar for 50:50 and 25:75, which shows that within this demand band modular uptake does not materially change the elasticity of port throughput with respect to sea-side demand.

For a given sea-freight level, both modular mixes process substantially more TEU than the conventional baseline. At 100% demand, total throughput increases from 1.12 million TEU in the 100:0 case to 1.26 million TEU in 50:50 and 1.31 million TEU in 25:75, gains of about 12% and 17% respectively. Similar uplifts appear at 80% and 120% demand, which confirms that the throughput advantage of modular mixes is robust across the tested sea-freight range.

There are two small second-order effects. Relative to their own 100% case, the modular mixes lose slightly less throughput when sea demand is relaxed to 80% (drops of 14.6% and 14.3% versus 14.9% in 100:0), reflecting the fixed inland component that remains unchanged in these experiments. Conversely, at 120% demand the 100:0 system shows a marginally steeper relative increase than the modular mixes (14.9% versus 13.6%), indicating that, at the upper end of the tested band, additional sea

calls translate into throughput slightly more strongly in the conventional system. These differences are small, and the dominant pattern remains that at all three sea-freight levels the mixed systems handle more TEU in absolute terms than the 100:0 baseline.

Table 7.8: Total port throughput under $\pm 20\%$ sea-freight demand, per barge–module mix. Throughput is reported in $\times 10^6$ TEU; percentage changes are relative to the 100% sea-freight case for the same mix.

	100:0		50:50		25:75	
Demand	Throughput	Δ vs. 100%	Throughput	Δ vs. 100%	Throughput	Δ vs. 100%
80%	0.95	-14.9%	1.07	-14.6%	1.12	-14.3%
100%	1.12	-	1.26	-	1.31	-
120%	1.29	$+14.9\%$	1.43	$+13.6\%$	1.48	$+13.6\%$

The inland component of throughput is summarised in Table 7.9. Across mixes, higher modular uptake consistently raises inland throughput and its share in the port total at any given sea-freight level. Combined with the total throughput results in Table 7.8, this confirms that the mixed systems handle more TEU in total and allocate a larger fraction of that volume to inland flows at every sea-freight level.

The Δ columns in Table 7.9 highlight how the inland share reacts when sea demand is relaxed or tightened relative to the 100% case for each mix. When sea-freight demand is reduced to 80%, the inland share increases by about 4.7 percentage points in the 100:0 baseline and by roughly 4.9 and 5.1 points in the 50:50 and 25:75 mixes. The modular configurations therefore benefit slightly more from relaxed sea demand, on top of already having higher inland shares and volumes. When sea-freight demand is raised to 120%, the inland share drops by about 3.4–3.6 percentage points in all three configurations, which is a very small difference across mixes. The main effect is that the mixed systems start from a higher inland share and retain that advantage throughout the demand range, rather than a fundamentally different sensitivity to sea-side intensity.

Table 7.9: Inland throughput and share under $\pm 20\%$ sea-freight demand, per barge:module mix. Inland throughput is reported in $\times 10^5$ TEU. "Share" is the inland share of total port throughput; Δ gives the change in share (percentage) relative to the 100% sea-freight case of the same mix.

	100:0			50:50			25:75		
Demand	Throughput	Share	Δ	Throughput	Share	Δ	Throughput	Share	Δ
80%	2.32	24.4%	+4.7%	3.52	32.7%	+4.9%	3.97	35.4%	+5.1%
100%	2.20	19.7%	-	3.49	27.8%	-	3.96	30.3%	-
120%	2.10	16.3%	-3.4%	3.46	24.2%	-3.6%	3.96	26.7%	-3.6%

Together, Tables 7.8 and 7.9 show that the tested increases in sea-freight intensity are absorbed without loss of service level and without displacing inland work. Modular mixes unlock higher total and inland throughput at every demand level, while the sensitivity of port volume to sea-side demand remains almost unchanged.

7.2.2. Terminal KPIs

From a terminal perspective, changes in sea-freight demand primarily affect how intensively quay metres and crane capacity are used. Effective berth utilisation is defined as the share of open berth-metre-hours that is occupied by vessels eligible to use the berth. As sea-freight demand increases from 80% to 120%, effective utilisation at Terminal B rises in all three mixes (100:0, 50:50, and 25:75; Table 7.10.

For a given demand level, effective utilisation is highest in the conventional 100:0 case, lower in the 50:50 mix, and lowest in 25:75. This pattern reflects the shorter and more modular inland calls in the mixed cases, which achieve higher throughput with fewer berth-metre-hours. Within each mix, increasing sea-freight demand from 80% to 120% raises average utilisation, but the change is much larger for the mixed systems (about 12–14 percentage points) than for the conventional baseline (about 5 percentage points). Most of this uplift occurs between 100% and 120% sea demand, where additional deep-sea and feeder calls tighten the schedule and reduce some of the idle gaps along the quay.

Table 7.10: Effective berth utilisation at Terminal B across sea-freight demand levels. Effective berth utilisation is defined as the share of quay metre-hours used, based on effective vessel length (LOA plus buffers).

Mix	Demand	Average effective utilisation (%)	Peak effective utilisation (%)
100:0	80%	76.5	98.4
	100%	80.3	99.2
	120%	81.6	99.5
50:50	80%	57.8	98.4
	100%	64.3	99.2
	120%	70.3	99.4
25:75	80%	45.9	97.5
	100%	52.8	98.3
	120%	59.8	99.2

Peak values remain close to full use in all cases and mainly reflect short intervals of near-saturation. Terminal B therefore operates in a high-utilisation but non-saturated regime across the tested demand band.

Crane utilisation at Terminal B shows a similar picture of tighter but still controlled operations. Table 7.11 reports the installed, peak, and average number of cranes in use for the general and module pools, together with the total duty factor (average cranes in use divided by installed capacity). In the 100:0 case only the 16 general cranes are counted as active capacity; the module pool is unused and is therefore recorded as zero. At 80% sea demand, an average of 11.24 general cranes is active, rising to 12.29 at 120%, so the duty factor increases from about 0.70 to 0.77.

Table 7.11: Crane capacity and utilisation at Terminal B across sea-freight demand levels.

		General			Module			Total		
Mix	Demand	Inst.	Peak Used	Avg. Used	Inst.	Peak Used	Avg. Used	Inst.	Avg. Used	Avg. vs Inst. (%)
100:0	80%	16	16.00	11.24	0	0.00	0.00	16	11.24	70.25%
	100%	16	16.00	12.01	0	0.00	0.00	16	12.01	75.06%
	120%	16	16.00	12.29	0	0.00	0.00	16	12.29	76.81%
50:50	80%	16	15.96	7.99	2	2.00	1.05	18	9.05	50.28%
	100%	16	15.98	9.08	2	2.00	1.06	18	10.14	56.33%
	120%	16	16.00	10.16	2	2.00	1.05	18	11.21	62.28%
25:75	80%	16	15.81	6.13	2	2.00	1.58	18	7.72	42.89%
	100%	16	15.93	7.25	2	2.00	1.58	18	8.83	49.06%
	120%	16	15.99	8.40	2	2.00	1.58	18	9.98	55.44%

In the mixed systems, two dedicated module cranes are available, so total installed capacity at Terminal B rises to 18 units. Within each mix, however, the average use of the module cranes is almost constant across the 80–120% band (about 1.0 crane for 50:50 and 1.6 cranes for 25:75). The main demand effect sits in the general crane pool: for 50:50, average general-crane use increases from 7.99 to 10.16 as sea-freight demand rises from 80% to 120%; for 25:75 it increases from 6.13 to 8.40 over the same range. This confirms that higher sea-side demand primarily loads the general crane pool, while module cranes run at a relatively stable average level determined by the inland mix.

For a fixed demand level, the total duty factor is highest in the conventional 100:0 case, lower for 50:50, and lowest for 25:75, consistent with the berth-utilisation results. Between 100% and 120% sea demand, the duty factor increases more strongly in the mixed systems than in the baseline, mirroring the larger rise in effective berth utilisation as additional deep-sea and feeder calls fill residual gaps in the schedule. Peak usage of general and module cranes is close to the installed level in all cases, but average-to-installed ratios remain between about 0.43 and 0.77. This indicates that crane capacity is

used intensively but not exhausted, and that both 50:50 and 25:75 can absorb $\pm 20\%$ changes in sea-freight demand without requiring additional crane installations, while keeping the general crane pool lighter loaded than in the conventional system.

7.2.3. Vessel KPIs

Turnaround time is composed of waiting at anchorage, berth and handling time (including minor set-ups), and short intra-port sailing legs. Table 7.12 summarises average turnaround time by vessel class for the three demand levels and the three mixes considered (100:0, 50:50, and 25:75).

Table 7.12: Average turnaround time by vessel class, mix, and sea-freight demand level (hours). Turnaround time is measured from arrival at the port area to departure.

	100:0			50:50			25:75		
Vessel class	80%	100%	120%	80%	100%	120%	80%	100%	120%
Deep-sea	27.8	28.5	30.0	26.0	26.6	27.3	25.5	26.0	26.4
Feeder	14.6	15.4	16.8	13.2	13.7	14.2	12.8	13.1	13.4
Barge	24.5	33.7	40.6	12.1	15.6	20.4	9.3	10.5	12.5
Module (mother)	-	-	-	1.1	1.4	2.2	1.6	2.1	3.6

The changes are modest for deep-sea and feeder vessels and much larger for barges. Across all mixes, deep-sea and feeder turnaround times increase gradually with sea-freight demand, by about 2 hours between 80% and 120%. This indicates that ship-to-shore capacity and berth priority rules protect deep-sea and feeder services effectively, even when demand is 20% above the baseline.

Barge turnaround times show a much stronger dependence on demand and on the barge–module mix. In the purely conventional case (100:0), average barge turnaround increases from 24.5 hours at 80% demand to 40.6 hours at 120%, an increase of about 16 hours. In the 50:50 mix, the corresponding increase is from 12.1 to 20.4 hours, and in the 25:75 mix from 9.3 to 12.5 hours. At 120% demand, barge turnaround in 25:75 remains well below the 80% barge turnaround in the conventional case. This shows that modularisation substantially stabilises inland service levels against higher sea-side intensity.

The breakdown of barge turnaround time clarifies the underlying mechanism. For each mix, the average time at terminals stays close to 8 hours and is almost insensitive to demand, whereas waiting time at anchorage or dolphins increases sharply as sea-freight load rises. The demand effect on barge turnaround is therefore driven by additional off-berth waiting. Per-class anchorage shares and anchorage waiting times for all mixes and demand levels are reported in Appendix F, Table F.2. These results confirm that higher sea-freight demand primarily increases anchorage for barges and modules, while deep-sea and feeder anchorage remains limited.

Table 7.13: Average barge waiting time by mix and sea-freight demand level (hours).

Mix	80%	100%	120%
100:0	15.61	23.62	29.73
50:50	3.49	6.72	11.05
25:75	0.58	1.84	3.97

In the conventional 100:0 case, barge waiting time almost doubles between 80% and 120% demand (from about 15.6 to 29.7 hours), while handling time at terminals remains close to 7.5–8 hours. In the 50:50 mix, barge waiting also increases with demand but remains well below the conventional case at every demand level. In the 25:75 mix, barge waiting time stays low, rising from only about 0.6 to 4.0 hours across the demand range. The modular mixes therefore absorb higher sea-freight intensity by keeping inland waiting times short and using the additional module calls to fill berth pockets more evenly.

For module vessels, the mother calls retain short turnaround times at moderate demand (around 1–1.5 hours in the 50:50 and 25:75 mixes at 80–100%), but their turnaround increases at 120% demand,

especially in the 25:75 mix (up to roughly 3.6 hours). This uplift reflects more frequent staging and synchronisation of module convoys around busy periods at the quay. Overall, the demand experiments confirm that sea-freight load primarily feeds into waiting at anchorage and that higher module shares stabilise barge performance without materially affecting deep-sea and feeder turnaround.

7.2.4. Overview of Demand Case Results

Total port throughput scales almost linearly with sea-freight demand (Table 7.8). For any given demand level, both mixed systems handle substantially more TEU than the conventional baseline: at 100% demand, total throughput rises from 1.12 million TEU in the 100:0 case to 1.26 million TEU in the 50:50 mix and 1.31 million TEU in 25:75, with similar uplifts at 80% and 120%. Inland throughput shows an even stronger contrast. At 100% demand, the inland component increases from 2.20×10^5 TEU in 100:0 to 3.49×10^5 TEU in 50:50 and 3.96×10^5 TEU in 25:75, and the inland share of total port throughput rises from 19.7% to 27.8% and 30.3% respectively (Table 7.9). Modular uptake therefore raises both total and inland throughput at every demand level, while the elasticity of port volume with respect to sea-side demand remains almost unchanged.

From a vessel perspective, deep-sea and feeder services remain robust across the tested demand band. Average deep-sea turnaround stays in the range 26–30 hours and feeder turnaround between roughly 13 and 17 hours for all mixes and demand levels (Table 7.12), with only a gradual increase as sea-freight intensity rises. In contrast, barge turnaround time is highly sensitive to both demand and mix. In the conventional 100:0 system, average barge turnaround increases from 24.5 hours at 80% to 40.6 hours at 120%, an uplift of about 16 hours. The 50:50 and 25:75 mixes sharply reduce both the level and the sensitivity of barge turnaround: in 50:50 the corresponding increase is from 12.1 to 20.4 hours, while in 25:75 it is limited to a rise from 9.3 to 12.5 hours. At 120% demand, barges in 25:75 still turn faster than barges in the 100:0 baseline at 80%. The decomposition in Table 7.13 shows that these differences are driven by off-berth waiting rather than by handling time at the quay: average time at terminals remains close to 8 hours in all mixes, while waiting time more than doubles with demand in the 100:0 case and remains much lower in the mixed systems.

On the terminal side, both berth and crane indicators confirm a tighter but non-saturated operating regime. At Terminal B, effective berth utilisation increases with sea-freight demand in all three configurations, but average values remain well below full use and peak values represent short intervals of near-saturation rather than sustained overload (Table 7.10). For a fixed demand level, the conventional 100:0 case uses the berth most intensively, the 50:50 mix occupies fewer berth-metre-hours for the same sea-freight load, and 25:75 operates at the lowest average utilisation while handling the highest total throughput. Crane utilisation at Terminal B shows the same ranking. Average general-crane use and overall duty factors increase with demand but stay in the range 0.4–0.7 of the installed base, and the mixed systems shift part of the inland workload to the module crane pool (Table 7.11). This means that both 50:50 and 25:75 absorb $\pm 20\%$ changes in sea-freight demand without requiring additional berth length or crane installations, and that the higher-module mix does so with lower average loading of the general crane pool than the conventional system.

Overall, the demand experiments show that within the 80–120% range the port operates in a high-utilisation yet non-congested regime. Modular mixes retain the same responsiveness of port throughput to sea-freight demand as the conventional system, while delivering higher total and inland throughput, shorter and more stable barge turnaround times, and lower average loading of terminal resources. The 50:50 configuration achieves most of these gains with moderate use of module capacity, whereas 25:75 offers the strongest protection of inland performance and the highest inland share of throughput.

7.3. Configuration Analysis: Cranes, Handling Speed, and Module Capacity

Implementation choices determine whether the module concept converts potential into sustained performance. This section studies three configuration choices in the module system; the handling speed of the dedicated module crane pool, the number of module cranes installed per terminal, and the capacity of individual modules. Their effects are evaluated under two inland mixes, the balanced 50:50 case and the higher-module 25:75 case. Demand patterns, berth plans, and operating policies are kept fixed, so

that observed changes in throughput, turnaround time, waiting, berth occupancy, and crane utilisation can be attributed to these configuration choices. The aim is to identify when additional module capacity, faster handling, or extra module cranes cease to yield meaningful operational benefits, and to reveal the resources that become limiting in each setting.

7.3.1. Case: Crane Configuration

The number of module cranes is varied while all other settings remain unchanged. General quay cranes are denoted by QC and module cranes by MC. Relative to a reference configuration, five alternative crane layouts are considered. Two swap cases partially convert quay capacity into module capacity: configuration $-2\mathrm{QC}/+2\mathrm{MC}$ replaces two QC with two MC, and $-1\mathrm{QC}/+1\mathrm{MC}$ replaces one QC with one MC. Three addition cases increase module capacity on top of the existing quay–crane pool without removing QC, labelled $+1\mathrm{MC}$, $+2\mathrm{MC}$, and $\infty\mathrm{MC}$, where the latter approximates an effectively unlimited MC pool.

All configurations are evaluated against the baseline crane layout for the mixed systems and interpreted in the context of the conventional 100:0 system without module cranes discussed earlier. These scenarios allow a clean separation of effects. If quay access and the QC pool are governing, adding MC capacity should exhibit diminishing returns and congestion attribution will remain predominantly quay–driven. Conversely, if module handling is constraining, additional MC capacity should lift throughput and reduce waiting, with a measurable shift in congestion attribution away from the quay.

Vessel KPIs

The time spent in the port in the 25:75 barge–module mix is summarised in Tables 7.14 and 7.15.

Table 7.14: Turnaround time in hours for the 25:75 barge:module mix across crane configurations.

Vessel class	Baseline	-2QC/+2MC	-1QC/+1MC	+1MC	+2MC	+3MC	∞MC
Deepsea	28.51	26.64	26.26	25.90	25.95	25.96	25.92
Feeder	15.39	13.70	13.36	13.04	13.06	13.08	13.06
Barge	33.65	14.35	12.33	10.39	10.45	10.89	10.66
Module (mother)	-	2.41	157.01	155.35	2.14	1.35	1.09

Table 7.15: Time spent waiting in the port by module vessels for the 25:75 barge:module mix (hours).

Vessel class	-2QC/+2MC	-1QC/+1MC	+1MC	+2MC	+3MC	∞MC
Module (child)	0.68	143.70	143.20	0.60	0.19	0.07

For barges, turnaround time is highest in the $-2\mathrm{QC}/+2\mathrm{MC}$ case (about 14.35 h), where quay capacity is most constrained, and drops once QC capacity is restored. In the addition-type configurations (+1MC, +2MC, ∞MC), barge turnaround stabilises in a narrow band around 10.4-10.9 h. Deep-sea and feeder vessels are comparatively insensitive to the crane layout: their turnaround times remain between roughly 25.9-26.6 h (deep-sea) and 13.0-13.7 h (feeders) across all cases.

Modules are highly sensitive to effective MC capacity. In the $-1\mathrm{QC}/+1\mathrm{MC}$ and baseline layouts, average module-mother turnaround jumps above 150 h. Table 7.15 shows that this behaviour is almost entirely due to waiting: module vessels spend about 143 h waiting off-berth, while time at terminals and sailing remain consistent. In contrast, the $-2\mathrm{QC}/+2\mathrm{MC}$, +1MC, +2MC and ∞MC configurations keep module-mother turnaround between about 1 and 2.5 h, with waiting below 1 h.

In the 25:75 mix, the module-crane pool therefore needs sufficient MC capacity to avoid severe module-side queues. Once MC capacity is adequately dimensioned (the configurations with short module waiting), module turnaround returns to normal and barge turnaround remains low, while deep-sea and feeder performance is essentially unaffected. Removing QC to create MC ($-2\mathrm{QC}/+2\mathrm{MC}$) helps modules but penalises barges, whereas addition-only layouts keep both barge and module performance in a favourable range.

Table 7.16: Turnaround time in hours for the 50:50 barge–module mix

Vessel class	Baseline	$-2\mathrm{QC}/+2\mathrm{MC}$	$-1\mathrm{QC}/+1\mathrm{MC}$	+1MC	+2MC	+3MC
Deepsea	28.51	27.91	27.15	26.70	26.64	26.49
Feeder	15.39	14.86	14.04	13.68	13.70	13.58
Barge	33.65	26.78	18.84	15.10	15.60	15.39
Module (mother)	-	1.85	36.76	36.15	1.38	1.20

Table 7.17: Time spent waiting in the port by module vessels for the 50:50 barge:module mix (hours).

Vessel class	$-2\mathrm{QC}/+2\mathrm{MC}$	$-1\mathrm{QC}/+1\mathrm{MC}$	+1MC	+2MC	+3MC
Module (child)	0.39	26.59	26.50	0.20	0.12

In the 50:50 mix, barge turnaround shows a similar dependence on crane capacity. In the $-2\mathrm{QC}/+2\mathrm{MC}$ case, average barge turnaround rises to about 26.8 h, more than ten hours longer than in the configurations with additional module capacity. Restoring one quay crane in the $-1\mathrm{QC}/+1\mathrm{MC}$ configuration reduces barge turnaround to 18.8 h, and once module handling is expanded (+1MC, +2MC, +3MC) barge turnaround settles in the range 15.1–15.6 h. Deep-sea and feeder turnaround times again remain comparatively stable, varying by less than about 1.5 h across all crane settings.

Module in the 50:50 system also exhibit a clear threshold behaviour. With reduced or baseline effective MC capacity combined with tight general-quay capacity ($-1\mathrm{QC}/+1\mathrm{MC}$ and +1MC), average module–mother turnaround increases to around 36 h. The segment breakdown in Table 7.17 shows that this uplift is almost entirely due to waiting (about 26.5–26.6 h), while time at terminals stays near 0.55 h and sailing near 0.9–1.0 h. In the $-2\mathrm{QC}/+2\mathrm{MC}$, +2MC, and +3MC cases, module–mother turnaround remains close to 1–2 h and waiting is below 0.4 h.

Overall, the 50:50 configuration remains governed by quay and module-handling capacity. Removing QC substantially degrades barge performance, whereas adding MC reduces barge turnaround and stabilises module mothers by shortening module-side queues. Deep-sea and feeder turnaround times are largely protected by their priority and by available ship-to-shore capacity, and vary only modestly across the crane configurations.

Terminal KPIs

Changing the number of cranes also affects terminal operations. This subsection examines how different crane configurations influence berth utilisation, crane utilisation, and anchorage behaviour for the 25:75 and 50:50 barge–module mixes.

Table 7.18 summarises effective berth utilisation at Terminal B for the different crane configurations. Effective berth utilisation is defined as the share of quay metre-hours used, based on effective vessel length (LOA plus buffers), consistent with the rectangle-based measure introduced in Section 5.4.1.

In both mixes, average effective utilisation stays comfortably below full use, between roughly 50% and 54% for 25:75 and between 62% and 67% for 50:50, while peak values approach full utilisation only for short intervals. The variation across crane configurations is modest: in 25:75 the span is about 4 percentage points; in 50:50 it is around 4–5%. This confirms that, at fixed demand, changing the QC/MC configuration does not fundamentally shift the berth from a high-utilisation to a saturated regime. Instead, the berth remains moderately loaded in 25:75 and more heavily but still controllably loaded in 50:50, and the main differences between configurations arise from how inland traffic uses the available windows rather than from large swings in the length–time occupancy ratio.

Table 7.18: Effective berth utilisation at Terminal B for different crane configurations.

Configuration	Avg. effective utilisation (%)	Peak effective utilisation (%)
Baseline (0 MC)	80.2	99.2
	25:75 mix	
$-2\mathrm{QC}/+2\mathrm{MC}$	53.3	98.3
$-1\mathrm{QC}/+1\mathrm{MC}$	51.0	98.3
+1MC	49.9	98.6
+2MC	52.8	98.3
+3MC	53.9	99.0
∞MC	53.6	99.3
	50:50 mix	
$-2\mathrm{QC}/+2\mathrm{MC}$	63.1	99.1
$-1\mathrm{QC}/+1\mathrm{MC}$	65.4	98.8
+1MC	62.0	99.2
+2MC	64.3	99.2
+3MC	64.5	99.3
∞MC	66.5	99.3

Crane utilisation at Terminal B is summarised in Table 7.19, which reports installed, peak, and average numbers of general quay cranes (QC) and module cranes (MC) in use, together with the average total duty factor.

Table 7.19: Crane capacity and utilisation at Terminal B for different crane configurations.

	General (QC)			Module (MC)			Total		
Configuration	Inst.	Peak	Avg.	Inst.	Peak	Avg.	Inst.	Peak	Avg. vs inst. (%)
Baseline	16	16.00	12.01	0	0.00	0.00	16	16.00	75.06%
				25:75 mix					
$-2\mathrm{QC}/+2\mathrm{MC}$	14	14.00	7.23	2	2.00	1.58	16	16.00	55.06%
$-1\mathrm{QC}/+1\mathrm{MC}$	15	14.99	7.22	1	1.00	0.98	16	15.98	51.25%
+1MC	16	15.96	7.08	1	1.00	0.98	17	16.87	47.41%
+2MC	16	15.93	7.25	2	2.00	1.58	18	17.83	49.06%
+3MC	16	15.99	7.44	3	3.00	1.58	19	18.79	47.47%
∞MC	16	15.93	7.39	100	10.28	1.58	116	23.80	7.73%
				50:50 mix					
$-2\mathrm{QC}/+2\mathrm{MC}$	14.00	13.99	8.70	2	2.00	1.05	16.00	15.99	60.94%
$-1\mathrm{QC}/+1\mathrm{MC}$	15.28	15.27	9.21	1	1.00	0.97	16.28	16.27	62.53%
+1MC	15.97	15.96	8.74	1	1.00	0.98	16.97	16.94	57.28%
+2MC	16.00	15.98	9.08	2	2.00	1.06	18.00	17.94	56.33%
+3MC	16.00	16.00	9.12	3	3.00	1.05	19.00	18.95	53.53%
∞MC	16.00	16.00	9.42	100	8.35	1.05	116.00	23.34	9.03%

Three patterns stand out.

1. Quay capacity remains governing; In both mixes, peak use of general quay cranes (QC) is close to the installed level in every variant (for example, 14.00 of 14 in the 25:75 case with $-2\mathrm{QC}/+2\mathrm{MC}$ and about 16 of 16 in the higher-capacity cases). Average QC use is about 7.1–7.4 cranes in the 25:75 mix and about 8.7–9.4 cranes in 50:50, corresponding to duty factors of roughly 45–55% and 55–65% respectively. Increasing the number of MCs does not raise the peak load on QCs, which confirms that quay access and ship-to-shore capacity remain the governing resource.

2. Single versus multiple module cranes; In the 25:75 mix, a single MC operates close to full duty on average (about 0.98 cranes busy when one is installed), indicating a hard single-server bottleneck.

Once two MCs are available, the average number in use is about 1.6 across the configurations, which removes persistent queues at the module-crane pool and spreads the workload more evenly. With three MCs, the fleet average remains near 1.6, so per-crane utilisation falls to roughly 50–55%, leaving substantial spare capacity. In the 50:50 mix the module workload is lighter: average MC use is close to one crane in all settings, so adding a second or third MC mainly reduces per-crane utilisation rather than the total number of busy cranes.

3. ∞MC exposes the ceiling; Installing a very large number of MCs increases the total installed count sharply, but the peak number of cranes actually used rises only modestly (to about 23–24 across QC and MC). The average share of installed cranes in use falls to below 10% in both mixes (a denominator effect), which underlines that quay access and scheduling, rather than the MC count, set the operational ceiling at Terminal B.

Anchorage behaviour reflects how berth occupancy and crane availability interact in the scheduler. Table 7.20 and Table 7.21 summarise the share of barge and module calls routed via anchorage and the associated average waiting times.

Table 7.20: Anchorage behaviour for barges and modules in the 25:75 mix, with the 100:0 baseline shown for comparison.

	Barges		Modules	
Configuration	Share via anchorage (%)	Avg. wait (h)	Share via anchorage (%)	Avg. wait (h)
Baseline	31.5	8.4	N/A	N/A
$-2\mathrm{QC}/+2\mathrm{MC}$	12.9	1.8	63.6	0.7
$-1\mathrm{QC}/+1\mathrm{MC}$	9.4	1.2	100.0	143.8
+1MC	6.1	0.5	100.0	143.0
+2MC	6.3	0.5	63.5	0.6
+3MC	7.0	0.7	19.9	0.2
∞MC	6.7	0.6	2.0	0.1

Table 7.21: Anchorage behaviour for barges and modules in the 50:50 mix, with the 100:0 baseline shown for comparison.

	Barges		Modules	
Configuration	Share via anchorage (%)	Avg. wait (h)	Share via anchorage (%)	Avg. wait (h)
Baseline	31.5	8.4	N/A	N/A
$-2\mathrm{QC}/+2\mathrm{MC}$	24.9	5.7	28.3	0.4
$-1\mathrm{QC}/+1\mathrm{MC}$	18.4	3.2	99.8	26.8
+1MC	15.3	2.0	99.8	26.7
+2MC	14.2	2.1	27.7	0.2
+3MC	14.6	2.1	7.1	0.1
∞MC	15.0	2.3	2.3	0.1

Deep-sea and feeder services show only small changes in anchorage use across configurations (not reported here), which is consistent with their protected access through priorities and look-ahead. The main effects are on inland traffic, and these effects differ by mix.

In the 25:75 mix, barge anchorage remains modest in all configurations (about 6–13% of calls, with average waiting below 2 h), while modules absorb most of the additional staging. When only one module crane is installed ($-1\mathrm{QC}/+1\mathrm{MC}$ and +1MC), virtually all module calls are routed via anchorage and average module waiting times explode to more than 140 h. This corresponds to the very long module turnaround times reported in Table 7.14 and the high waiting component in Table 7.15. Once a second module crane is added (+2MC), module anchorage shares and waiting times collapse to levels similar to the $-2\mathrm{QC}/+2\mathrm{MC}$ case, and moving to +3MC or ∞MC progressively reduces both further. This shows that under the 25:75 workload a single MC creates a hard single-server bottleneck, while two MCs are sufficient to restore short module queues and keep barge anchorage low.

In the 50:50 mix, barge anchorage is higher overall (about 14–25% of calls) and more sensitive to general-quay capacity, but still only weakly affected by additional module cranes. The most pronounced

changes again appear in module behaviour. With one MC ($-1\mathrm{QC}/+1\mathrm{MC}$ and +1MC), almost all module calls are staged at anchorage and wait on average about 27 h before entering the berth. With two MCs (+2MC), the module anchorage share drops back to around 28% with negligible waiting, and with +3MC or ∞MC it falls below 10%. The berth and crane results in Tables 7.18 and 7.19 show that in these high-anchorage cases the quay remains below saturation but the module crane operates near full duty, so the anchorage spike is driven by module-handling scarcity rather than by a lack of berth-metre-hours.

Overall, the anchorage patterns confirm that higher anchorage shares in the crane experiment are driven by the interaction of quay access and module-crane availability. At fixed demand, berth utilisation at Terminal B stays in a high but non-saturated regime, while insufficient MCs create single-server queues that spill into anchorage. Providing at least two module cranes per terminal is therefore important to prevent excessive module staging and to keep barge anchorage at manageable levels, especially in the high-module 25:75 configuration.

Port-wide KPIs

Table 7.22 and Table 7.23 summarise total port throughput and the inland component for the crane configurations in the 50:50 and 25:75 mixes. For reference, the conventional 100:0 baseline processes about 1.12×10^6 TEU, of which 2.20×10^5 TEU (19.7%) are inland barge moves.

Table 7.22: Port throughput and inland share for the 50:50 barge–module mix under different QC/MC configurations.

	QC/MC configuration					
	-2QC/+2MC	-1QC/+1MC	+1MC	+2MC	+3MC	∞MC
Total throughput ($\times 10^6$ TEU)	1.25	1.24	1.25	1.26	1.26	1.26
Inland throughput ($\times 10^5$ TEU)	3.44	3.32	3.35	3.49	3.49	3.49
Inland share (%)	27.5	26.8	26.9	27.8	27.8	27.8

Table 7.23: Port throughput and inland share for the 25:75 barge–module mix under different QC/MC configurations.

	QC/MC configuration					
	-2QC/+2MC	-1QC/+1MC	+1MC	+2MC	+3MC	∞MC
Total throughput ($\times 10^6$ TEU)	1.30	1.18	1.18	1.31	1.31	1.31
Inland throughput ($\times 10^5$ TEU)	3.95	2.73	2.73	3.96	3.97	3.97
Inland share (%)	30.4	23.1	23.1	30.3	30.4	30.4

Across all realistic QC/MC configurations, both mixed systems deliver higher port throughput than the conventional 100:0 baseline. In the 50:50 mix, total throughput lies between 1.24 and 1.26×10^6 TEU, roughly 11–12% above the baseline. The 25:75 mix achieves about 1.31×10^6 TEU when at least two module cranes are installed, an uplift of roughly 16–17% relative to 100:0. The only exception is the 25:75 cases with a single MC ($-1\mathrm{QC}/+1\mathrm{MC}$ and +1MC), where throughput drops to about 1.18×10^6 TEU, consistent with the very long module waiting times reported in Section 7.3.1.

The inland component shows an even stronger contrast with the conventional system. In 50:50, inland throughput is around 3.3–3.5×10^5 TEU, corresponding to an inland share of 27–28% of total port moves, compared with 19.7% in 100:0. In the well-dimensioned 25:75 configurations (at least two MCs), inland throughput reaches about 4.0×10^5 TEU and the inland share stabilises near 30%. The single-MC 25:75 cases again stand out: inland throughput falls to about 2.7×10^5 TEU and the inland share to roughly 23%, showing that insufficient module-handling capacity suppresses inland performance.

The ∞MC variants do not produce additional gains in total or inland throughput compared with the +2MC and +3MC cases. This confirms that, once two or three module cranes are available, port-wide throughput is governed by quay access and general-crane capacity rather than by the exact size of the module-crane pool.

Overview of Crane Configuration Effects

The results across the 25:75 and 50:50 mixes indicate a system that is governed primarily by quay access, with a clear threshold in module-handling capacity:

- Binding resource. Peak use of general quay cranes (QC) at Terminal B is close to the installed level in all variants, while average duty factors remain in the ranges reported in Table 7.19 (about 45–55% in 25:75 and 55–65% in 50:50). Changing the MC configuration does not raise the peak load on QCs. Together with the berth utilisation levels in Table 7.18, this confirms a high-utilisation but non-saturated regime that is governed by quay access and ship-to-shore capacity rather than by the size of the module-crane pool.

- Threshold at two MCs in 25:75. In the 25:75 mix a single MC acts as a hard single-server bottleneck. Module mothers experience extreme turnaround times, on the order of 150–160 h in the $-1\mathrm{QC}/+1\mathrm{MC}$ and +1MC cases, almost entirely due to waiting (Tables 7.14 and 7.15). Module anchorage shares reach 100% with average waits of roughly 144 h (Table 7.20), and inland throughput and share are depressed to about 2.7×10^5 TEU and 23% (Table 7.23). Once two MCs are available (+2MC and beyond), module queues collapse, module and barge anchorage return to low levels, and inland throughput recovers to about 4.0×10^5 TEU with a 30% inland share. A third MC mainly creates spare capacity and yields only marginal additional gains.

- Balanced-mix behaviour in 50:50. In the 50:50 mix the module workload per crane is lower. Removing QCs strongly affects barges (average turnaround rising to about 27 h in the $-2\mathrm{QC}/+2\mathrm{MC}$ case), whereas adding MCs reduces barge turnaround to about 15–16 h and stabilises module mothers around 1–2 h, see Tables 7.16 and 7.17. Module waiting times spike to about 27 h only in the single-MC cases ($-1\mathrm{QC}/+1\mathrm{MC}$ and +1MC), with almost all modules routed via anchorage (Table 7.21). Deep-sea and feeder turnaround remains comparatively stable across configurations and stays within a narrow band of about 26–27 h and 13–15 h, confirming that their service is effectively protected.

- Anchorage as a symptom of module scarcity. For both mixes, deep-sea and feeder anchorage remains limited, while inland anchorage responds strongly to module-capacity shortages. In 25:75, barges are kept largely away from anchorage and modules absorb most of the staging; in 50:50, barge anchorage is higher overall but still small relative to the extreme module anchorage in the single-MC cases. The anchorage patterns in Tables 7.20 and 7.21 mirror the module waiting times and show that high anchorage is driven by single-server effects in the module-crane pool rather than by a lack of total quay length.

- Berth occupancy versus throughput. Effective berth utilisation at Terminal B varies only modestly across QC/MC configurations in each mix (about 50–54% in 25:75 and 62–67% in 50:50; Table 7.18). At fixed demand, crane reconfigurations therefore change how calls are packed into the available windows rather than pushing the berth into genuine saturation. This confirms that berth occupancy alone is not a reliable proxy for throughput in these experiments; short module calls can increase moves per occupied hour without large changes in the length–time occupancy ratio.

- Throughput linkage and diminishing returns. Both modular mixes deliver higher port throughput than the conventional 100:0 baseline for all realistic QC/MC configurations. In 50:50, total throughput is about 1.24–1.26×10^6 TEU with an inland share of 27–28% (Table 7.22). In well-dimensioned 25:75 configurations with at least two MCs, total throughput reaches about 1.30–1.31×10^6 TEU and the inland share about 30% (Table 7.23). The single-MC 25:75 cases are the clear outliers, with lower total and inland throughput. The ∞MC stress tests produce no further port-wide gains compared with the +2MC and +3MC cases and simply lower average utilisation as a share of installed capacity, which reinforces that quay access and scheduling set the ceiling once module scarcity is removed.

For the high-module 25:75 configuration, at least two MCs per terminal are required to avoid persistent single-server queues and excessive module anchorage. In the balanced 50:50 mix, one or two MCs already suffice under the tested demand, and additional cranes mainly provide robustness rather than higher throughput. Beyond these thresholds, diminishing returns are observed, since berth access and general QC capacity, rather than the MC count, become the governing constraints.

7.3.2. Case: Handling Speed

This section analyses how the speed of the dedicated module crane pool affects system performance for two modular uptake levels. Starting from the 50:50 and 25:75 barge:module mixes, the module handling time in the dedicated crane pool is varied between 1.5 and 3.5 min/container (runs sp1.5 to

sp3.5), interpreted as the average time to handle a single container move, while the vessel mix, demand pattern and number of module cranes are kept constant. The 100:0 case serves as a conventional barge baseline.

Vessel KPIs

Vessel level results show that changes in module crane speed mainly play out inside the modular system, while conventional and seagoing services remain comparatively robust. Tables 7.24, and 7.25 summarise the effects.

Introducing modules already transforms barge performance: in both the 50:50 and 25:75 mixes, barge turnaround times drop well below the conventional 100:0 baseline. Once this modular routing is in place, varying the module crane speed has only a secondary effect on barges. Across the sp1.5 to sp3.5 range, barge turnaround in each mix stays within a relatively narrow band around its new, lower level. This indicates that the main barge improvement comes from the presence of modular services and the changed berth allocation logic, not from the exact speed of the module crane pool.

Deepsea and feeder vessels are even less sensitive. Their turnaround times remain close to the baseline values across all speed settings and both mixes. Priorities and look ahead ensure that seagoing calls continue to see similar access to quay and cranes, so the dedicated module crane pool can be slowed down considerably before it has any noticeable effect on deepsea and feeder turnaround.

The modular services themselves behave very differently. Module mothers and children react in a strongly nonlinear way once the crane pool becomes heavily loaded, especially in the 25:75 mix where most inland demand is routed via modules. For fast handling (around 1.5 min/container), module mothers complete a visit in roughly the same order of magnitude as a short call: terminal time plus one shuttle cycle. As handling slows, the system initially absorbs the change through slightly longer dwell. Beyond a certain point, however, the dedicated crane pool saturates and a classic single server effect emerges. In the 50:50 mix this appears first as a gradual lengthening of module waiting times, then as a sharp jump in mother turnaround at the slowest speeds. In the 25:75 mix the transition is much steeper: once handling moves beyond roughly 2.0 to 2.5 min/container, module waiting explodes into tens and then hundreds of hours.

The breakdown of module child time components confirms the mechanism. In both mixes, terminal and sailing times remain close to one hour over the entire speed range. Almost all additional time is accumulated as waiting in front of the module cranes. The speed experiment therefore does not change what a module does while it is alongside or sailing, but it changes how often modules must queue for access to the limited number of module cranes. When utilisation of that pool approaches one, even modest further slowdowns cause very large increases in modular waiting, while barge and seagoing turnaround times remain comparatively stable.

Table 7.24: Average turnaround time by vessel class for different handling speeds with barge:module ratios 50:50 and 25:75 (hours).

		Handling speed scenario				
Vessel class	Baseline	sp1.5	sp2	sp2.5	sp3	sp3.5
		Ratio 50:50				
Deepsea	28.82	26.71	26.74	26.51	26.85	26.53
Feeder	15.70	13.78	13.70	13.72	13.70	13.60
Barge	33.64	16.21	15.02	16.32	15.97	14.74
Module (mother)	-	1.38	2.16	3.00	7.56	37.79
		Ratio 25:75				
Deepsea	28.82	25.97	25.88	25.87	26.12	25.96
Feeder	15.70	13.08	13.05	13.06	13.14	13.10
Barge	33.64	10.56	10.56	10.76	11.21	10.97
Module (mother)	-	2.10	13.09	73.02	118.11	152.51

Table 7.25: Average time spent on different segments in the port by module children (hours).

Mix	Segment	sp1.5	sp2.0	sp2.5	sp3.0	sp3.5
50:50	Terminals	0.55	0.74	0.92	1.10	1.29
	Waiting	0.20	0.48	0.85	2.87	27.57
25:75	Terminals	0.55	0.73	0.92	1.10	1.29
	Waiting	0.60	7.23	65.02	109.98	144.16

Terminal KPIs

Terminal level indicators show that the handling speed experiment does not push Terminal B into a fundamentally different operating regime. Effective berth utilisation adjusts to the changed module handling, but remains below the levels observed in the conventional system.

Table 7.26 reports average and peak effective berth utilisation at Terminal B for the baseline and for different module handling speeds in the 25:75 and 50:50 mixes. In both mixes, modular uptake reduces average effective utilisation compared with the 100:0 baseline. Shorter, more frequent inland calls and the redistribution of work across terminals mean that the same or higher throughput is achieved with fewer metre-hours in use. Within each mix, changing the module crane speed only nudges these utilisation indicators by a few percentage points. Peak utilisation still reaches values close to full use in all runs, so the terminal experiences short periods of near full loading, but on average it operates with more slack than in the conventional barge system.

Table 7.26: Average and peak effective berth utilisation at Terminal B for different module handling speeds.

Scenario	Avg. eff. berth util. (%)	Peak eff. berth util. (%)
Baseline	80.3	99.2
	25:75 mix	
1.5 min/container	52.8	98.3
2.0 min/container	54.1	98.9
2.5 min/container	55.3	98.6
3.0 min/container	53.0	98.0
3.5 min/container	52.5	98.5
	50:50 mix	
1.5 min/container	64.3	99.2
2.0 min/container	63.4	99.3
2.5 min/container	65.6	99.3
3.0 min/container	67.2	99.1
3.5 min/container	64.0	98.8

Crane utilisation at Terminal B follows a similar pattern (Table 7.27). Introducing modules lowers the average use of the general quay crane pool compared with the 100:0 baseline and shifts part of the workload to the dedicated module cranes. As module handling slows, the module cranes become busier, while the general pool adjusts only moderately. The last column shows average total crane use as a share of installed capacity; this remains below the baseline level in all modular runs, especially in the 25:75 mix where general crane use is clearly lower than in the conventional system. From the perspective of this representative terminal, slower module handling therefore manifests as a rebalancing of work between crane types rather than as a clear move into a crane-saturated regime.

The bottleneck counters in Table 7.28 confirm that most of the time Terminal B still operates in a slack or lightly loaded state. Slower module handling moves a small fraction of periods into a crane limited state, particularly in the 50:50 mix, but the share of time where cranes bind remains low and quay limitations are negligible. The very large increases in modular turnaround times at high modular shares are therefore not explained by persistent, local congestion at Terminal B alone. They arise from the cumulative effect of a heavily loaded module crane pool across the network, which generates long queues and anchorage staging for modules even though individual terminals, viewed in isolation, still

Table 7.27: Crane capacity and utilisation at Terminal B for the baseline and module handling speed variants. General cranes handle all barge and seagoing work; module cranes are dedicated to modules.

	General (QC)			Module (MC)			Total			
Scenario	Inst.	Peak Used	Avg. Used	Inst.	Peak Used	Avg. Used	Inst.	Peak Used	Avg. Used	Avg. vs Inst. (%)
100:0 baseline	16	16.00	12.01	2	0.00	0.00	16	16.00	12.01	75.04%
25:75 mix										
1.5 min/container	16	15.93	7.25	2	2.00	1.58	18	17.83	8.83	49.06%
2.0 min/container	16	15.93	7.35	2	2.00	1.94	18	17.80	9.29	51.62%
2.5 min/container	16	15.94	7.58	2	2.00	1.95	18	17.84	9.53	52.94%
3.0 min/container	16	15.91	7.26	2	2.00	1.96	18	17.78	9.23	51.26%
3.5 min/container	16	15.93	7.21	2	2.00	1.96	18	17.80	9.17	50.92%
50:50 mix										
1.5 min/container	16	15.98	9.08	2	2.00	1.06	18	17.94	10.14	56.33%
2.0 min/container	16	16.00	8.88	2	2.00	1.31	18	17.97	10.20	56.65%
2.5 min/container	16	15.99	9.17	2	2.00	1.58	18	17.99	10.75	59.71%
3.0 min/container	16	16.00	9.35	2	2.00	1.82	18	17.97	11.16	62.03%
3.5 min/container	16	15.98	8.81	2	2.00	1.96	18	17.95	10.77	59.81%

appear reasonably well dimensioned.

Table 7.28: Share of time limited by cranes, policy or slack at Terminal B for the baseline and module handling speed variants.

Scenario	Crane limited	Policy limited	Slack or no queue
100:0 baseline	0.00%	0.38%	99.62%
50:50 mix			
1.5 min/container	3.78%	0.32%	95.90%
2.0 min/container	4.79%	0.31%	94.89%
2.5 min/container	8.10%	0.31%	91.60%
3.0 min/container	10.46%	0.31%	89.24%
3.5 min/container	8.57%	0.32%	91.12%
25:75 mix			
1.5 min/container	1.85%	0.18%	97.97%
2.0 min/container	3.24%	0.18%	96.59%
2.5 min/container	3.46%	0.18%	96.37%
3.0 min/container	3.13%	0.18%	96.69%
3.5 min/container	2.85%	0.18%	96.98%

Port-wide KPIs

Port wide indicators summarise how changes in module crane speed propagate from local waiting times into total throughput and the balance between inland and seagoing flows. Table 7.29 reports total processed TEU, inland throughput and the inland share for the baseline and the speed variants.

The conventional 100:0 system provides the reference point. Introducing modules in a 50:50 mix already lifts total throughput and inland share substantially. Within this moderate modular regime, the port is quite forgiving to slower module handling. Across the sp1.5 to sp3.5 range, total throughput and inland throughput move only modestly and the inland share remains close to a stable level of about 28%. From a system perspective, the module crane pool mainly affects the level of service experienced by modules, not the number of inland units that can be processed.

At high modular uptake (25:75), the picture changes. With fast module handling, the port achieves both a higher total throughput and a much larger inland share than in the conventional baseline. As module handling slows, throughput gains are gradually eroded. Once the crane pool is slow enough

that modular queues dominate, total throughput starts to drop back towards the baseline and the inland share falls noticeably. This deterioration is almost entirely driven by reduced module volumes; barge, deepsea and feeder flows stay close to their respective baselines because their arrival patterns and priorities are unchanged. In other words, slow module cranes do not push seagoing traffic out of the system, but they do prevent the modular inland concept from realising its potential throughput.

Table 7.29: Throughput and inland share for the baseline and module crane pool speed variants. Total throughput is reported in $\times 10^6$ TEU, inland throughput (barge + module) in $\times 10^5$ TEU.

Scenario	Total TEUs ($\times 10^6$)	Change vs 100:0	Inland TEUs ($\times 10^5$)	Inland share
100:0 baseline	1.12	-	2.20	19.7%
		25:75 mix		
1.5 min/container	1.31	+16.6%	3.96	30.3%
2.0 min/container	1.30	+16.0%	3.92	30.2%
2.5 min/container	1.25	+11.4%	3.41	27.3%
3.0 min/container	1.21	+8.3%	3.02	24.9%
3.5 min/container	1.18	+5.3%	2.72	23.1%
		50:50 mix		
1.5 min/container	1.26	+12.3%	3.49	27.8%
2.0 min/container	1.25	+12.1%	3.48	27.8%
2.5 min/container	1.25	+11.6%	3.49	27.9%
3.0 min/container	1.26	+12.3%	3.48	27.7%
3.5 min/container	1.24	+10.6%	3.35	27.0%

Overview of Handling Speed Case Results

Across vessel, terminal and port levels, the handling speed experiment highlights three main dynamics.

First, modular uptake is the primary driver of barge and inland performance, while module crane speed is a secondary tuning parameter. Moving from a conventional 100:0 system to a mixed fleet roughly halves barge turnaround in the 50:50 mix and cuts it to about one third in the 25:75 mix. These gains remain largely intact when module handling is slowed within the tested range. Deepsea and feeder services are even more robust, with turnaround times and volumes staying close to their baseline levels.

Second, the dedicated module crane pool introduces a distinct single server sensitivity for the modular traffic. As long as the pool operates comfortably below full utilisation, changes in handling speed mainly translate into small adjustments in module dwell. Once the pool approaches saturation, further slowdowns trigger a strong, nonlinear increase in module waiting and mother turnaround, especially in the 25:75 mix where most inland demand is modular. The time breakdowns show that almost all of this additional delay is pure waiting at the crane, not longer service or sailing.

Third, the terminal and port level indicators suggest that these effects arise from how modular work is packed around existing sea freight, rather than from a general lack of quay or crane capacity. Terminal B operates with lower average berth occupancy and lower total crane use than in the conventional system, even for the slowest module speeds. The share of crane limited periods increases, but slack still dominates the time distribution. At port level, a moderate 50:50 modular uptake delivers robust throughput gains and a higher inland share across all speed settings, whereas very slow module handling in the 25:75 mix progressively erodes the modular throughput advantage.

Overall, the speed experiment suggests that moderate modular adoption is relatively forgiving to slower module handling: the port still benefits from additional inland windows and reduced barge waiting, even if module cranes are not extremely fast. High modular shares, by contrast, require a sufficiently fast and well dimensioned module crane pool to avoid severe queueing in the modular subsystem and to sustain the intended uplift in inland throughput.

Table 7.30: Qualitative summary of handling speed effects by modular mix.

Mix and speed	Modular system	Barge and seagoing traffic
50:50, fast handling (1.5 to 2.0 min/cont.)	Short module queues, mother calls close to a shuttle cycle; module waiting remains limited	Barge turnaround roughly halves vs 100:0; deepsea and feeder times remain near baseline
50:50, slow handling (2.5 to 3.5 min/cont.)	Module crane pool becomes heavily used; waiting grows, with a sharp spike only at the slowest speed	Barge and seagoing turnaround remain within a narrow band; throughput and inland share stay stable
25:75, fast handling (1.5 to 2.0 min/cont.)	High modular throughput with short queues; module calls fit into residual windows	Strong throughput gains and inland share around 30%; barge times remain low
25:75, slow handling (2.5 to 3.5 min/cont.)	Module crane pool saturates; module queues and mother turnaround grow to very high values	Total and inland throughput drop back towards the 100:0 baseline; deepsea and feeder flows are largely unchanged

7.3.3. Case: Module Capacity

This experiment varies the module capacity across 12, 24, 36 and 48 TEU for the 50:50 and 25:75 traffic mixes. To keep the annual module workload constant, the number of module missions is adjusted when capacity changes, and the module length is updated accordingly. The exact scaling of mission counts and vessel dimensions is given in Appendix D.9.

Vessel KPIs

Table 7.31: Average turnaround time in hours for the 25:75 barge–module mix.

Vessel class	12 TEU	24 TEU	36 TEU	48 TEU
Deep-sea	25.92	25.95	26.00	25.92
Feeder	13.11	13.06	13.12	13.07
Barge	11.18	10.45	11.28	10.34
Module (mother)	9.52	2.14	2.16	6.94

Table 7.32: Average time at terminals and in waiting for modules (children) in the 25:75 barge–module mix (hours).

Segment	12 TEU	24 TEU	36 TEU	48 TEU
Terminals	0.55	0.55	0.80	1.05
Waiting	0.20	0.60	0.49	2.73

With annual module workload held constant, changing module capacity in the 25:75 mix leaves the other vessel classes largely unchanged. Barge, deep-sea and feeder turnaround in Table 7.31 varies by less than about one hour across the four capacities and shows no clear monotonic trend, which suggests no systematic spillover from the module configuration to these classes.

By contrast, module-mother turnaround shows a pronounced non-linear response. At 12 TEU, average module-mother turnaround is about 9.5 h, indicating that very small modules and the associated high mission frequency create substantial module-side pressure. Increasing capacity to 24 TEU and 36 TEU sharply reduces module-mother turnaround to about 2.1–2.2 h. When capacity is pushed further to 48 TEU, the average module-mother turnaround rises again to almost 7 h.

The segment breakdown for module children in Table 7.32 clarifies the mechanism. As capacity increases, terminal dwell grows steadily from about 0.55 h at 12–24 TEU to 0.80 h at 36 TEU and 1.05 h at 48 TEU, reflecting more handling work per call. Waiting remains modest at intermediate capacities (around 0.5–0.6 h at 24–36 TEU) but jumps to roughly 2.7 h at 48 TEU, indicating more frequent staging and queueing for suitable windows and crane slots when very large module moves arrive in fewer, heavier pulses. Taken together, these results suggest that under a high module share there is a “sweet spot’’ around 24–36 TEU: very small modules lead to many short but intrusive visits, while very large modules reduce call frequency but create long waits and heavier bursts of work.

Table 7.33: Average turnaround time in hours for the 50:50 barge–module mix.

Vessel class	12 TEU	24 TEU	36 TEU	48 TEU
Deep-sea	26.64	26.64	26.66	26.68
Feeder	13.65	13.70	13.57	13.75
Barge	15.35	15.60	14.73	15.47
Module (mother)	1.23	1.38	1.80	2.32

Table 7.34: Average time at terminals and in waiting for modules (children) in the 50:50 barge–module mix (hours).

Segment	12 TEU	24 TEU	36 TEU	48 TEU
Terminals	0.30	0.55	0.80	1.05
Waiting	0.27	0.20	0.28	0.38

Under the balanced 50:50 mix, the behaviour is smoother but qualitatively similar. Barge, deep-sea and feeder turnaround times in Table 7.33 remain within a narrow band across capacities and show no systematic trend; differences are on the order of a few tenths of an hour. Module-mother turnaround, however, increases monotonically with module size, from about 1.2 h at 12 TEU to 1.4 h at 24 TEU, 1.8 h at 36 TEU and roughly 2.3 h at 48 TEU.

The module-child decomposition in Table 7.34 shows that the additional time is mainly taken up at the interface rather than in transit. Terminal dwell rises from about 0.30 h at 12 TEU to 0.55 h, 0.80 h and 1.05 h as capacity increases, while waiting remains relatively low, between roughly 0.2 h and 0.4 h, with only a modest uptick at 48 TEU. Fewer, larger module calls therefore concentrate more handling per visit and slightly extend average module turnaround without materially affecting other vessel classes at the current load.

Comparing the mixes, module mothers complete much faster in the 50:50 case than in 25:75 at the same capacity, particularly at the extremes. At 12 TEU the difference is sizeable (about 9.5 h in 25:75 versus 1.2 h in 50:50), and at 48 TEU module-mother turnaround is roughly 6.9 h in 25:75 against 2.3 h in 50:50. At intermediate capacities of 24–36 TEU, the gap narrows and both mixes keep module-mother turnaround near 2 h. Since the annual module TEU is held constant by construction, these patterns indicate that capacity choices primarily trade call frequency against per-call handling and waiting, with the high-module 25:75 configuration more sensitive to very small and very large capacities.

In practical terms, the results suggest that module capacity should be chosen jointly with the traffic mix and crane configuration. Larger modules simplify scheduling in the sense of fewer calls, but at high module shares this can come at the cost of longer waits for suitable quay windows and module-crane capacity, as seen in the 48 TEU case for 25:75. Mid-range capacities around 24–36 TEU offer a compromise: they keep module turnaround short in both mixes, maintain stable barge, deep-sea and feeder performance, and avoid the extreme waiting effects observed for very small or very large module sizes.

Terminal KPIs

This section shifts from vessel turnaround to terminal-facing indicators. It examines anchorage, berth fit utilisation and crane utilisation under constant annual module TEU, comparing capacities of 12, 24, 36 and 48 TEU across the 25:75 and 50:50 mixes to identify whether observed changes stem from access to the quay, from packing of berth metres, or from crane availability. Terminal B is used as a representative example because its installed crane counts are typical within the terminal portfolio.

Anchorage behaviour is summarised in Tables 7.35 and 7.36, which report, for barges and modules, the share of calls routed via anchorage and the average anchorage waiting time by module capacity.

Deep-sea and feeder services show only minor changes in anchorage use across capacities. In the 25:75 mix, deep-sea anchorage remains around 8 to 9 % of calls and feeders around 3 to 5 %; in the 50:50 mix the corresponding ranges are roughly 11 to 12 % for deep-sea and 4 to 5 % for feeders, with

Table 7.35: Anchorage behaviour for barges and modules under different module capacities (25:75 mix).

Capacity	Barges		Modules	
	Share via anchorage (%)	Avg. wait (h)	Share via anchorage (%)	Avg. wait (h)
12 TEU	7.2	0.78	97.2	4.92
24 TEU	6.4	0.54	63.5	0.57
36 TEU	7.1	0.81	51.5	0.46
48 TEU	6.9	0.53	86.1	2.71

Table 7.36: Anchorage behaviour for barges and modules under different module capacities (50:50 mix).

Capacity	Barges		Modules	
	Share via anchorage (%)	Avg. wait (h)	Share via anchorage (%)	Avg. wait (h)
12 TEU	14.5	2.10	46.1	0.26
24 TEU	14.2	2.14	27.7	0.20
36 TEU	14.1	1.91	20.6	0.25
48 TEU	14.6	2.13	34.4	0.38

average waiting times near 0.3 h for deep-sea and below 0.2 h for feeders. This is consistent with their protected access through priorities and look-ahead.

The main effects appear in inland traffic. For barges, anchorage shares remain modest and fairly stable. In 25:75, between 6 and 7 % of barge calls are routed via anchorage with average waiting below 1 h for all capacities; in 50:50, barge anchorage stays near 14 to 15 % with average waiting about 2 h.

Module calls are far more sensitive to capacity and drive most of the variation in inland anchorage. In the high-module 25:75 mix, nearly all module calls are staged at anchorage at 12 TEU, with average waiting close to 5 h. Increasing capacity to 24 or 36 TEU sharply reduces module anchorage shares to about 64 and 52 %, and brings average module waiting down to roughly 0.5 h. At 48 TEU, module anchorage increases again to more than 86 % of calls and waiting rises to about 2.7 h, indicating that very large modules create longer, less frequent windows that require more staging. In the balanced 50:50 mix, module anchorage levels are lower overall but show a similar pattern: shares fall from about 46 % at 12 TEU to 28 and 21 % at 24 and 36 TEU, with short waiting times below 0.3 h, before rising again to roughly 34 % and 0.4 h at 48 TEU.

These patterns suggest that, under constant annual module TEU, capacities around 24 to 36 TEU offer a compromise between call frequency and per-call duration. Smaller modules generate many short visits and frequent pulses that must be staged, while very large modules reduce call counts but create long visits that block windows and trigger additional anchorage. Intermediate capacities reduce both call frequency and per-visit blocking and therefore produce the lowest module anchorage and the shortest module waiting times. Barge anchorage remains largely governed by the interaction of berth access rules and sea-side protection, and is only weakly affected by the module capacity choice.

Effective berth utilisation at Terminal B is summarised in Table 7.37. Effective berth utilisation is defined as the share of quay metre-hours used, based on effective vessel length (LOA plus buffers), consistent with the rectangle-based measure introduced in Section 5.4.1.

In both mixes, average effective utilisation at Terminal B remains on a stable plateau as capacity changes. In the 25:75 mix, the average ranges from about 53 to 54 %, with peak utilisation close to full use for short intervals in all cases. In the 50:50 mix, the berth operates in a higher but still controlled regime, with average effective utilisation around 64 to 66 % and peak utilisation again near 99 %. The variation across capacities is modest, on the order of one to two percentage points.

These results indicate that, at this terminal, changing module capacity under constant annual module TEU does not move the berth from a moderately loaded to a saturated state. Differences between capacities arise more from how inland calls are scheduled and packed into the available windows than from large shifts in the underlying length–time occupancy ratio. Capacities around 36 TEU show slightly

Table 7.37: Effective berth utilisation at Terminal B for different module capacities.

Capacity	Avg. effective utilisation (%)	Peak effective utilisation (%)
25:75 mix		
12 TEU	53.7	98.7
24 TEU	52.8	98.3
36 TEU	53.3	99.3
48 TEU	54.4	98.8
50:50 mix		
12 TEU	64.3	99.3
24 TEU	64.3	99.2
36 TEU	66.1	99.3
48 TEU	64.3	99.3

higher average effective utilisation, consistent with the improved anchorage and packing behaviour observed at intermediate capacities, but the effect size at this single terminal remains limited.

Crane utilisation and congestion attribution at Terminal B are reported in Table 7.38. The table shows installed, peak and average numbers of general and module cranes in use and the average duty factor relative to installed crane capacity.

Table 7.38: Crane capacity and utilisation at Terminal B for different module capacities.

	General			Module			Total	
Capacity	Inst.	Peak	Avg.	Inst.	Peak	Avg.	Peak vs inst. (%)	Avg. vs inst. (%)
25:75 mix								
12 TEU	16	15.98	7.24	2	2	1.93	99.4	50.9
24 TEU	16	15.93	7.25	2	2	1.58	99.1	49.1
36 TEU	16	16.00	7.35	2	2	1.44	99.3	48.9
48 TEU	16	15.97	7.28	2	2	1.83	98.7	50.6
50:50 mix								
12 TEU	16	15.99	8.96	2	2	1.33	99.9	57.1
24 TEU	16	15.98	9.08	2	2	1.06	99.7	56.3
36 TEU	16	15.99	9.39	2	2	0.96	99.8	57.5
48 TEU	16	16.00	8.98	2	2	1.22	99.8	56.7

Across all capacities and both mixes, peak use of general cranes is essentially at the installed level (about 16 cranes) and average use is moderate: roughly 7.2 to 7.4 general cranes busy in the 25:75 mix and about 9 to 9.4 in the 50:50 mix. Module cranes operate as a small but active pool. In 25:75, average module-crane use ranges between about 1.4 and 1.9 of two cranes; in 50:50 it lies between roughly 1.0 and 1.3. The average share of installed cranes in use remains around 49 to 51 % in 25:75 and 56 to 58 % in 50:50, with peak total use near the installed count in all cases.

Congestion attribution metrics confirm that cranes are rarely binding. At Terminal B, the fraction of time classified as crane-limited remains between about 0.02 and 0.05 % across all capacities and mixes, with no periods where the quay itself is binding and the remaining time almost entirely in slack or no-queue states. Policy-related limitations account for only about 0.002 to 0.003 % of time.

Taken together, the terminal-level indicators show that changes in module capacity primarily reallocate crane hours and modify how module calls are staged at anchorage, without creating sustained crane scarcity. General-crane loading and effective berth utilisation at Terminal B are comparatively insensitive to capacity under constant annual module TEU. Capacity choices should therefore be guided by berth-fit considerations and by the trade-off between module call frequency and per-call duration,

rather than by expectations of relieving crane constraints at the terminal.

Port-wide KPIs

Tables 7.39 and 7.40 summarise total port throughput and the inland component for the module-capacity experiments in the 50:50 and 25:75 mixes. Total throughput is reported in $\times 10^6$ TEU and inland throughput (barges plus modules) in $\times 10^5$ TEU. For reference, the conventional 100:0 baseline processes about 1.12×10^6 TEU, of which 2.20×10^5 TEU (19.7%) are inland barge moves.

Table 7.39: Port throughput and inland share for the 50:50 barge–module mix under different module capacities.

Capacity	Total throughput ($\times 10^6$ TEU)	Inland throughput ($\times 10^5$ TEU)	Inland share (%)
12 TEU	1.26	3.50	27.8
24 TEU	1.26	3.49	27.8
36 TEU	1.26	3.47	27.6
48 TEU	1.33	4.21	31.6

Table 7.40: Port throughput and inland share for the 25:75 barge–module mix under different module capacities.

Capacity	Total throughput ($\times 10^6$ TEU)	Inland throughput ($\times 10^5$ TEU)	Inland share (%)
12 TEU	1.30	3.92	30.2
24 TEU	1.31	3.96	30.3
36 TEU	1.30	3.94	30.3
48 TEU	1.41	5.03	35.6

Across all tested capacities, both mixed systems deliver higher port throughput than the conventional 100:0 baseline. In the 50:50 mix, total throughput remains close to 1.26×10^6 TEU for capacities between 12 and 36 TEU, roughly 11–12% above the baseline, while the inland component is about 3.5×10^5 TEU, corresponding to an inland share of 27–28%. The 25:75 mix achieves about 1.30–1.31×10^6 TEU over the same capacity range, an uplift of around 16–17%, with inland throughput near 4.0×10^5 TEU and an inland share of approximately 30%. These results confirm that, for similar annual module volumes, changing capacity from 12 to 36 TEU mainly trades call frequency against per-call handling without materially altering the port-wide split between sea-going and inland moves.

The 48 TEU scenarios combine larger units with a higher module workload, which lifts both total and inland throughput beyond the 12–36 TEU cases. In the 50:50 mix, total throughput increases to about 1.33×10^6 TEU and the inland share rises to 31.6%. In the 25:75 mix, total throughput reaches roughly 1.41×10^6 TEU and the inland share about 35.6%. These outcomes illustrate that increasing the module share in the traffic mix, alongside higher per-module capacity, shifts the port-wide composition towards inland container flows. Taken together with the turnaround and terminal indicators, the results suggest that capacity choices primarily influence how a given module workload is packaged and scheduled, while larger jumps in inland share arise when the overall module volume in the system is increased.

Summary of Module Capacity Effects

Taken together, the module-capacity experiments show that capacity choices mainly reshape how a given module workload is packaged into calls and handled at the interface. For capacities between 12 and 36 TEU the annual module TEU is held approximately constant, so changes primarily trade call frequency against per-call handling and waiting. Under the high-module 25:75 mix, module-mother turnaround responds non-linearly: very small modules (12 TEU) lead to long module-mother turnaround, intermediate capacities of 24–36 TEU reduce it to around 2 h, and very large modules (48 TEU) push it back towards 7 h. In the balanced 50:50 mix, module-mother turnaround increases more smoothly with capacity from about 1.2 h at 12 TEU to roughly 2.3 h at 48 TEU, while barge, deep-sea and feeder turnaround times remain within a narrow band and show no systematic trend in either mix.

Module-child indicators and anchorage behaviour point to a mid-range capacity “sweet spot” around 24–36 TEU. As capacity increases, terminal dwell for module children rises steadily, while waiting

stays modest for 12–36 TEU and becomes much longer at 48 TEU in the 25:75 mix. Anchorage results reinforce this pattern. In 25:75, almost all module calls are routed via anchorage at 12 TEU with long waits, anchorage shares and waiting drop sharply at 24–36 TEU, and both increase again at 48 TEU. The 50:50 mix exhibits the same shape at lower overall levels: module anchorage shares are highest at 12 and 48 TEU and lowest at 24–36 TEU, while barge anchorage remains modest and comparatively stable in both mixes.

Terminal-level indicators confirm that changing module capacity does not move the representative terminal into a sustained berth- or crane-saturated regime. At Terminal B, average effective berth utilisation remains on a stable plateau as capacity varies, at about 53–54 % in the 25:75 mix and 64–66 % in the 50:50 mix, with peak utilisation close to 99 % for all capacities. Crane metrics show peak general-crane use near the installed level but only about half of total crane capacity in use on average, and congestion attribution measures indicate that crane-limited periods are rare, with slack or no-queue states dominating.

Port-wide throughput results reinforce the view that, for similar annual module volumes, capacity mainly repackages workload rather than altering the overall split between sea-going and inland flows. For 12–36 TEU, both mixes deliver higher total throughput than the 100:0 baseline: around 1.26×10^6 TEU with an inland share of 27–28 % in the 50:50 mix, and about 1.30–1.31×10^6 TEU with an inland share near 30 % in 25:75. The 48 TEU scenarios, which combine larger units with a higher module workload, lift both total throughput and inland share beyond these levels, particularly in the 25:75 mix. Overall, capacities in the 24–36 TEU range offer a robust compromise: they keep module-mother turnaround short in both mixes, minimise module anchorage and waiting, maintain stable performance for other vessel classes, and avoid sustained berth or crane scarcity. Very small modules create many short but intrusive visits, whereas very large modules reduce call frequency at the cost of longer waits and heavier bursts of work, especially when modules account for a high share of inland volume.

7.3.4. Overview of Configuration Effects

Taken together, the configuration experiments show that modular performance is shaped by a combination of thresholds and sweet spots, rather than increasing monotonically with added capacity. Across all cases, quay access and general crane capacity remain the governing resources, while the dedicated module system introduces additional constraints that must be dimensioned carefully.

For crane configuration, at least two module cranes per terminal are required in the high-module 25:75 mix to avoid single-server queues, excessive module anchorage and depressed inland throughput, whereas additional module cranes beyond two or three yield only modest gains because quay access becomes limiting. Handling-speed experiments indicate that moderate modular uptake (50–50) is comparatively robust to slower module cranes, sustaining higher throughput and inland shares than the conventional 100:0 baseline, while high modular shares (25–75) become highly sensitive: once the module crane pool is heavily loaded, small further slowdowns generate very large increases in modular waiting and erode inland volumes.

Module-capacity tests reveal a mid-range sweet spot around 24–36 TEU. For a given annual module workload, these capacities keep module turnaround short, minimise module anchorage and waiting, and maintain stable performance for barges, deep-sea vessels and feeders, without pushing representative terminals into berth or crane saturation. Very small modules produce many short but intrusive visits, whereas very large modules reduce call frequency at the cost of longer waits and heavier bursts of work, especially at high modular shares. Well-dimensioned modular systems therefore combine sufficient module cranes, adequately fast handling and mid-range module sizes to unlock port-wide throughput and inland-share gains without creating new bottlenecks.

7.4. Early Modular Adoption Scenario

This section analyses an early modular adoption scenario that reflects a modest near-term implementation. Relative to the 100:0 baseline, annual inland throughput is increased by approximately 10%, and all additional inland TEU are handled by modular vessels rather than by expanding the conventional barge fleet. Sea-side demand remains unchanged, so that the 100:10 futures represent incremental changes around the baseline system rather than a fully reoptimised port.

Within the 100:10 inland mix, two levers are varied. The first is whether modules receive dedicated handling capacity through a small module-crane pool at key inland terminals. Two dedicated-module configurations are considered: *+1MC* (configuration `mod1`) and *+2MC* (configuration `mod2`), both operating at a handling speed of 1.5 min/container. The second lever concerns cases where modules do not have dedicated cranes but are processed on the general quay crane pool. For these "modules on quay cranes" configurations, module handling is set either to 1.5 min/container (same move speed as the module cranes) or to 3.0 min/container (same as the barge handling time). For each speed, three quay-crane capacities are examined: *0QC* (configuration `gen0`, no new quay cranes), *+1QC* (`gen1`) and *+2QC* (`gen2`), in which one or two quay cranes are added at selected inland terminals.

7.4.1. Vessel KPIs

Vessel-level indicators describe how these growth configurations affect the time ships spend in the port system. Table 7.41 reports the average turnaround time per vessel type for the 100:0 baseline and the 100:10 futures. Turnaround time is measured from arrival at the port area to departure and includes both waiting and service. Table 7.42 shows the corresponding average waiting time at terminals, that is, the time vessels spend queued or idle before handling starts.

Table 7.41: Average turnaround time in hours by vessel type for the 100:0 baseline and early modular adoption futures (100:10). For 100:10, "+1MC" and "+2MC" denote dedicated module cranes with 1.5 min/container handling; "0QC", "+1QC" and "+2QC" denote the number of additional quay cranes when modules are handled on the general crane pool at either 1.5 or 3.0 min/container.

	100:0	100:10, on MC (sp1.5)		100:10, on QC (sp1.5)			100:10, on QC (sp3.0)		
Vessel type	Baseline	+1MC	+2MC	0QC	+1QC	+2QC	0QC	+1QC	+2QC
Barge	33.65	33.89	33.40	36.77	29.14	23.64	33.16	29.30	25.08
Deepsea	28.51	28.48	28.72	29.11	28.22	27.48	28.91	28.03	27.36
Feeder	15.39	15.36	15.69	15.78	15.08	14.49	15.76	14.95	14.26
Module (mother)	–	1.24	1.46	26.24	18.72	14.02	23.79	20.08	15.79

Table 7.42: Average waiting time in hours at terminals by vessel type for the 100:0 baseline and early modular adoption futures (100:10). For 100:10, "+1MC" and "+2MC" denote dedicated module cranes with 1.5 min/container handling; "0QC", "+1QC" and "+2QC" denote the number of additional quay cranes when modules are handled on the general crane pool at either 1.5 or 3.0 min/container.

	100:0	100:10, on MC (sp1.5)		100:10, on QC (sp1.5)			100:10, on QC (sp3.0)		
Vessel type	Baseline	+1MC	+2MC	0QC	+1QC	+2QC	0QC	+1QC	+2QC
Barge	23.62	24.15	24.64	25.59	19.80	14.64	23.29	19.19	16.00
Deepsea	0.74	0.73	0.89	0.92	0.89	0.71	1.08	0.71	0.65
Feeder	0.36	0.30	0.40	0.47	0.44	0.39	0.56	0.39	0.23
Module (child)	–	0.13	0.21	9.52	7.07	5.14	8.96	7.18	5.63

For deepsea and feeder vessels, average turnaround times change only marginally across the futures. Deepsea calls remain close to their baseline value, moving within roughly one hour across all 100:10 configurations, and feeders stay in a narrow band around 15 h. The waiting-time results in Table 7.42 show that these turnaround times are still dominated by service rather than by queueing: deepsea and feeder waiting times remain well below 1 h on average in every scenario. This is consistent with the terminal-level indicators, which show only limited changes in berth occupancy and crane utilisation at the sea-freight-dominated terminals. In other words, the introduction of modular inland services and the associated crane investments do not materially disrupt seagoing traffic in this experiment.

Barge turnaround times are more sensitive to how additional inland work is resourced. When dedicated module cranes are installed (*+1MC* or *+2MC*), average barge turnaround time remains close to the baseline level and barge waiting times change only slightly. The extra modular volume is largely absorbed by the module cranes, so the barge experience at the quay is similar to the 100:0 system. When modules instead draw on the existing quay crane pool without extra capacity (100:10, 0QC), barge turnaround time increases and waiting time rises to around 25–26 h, reflecting a higher share

of periods in which cranes become the bottleneck. Adding one or two quay cranes at the inland terminals (*+1QC* and *+2QC*) reverses this effect: barge turnaround times fall below the baseline, and barge waiting drops to roughly 20 h and 15–16 h respectively. The same pattern appears at both 1.5 and 3.0 min/container module speeds, although the gains are slightly smaller in the slower-handling cases.

Module-mother and child results are highly configuration dependent. With dedicated module cranes (*+1MC* and *+2MC*), mother vessels complete their visit in about 1.2–1.5 h on average, and child modules see waiting times of only a few tenths of an hour, indicating that the small module crane pool almost always has capacity available. When modules are handled on quay cranes with no added capacity (0QC), mother turnaround times rise into the tens of hours and child-module waiting jumps to around 9–10 h, showing that modular calls frequently queue behind other work on the general crane pool. Quay-crane investments reduce these delays: moving from 0QC to +1QC and +2QC cuts child waiting from roughly 9.5 h to about 7 h and 5 h respectively, and the corresponding mother turnaround times fall but remain substantially higher than in the dedicated-crane cases. The slower 3.0 min/container setting follows the same pattern, with somewhat lower mother turnaround times but still much longer queues than under dedicated module handling.

Taken together, the turnaround time and waiting-time indicators show that the main vessel-level impacts of the 100:10 growth configurations are concentrated in inland services. Seagoing calls see only small changes in total turnaround time and very limited queueing, while barge and especially module performance is strongly shaped by whether modules draw on dedicated cranes or compete for the existing quay crane pool, and by whether quay crane capacity is expanded alongside the growth in inland demand.

7.4.2. Terminal KPIs

Terminal level indicators summarise how the early modular adoption futures affect the use of quay length and crane capacity. For Terminal B, Table 7.43 reports average and peak effective berth utilisation. Effective utilisation is measured as the share of the available quay length that is occupied by vessels, based on effective vessel length. In the 100:0 baseline, Terminal B uses about 80% of its effective quay length on average, with short peaks where almost the full length is occupied (around 99%). Across all 100:10 configurations, average effective utilisation remains in a narrow band between roughly 78 and 81%, with peak values consistently close to full occupancy. This indicates that growth is absorbed mainly by filling previously idle gaps in the berth schedule rather than by pushing the terminal into persistently higher average utilisation.

Table 7.43: Effective berth utilisation at Terminal B for the 100:0 baseline and early modular adoption futures (100:10).

Scenario	Avg. effective utilisation (%)	Peak effective utilisation (%)
Baseline	80.3	99.2
100:10, dedicated module cranes (1.5 min/container)		
+1MC	79.3	99.1
+2MC	78.8	99.2
100:10, modules on quay cranes (1.5 min/container)		
0QC	80.1	99.3
+1QC	79.5	99.4
+2QC	79.8	99.3
100:10, modules on quay cranes (3.0 min/container)		
0QC	78.3	99.2
+1QC	80.9	99.2
+2QC	80.0	99.4

Short module calls reduce average service time per call and can improve the packing of work across the day. This allows the terminal to handle more moves while keeping total occupied metre-hours almost constant. Both dedicated-module configurations (+1MC and +2MC) record slightly lower effective utili-

sation than the baseline, consistent with modular calls inserting into residual windows between barge and seagoing work. When modules share the existing quay crane pool (0QC), effective utilisation returns close to the baseline value. Adding one or two quay cranes in the +1QC and +2QC configurations restores some slack and keeps effective utilisation close to its original level. The 3.0 min/container variants show very similar patterns; for instance, the +1QC (3.0) case yields effective utilisation comparable to the baseline with a modest increase in average utilisation, suggesting that once handling speeds are aligned, modular calls primarily change how work is distributed along the quay rather than pushing the terminal into sustained congestion.

Crane indicators reveal a different aspect of the same growth path. Table 7.44 reports the split between general and module cranes and the degree to which the installed capacity is used. In the 100:0 baseline, Terminal B has 16 general and 2 module cranes installed. On average, about 12 general cranes are active while module cranes are effectively idle, so roughly two thirds of the total capacity are used and peak total use reaches about 89% of installed capacity.

Table 7.44: Crane capacity and utilisation at Terminal B for the 100:0 baseline and early modular adoption futures (100:10). General cranes handle all barge and seagoing work; module cranes are dedicated to modules in the +1MC and +2MC configurations.

	General (QC)			Module (MC)			Total			
Scenario	Inst.	Peak Used	Avg. Used	Inst.	Peak Used	Avg. Used	Inst.	Peak Used	Avg. Used	Avg. vs Inst. (%)
Baseline	16	16.00	12.01	0	0.00	0.00	16	16.00	12.01	75.06%
100:10, dedicated module cranes (1.5 min/container)										
+1MC	16	16.00	11.75	1	1.00	0.21	17	17.00	11.96	70.32%
+2MC	16	15.99	11.71	2	2.00	0.21	18	17.94	11.92	66.22%
100:10, modules on quay cranes (1.5 min/container)										
0QC	16	16.00	12.04	0	0.00	0.00	16	16.00	12.04	75.24%
+1QC	17	17.00	12.18	0	0.00	0.00	17	17.00	12.18	71.65%
+2QC	18	18.00	12.48	0	0.00	0.00	18	18.00	12.48	69.34%
100:10, modules on quay cranes (3.0 min/container)										
0QC	16	16.00	11.80	0	0.00	0.00	16	16.00	11.80	73.75%
+1QC	17	17.00	12.58	0	0.00	0.00	17	17.00	12.58	74.01%
+2QC	18	18.00	12.60	0	0.00	0.00	18	18.00	12.60	70.01%

In the dedicated-module cases (+1MC and +2MC), the total installed crane capacity is similar to or slightly below the baseline, and average total use stays around 66–70% of installed capacity. Modules draw only a small fraction of this capacity on average (about 0.2 cranes), so the additional inland volume is absorbed with little change in total duty factor. When modules are handled exclusively by the general quay cranes (0QC, +1QC, +2QC), the module cranes are removed and the installed capacity consists only of general cranes. With 0QC, average total use rises to about 75% of installed capacity, and even with +1QC and +2QC it remains around 69–72%. The 3.0 min/container variants show almost identical duty factors: slower module handling slightly shifts work among scenarios but does not fundamentally change the average load. From the perspective of this terminal, slower module handling therefore manifests primarily as a rebalancing of work between crane types and a higher duty factor in the “modules on QC” futures, rather than as a clear move into a globally crane-saturated regime.

The bottleneck counters in Table 7.45 clarify how these utilisation patterns translate into operational constraints. In the 100:0 baseline, Terminal B is almost always in a slack state: cranes are binding in essentially none of the periods, and slack or no queue is recorded in about 99.6% of the time steps. The dedicated-module futures remain close to this regime. In the +2MC case, cranes limit operations in only about 0.5% of the time; in +1MC the crane-limited share rises to around 6–7%, but slack still accounts for more than 90% of time.

When modules share the quay crane pool without extra capacity (0QC), the picture changes markedly.

In the 1.5 min/container 0QC configuration, the share of time that is crane limited increases to roughly 35%, and slack falls to about 64%. Adding quay cranes in the +1QC and +2QC configurations reduces the crane-limited share to about 28% and 25% respectively, and slack recovers to around 72–75%, at the cost of higher installed crane capacity. The 3.0 min/container variants exhibit very similar crane-limited shares to their 1.5 min/container counterparts, even though module handling times match barge handling times. Across all configurations, quay limitations remain negligible at Terminal B, so cranes are the dominant resource that transitions from a non-binding capacity in the baseline and dedicated-module cases to a frequent bottleneck in the "modules on QC" growth scenarios.

Table 7.45: Share of time limited by cranes, policy or slack at Terminal B for the 100:0 baseline and early modular adoption futures (100:10).

Scenario	Crane limited	Policy limited	Slack or no queue
Baseline	0.0%	0.4%	99.6%
100:10, dedicated module cranes (1.5 min/container)			
+1MC	6.7%	0.4%	92.9%
+2MC	0.5%	0.4%	99.2%
100:10, modules on quay cranes (1.5 min/container)			
0QC	35.4%	0.4%	64.2%
+1QC	27.6%	0.4%	72.0%
+2QC	24.6%	0.5%	74.9%
100:10, modules on quay cranes (3.0 min/container)			
0QC	32.2%	0.4%	67.4%
+1QC	32.4%	0.4%	67.2%
+2QC	25.1%	0.5%	74.4%

Taken together, the berth and crane indicators show that moderate growth in modular traffic at Terminal B does not lead to dramatic changes in time-based berth occupancy or length-based berth fit. Instead, the main effect of the 100:10 growth configurations is a shift in the utilisation of crane capacity. Dedicated-module configurations keep the terminal in a largely slack regime, whereas futures in which modules are processed on the general quay crane pool, especially without additional quay cranes, move the terminal into a state where cranes become the primary operational constraint.

7.4.3. Port-wide KPIs

Port-wide indicators summarise how the early modular adoption futures affect total throughput and the balance between inland and seagoing traffic. In this scenario, an additional 10% of annual inland demand is routed via modules in a 100:10 barge:module mix, while deepsea and feeder demand remain unchanged. Table 7.46 reports the total number of containers processed across all vessel types and terminals, together with inland throughput and the inland share.

In the 100:0 baseline, the port handles approximately 1.12×10^6 TEU in total, of which about 2.20×10^5 TEU (19.7%) move via inland barges. Introducing modular inland services with a 100:10 mix increases total throughput by roughly 3% to 6%, depending on how modules are handled. The configuration with two dedicated module cranes (+2MC, 1.5 min/container) delivers about 4.4% more throughput than the baseline, while the most capacity intensive quay-crane expansion (+2QC, 3.0 min/container) reaches an uplift of around 5.4%.

These gains are driven mainly by inland traffic. Inland throughput (barges and modules combined) rises from about 2.20×10^5 TEU in the baseline to between 2.55×10^5 and 2.79×10^5 TEU in the 100:10 futures, an increase on the order of 16% to 26%. Over the same configurations, the inland share of total port moves increases from 19.7% to around 23% to 24%. Deepsea and feeder volumes remain close to their baseline levels, so the extra barge and module moves are largely additive rather than displacing seagoing calls.

The crane configurations modulate how much of this potential is realised. With two dedicated module

Table 7.46: Port-wide throughput and inland share for the 100:0 baseline and early modular adoption futures (100:10). Total throughput is reported in $\times 10^6$ TEU, inland throughput (barge + module) in $\times 10^5$ TEU.

Scenario	Total TEUs ($\times 10^6$)	Change vs 100:0	Inland TEUs ($\times 10^5$)	Inland share
Baseline	1.12	-	2.20	19.7%
100:10, dedicated module cranes (1.5 min/container)				
+1MC	1.17	+4.4%	2.66	22.8%
+2MC	1.17	+4.4%	2.66	22.7%
100:10, modules on quay cranes (1.5 min/container)				
0QC	1.16	+3.2%	2.55	22.1%
+1QC	1.17	+4.5%	2.68	22.9%
+2QC	1.19	+5.8%	2.79	23.5%
100:10, modules on quay cranes (3.0 min/container)				
0QC	1.16	+3.2%	2.55	22.1%
+1QC	1.17	+4.9%	2.67	22.7%
+2QC	1.18	+5.4%	2.76	23.4%

cranes (+2MC), the additional inland work is absorbed with limited impact on the general quay crane pool, yielding a throughput uplift of about 4.4% and an inland share of 22.7%. When modules are handled on the existing quay cranes without extra capacity (0QC), total throughput increases by only about 3.2% and inland share sits at the lower end of the 100:10 range, reflecting the crane bottlenecks observed at key inland terminals. Adding quay cranes in the +1QC and +2QC futures allows more modular calls to be processed without prolonged crane constraints and pushes both total throughput and inland share towards the upper end of the range.

The 3.0 min/container variants isolate the effect of quay use from the effect of faster module handling. Their throughput levels are very close to those of the corresponding 1.5 min/container cases, both in total and for inland traffic. This suggests that the primary port-wide impact of the modular configurations in this study is not an artefact of faster crane cycles, but stems from the additional inland calls and the way they are fitted into existing berth and crane capacity.

7.4.4. Overview of Early Adoption Case Results

Table 7.47 summarises the main effects of the 100:10 futures relative to the 100:0 baseline. Across all 100:10 configurations total port throughput increases by about 3.2% to 5.8%. Inland throughput (barges and modules combined) rises by roughly 16% to 26%, which shifts the inland share of total throughput from 19.7% in the baseline to between 22.1% and 23.5%. Deepsea and feeder turnaround times move only a few percent around their baseline values, which confirms that the adoption of modular inland services in these scenarios does not materially disrupt seagoing traffic.

The main vessel-level impacts are concentrated in inland services. Across all 100:10 futures, average barge turnaround time ranges from about 9% higher than the baseline when modules share existing quay cranes without extra quay capacity (100:10 0QC) to roughly 30% lower when two additional quay cranes are installed (100:10 +2QC). The corresponding changes in barge waiting times span from an increase of around 8% to a reduction of nearly 38%. Configurations with dedicated module cranes (100:10 mod1 and 100:10 +2MC) keep barge port and waiting times very close to the baseline while giving modules very short turnaround. When modules are instead handled by the general quay crane pool (0QC, +1QC, +2QC), the extent to which the additional inland demand is absorbed without long queues depends critically on whether extra quay cranes are added.

Overall, a ten percent modular share delivers higher port throughput and a larger inland share without major adverse effects on seagoing vessels. The choice of crane configuration determines whether the additional inland traffic translates into longer barge queues, as in the 100:10 0QC layouts where modules compete for existing quay cranes, or into shorter barge and module turnaround times, as in the 100:10 +1QC and +2QC futures where crane capacity is expanded. This pattern implies that, for modest levels of modular uptake, the main design lever is how crane capacity is organised around

Table 7.47: Summary of changes in key KPIs across the 100:10 futures relative to the 100:0 baseline.

KPI	Change vs 100:0
Total port throughput	$+3.2\%$ to $+5.8\%$
Inland share of throughput	$+2.4$ to $+3.8$ percentage points
Barge average turnaround time	about $+9\%$ to -30%
Barge average waiting time	$+8.3\%$ to -38.0%
Deepsea average turnaround time	about -4% to $+2\%$
Feeder average turnaround time	about -8% to $+3\%$

inland traffic rather than whether modular services are deployed at all.

7.5. Conclusion

The simulation results show that autonomous modules deliver sustained operational benefits primarily through improved temporal fit at the quay. Short, concentrated module calls convert previously idle micro-intervals into productive handling without requiring additional quay length. As the inland fleet becomes more modular, total and inland throughput increase and average turnaround times fall for barges, while deep-sea and feeder services see smaller but consistent improvements. Across the mix experiments, intermediate modular shares are most attractive. The balanced 50:50 and higher-module 25:75 cases capture most of the gains in turnaround, anchorage and throughput. The fully modular 0:100 configuration still lifts total and inland throughput, but at the cost of very long module queues and high anchorage exposure for module children once all inland work is routed via modules.

Across the configuration experiments, quay access and general cranes remain the governing resources. Peak use of general (quay) cranes is close to the installed level in almost all variants, while average duty factors move only moderately. Adding module cranes beyond the level required to avoid a single-server queue yields limited aggregate benefit. In the 25:75 mix at least two module cranes per terminal are needed to stabilise module turnaround and keep module anchorage at manageable levels; a third crane mainly creates slack. In the 50:50 mix the module workload per crane is lower and one to two module cranes suffice under the tested demand. The stress tests with a very large notional pool of module cranes increase the peak number of cranes actually used only marginally and sharply reduce average utilisation relative to installed capacity. This reveals a ceiling set by berth access and scheduling rather than by module-crane inventory. Varying module capacity does not alter this finding: module cranes are busy under high module shares, but crane-driven congestion at the representative terminal remains rare.

Berth occupancy and anchorage respond primarily to call structure. Effective berth utilisation at Terminal B declines systematically as the module share increases, even though throughput rises, because shorter inland calls pack more efficiently into small berth pockets. For deep-sea and feeders, anchorage shares and waiting times stay low and are almost insensitive to crane configuration, capacity or module speed, reflecting their protected access in the scheduling rules. Inland anchorage, by contrast, reacts strongly when the module system is poorly dimensioned. With only one module crane in high-module mixes, module mothers and children experience very long waits and almost all module calls route via anchorage, while barges remain largely protected. Moving from one to two module cranes collapses these single-server queues; gains beyond two are modest. Under constant annual module TEU, capacities around 24–36 TEU minimise module anchorage and waiting. Very small modules create many short but intrusive calls that must be staged, whereas very large modules reduce call frequency but generate heavier pulses of work, which again increase module anchorage and waiting, especially in the high-module 25:75 configuration.

Demand experiments with sea-freight intensity scaled between 80%, 100% and 120% confirm that the modular concepts are robust within the tested range. For all three mixes (100:0, 50:50 and 25:75) the throughput ratio remains close to one, so almost all requested moves are processed. Total port throughput scales approximately linearly with sea-side demand, and at any given demand level the mixed systems handle substantially more TEU than the conventional system, particularly on the inland side. Deep-sea and feeder turnaround times increase only slightly as demand rises and remain close

to their baseline levels across all mixes. Barge turnaround is much more sensitive to demand in the purely conventional 100:0 case, where higher sea-freight load translates into long additional waiting at anchorage and dolphins. In the 50:50 and 25:75 mixes, barge turnaround remains much lower and much less sensitive to the 20% demand swing, because modular calls fill berth pockets without forcing barges back into long anchorage queues. At Terminal B, average berth utilisation and crane duty factors increase with sea-freight intensity but remain well below saturation, and crane- and quay-limited periods account for only a small share of time.

The handling-speed experiments show that slower module cranes mainly affect modular services once the dedicated pool becomes heavily utilised. For moderate modular uptake (50:50), changing the module handling time between 1.5 and 3.5 min/TEU leaves barge, deep-sea and feeder turnaround and port-wide throughput almost unchanged, while module mothers and children see longer waits only at the slowest speeds. For high modular uptake (25:75), the same slowdown produces a pronounced, non-linear deterioration in module turnaround and a visible erosion of inland throughput and inland share as the module crane pool approaches its capacity limits. These runs underline that module speed is a second-order design choice at moderate modular shares, but becomes critical when most inland work is routed via modules.

Together, the configuration and capacity experiments indicate that modular performance is shaped by a combination of thresholds and sweet spots rather than by monotonic benefits from added capacity. For high modular shares, at least two module cranes per terminal and mid-range module capacities around 24–36 TEU are required to avoid single-server effects and excessive module staging. Beyond these thresholds, additional module cranes or faster handling mainly provide local robustness, because berth access and general-crane capacity become the binding resources. The representative inland terminal remains in a high-utilisation but non-saturated regime across all tested configurations, and crane-limited periods at that terminal explain only part of the waiting observed in the module system.

A forward-looking 100-10 scenario, which introduces a modest modular share on top of the current barge system, illustrates a pragmatic adoption step. Across the 100-10 configurations, total port throughput increases by about 3% to 6% and inland throughput by roughly 16% to 26% relative to the conventional 100:0 baseline. The inland share of throughput rises by about 2 to 4 percentage points, while deep-sea and feeder turnaround times remain close to baseline values. The main differences between 100-10 variants arise from how cranes are organised. When modules share the existing quay crane pool without additional capacity, crane-limited periods at inland terminals become frequent and barge waiting worsens. Configurations with dedicated or additional cranes keep berth occupancy close to baseline levels, reduce barge waiting relative to the conventional case and still deliver higher inland throughput. This pattern suggests headroom for further modular uptake at current quay inventory, provided that crane capacity around inland work is dimensioned and scheduled explicitly.

Implications for Design and Operations

- Primary lever: berth access and planning. Since quay access governs in almost all experiments, the most effective levers are berth-window design, look-ahead settings and pocket-fit optimisation. Module-crane additions beyond the level needed to remove single-server effects deliver diminishing returns under the tested demand.
- Module-crane sizing and mix choice. In module-heavy mixes such as 25:75, at least two module cranes per terminal are required to avoid persistent module queues and excessive staging; three or more cranes mainly provide robustness against local surges. At moderate modular shares, such as 50:50 or early 100-10 adoption, one or two module cranes per terminal combined with careful berth planning already capture most of the benefits while keeping crane utilisation in a comfortable range.
- Module capacity and speed. Capacity should be selected on operational grounds, trading call frequency against per-call dwell and considering berth geometry. Mid-range capacities around 24–36 TEU minimise module waiting and anchorage under constant annual module TEU, while very small or very large modules increase staging requirements. Module-handling speed is of secondary importance at moderate modular shares but becomes a key design parameter when most inland work is routed via modules.

- Anchorage management. Where anchorage spikes arise in single-crane configurations or for extreme module sizes, design responses include upgrading to two module cranes per terminal, adjusting berth-window rules, or introducing short micro-windows and near-berth holding for modules to smooth staging without compromising sea-side predictability.
- Measurement. Time-based capacity indicators such as berth occupancy and cranes-in-use should be paired with productivity measures such as moves per berth-hour and moves per active crane-hour. This avoids misinterpreting higher occupancy as congestion or lower occupancy as slack and provides a clearer link between resource use and throughput.

The conclusions are drawn for the tested mixes, berth plans and operating policies. Sensitivity to alternative admission rules, different convoy recombination patterns, asymmetric terminal layouts and other port designs warrants targeted follow-up. Including confidence intervals or other measures of statistical uncertainty around key KPIs would add precision to small differences between close-performing cases and help distinguish genuine performance effects from sampling variation.

8

Discussion

This chapter reflects on the findings of the simulation study and discusses their implications for operational efficiency and coordination within port-wide vessel–terminal operations. It places the results in context, outlines the limitations of the current research, and identifies directions for future development.

Interpretation of Results

The results presented in Chapter 7 demonstrate that the composition of the vessel fleet and the configuration of terminal resources have a pronounced influence on port-wide performance. The comparison of vessel ratios reveals that a balanced mix of conventional barges and modular or autonomous vessels achieves the highest overall efficiency. The intermediate scenario (50:50 ratio) yields the lowest average time in port and the most balanced quay utilisation, suggesting that hybrid operations may offer a pragmatic transition pathway for ports introducing modular concepts alongside existing barge services. A fully modular system further reduces waiting time but introduces a higher number of vessel calls, which can increase berth fragmentation if scheduling coordination is insufficient. Under constant annual module TEU, the capacity tests show that larger modules substitute fewer, longer calls for more, shorter ones. Module-mother visits therefore lengthen mainly through additional terminal dwell, while barge, feeder and deep-sea times remain broadly unchanged. Fewer insertions reduce window fragmentation and lower anchorage for inland traffic, whereas deep-sea anchorage stays near 7–8% by design. Effects on berth-fit are terminal specific: where short calls previously created many unusable gaps, fewer, larger calls improve packing; where geometry and long blocks dominate, longer visits can leave remainder gaps that reduce effective fit. Across both mixes, crane-driven congestion remains negligible; module-crane productivity rises slightly with capacity as changeovers fall, but cranes are not the binding constraint.

The results also show that improvements can be realised almost immediately after implementation. The rapid-shift scenario (100–10) demonstrates that even a modest initial share of modular vessels substantially reduces average waiting times and quay congestion. This finding illustrates that ports can begin to capture operational benefits well before full system conversion, which is particularly relevant for incremental innovation under real-world constraints such as limited fleet renewal or phased policy rollout.

From an operational perspective, the findings indicate that ports can achieve meaningful efficiency gains without major infrastructure expansion, provided that berth allocation and crane scheduling are adapted to handle smaller, more frequent arrivals. Modular and autonomous units distribute quay demand more evenly and sustain continuous crane activity by reducing idle intervals between vessel exchanges. This aligns with the logic of resource pooling used in multi-user terminals and reinforces the value of flexible scheduling for congestion management. Although overall crane productivity remained relatively stable across scenarios, utilisation became more uniform under mixed and modular conditions, with fewer interruptions due to vessel waiting. This implies a shift in the optimal operational setup: from peak-intensive deployment to steady, balanced utilisation. For terminal operators, such a

shift translates into higher predictability of labour and energy use and calls for adaptive scheduling and predictive planning tools, particularly in ports where quay cranes are managed under separate concession holders. The capacity experiments reinforce this: changing module capacity primarily rebalances call frequency and per-call dwell without raising persistent crane pressure. Module-crane hours become slightly more productive with larger calls, but quay access and berth-window design continue to govern system performance.

These implications are particularly relevant for the Port of Rotterdam, which operates under a landlord governance model in which the port authority manages shared infrastructure while private terminal operators hold individual concessions. The simulation outcomes highlight the importance of coordination mechanisms that transcend terminal boundaries, ensuring that aggregate berth demand remains synchronised with overall port capacity. Initiatives such as *Nextlogic*, which already facilitate integrated barge planning and scheduling across terminals, exemplify this need. The results reinforce the strategic value of such digital coordination systems, showing that synchronised scheduling can reduce anchorage times and enhance berth-fit efficiency without requiring new quay development. In particular, when annual module TEU is held constant, larger per-call capacity reduces the number of module insertions and thus the fragmentation of short windows, which lowers inland anchorage even as per-visit dwell increases. This clarifies why coordination tools that smooth the timing of short calls are effective levers alongside physical capacity.

At Maasvlakte 2, where terminals such as APMT and RWG employ highly automated cranes and guided vehicles, the model's insights suggest that modular or autonomous vessels could be integrated with minimal physical adaptation on the quayside. The primary challenge would lie in operational synchronisation rather than infrastructure readiness. A modular arrival structure aligns well with the continuous handling cycles of automated terminals, which are designed for predictable, high-frequency service. For the port authority, these findings imply that the transition toward modular and autonomous operations should be accompanied by increased data sharing, common scheduling frameworks, and adaptive berth allocation policies to balance utilisation across the port's waterway system.

In a broader perspective, the results underline that modular and autonomous vessel concepts can strengthen the resilience and adaptability of port operations. The distributed planning approach observed in Rotterdam proves particularly advantageous under uncertain or fluctuating conditions, as it allows terminals to adjust locally while preserving system-wide coherence. These findings suggest that the transition towards modularisation is not only a technological evolution but also an organisational one, requiring ports to combine digital integration with procedural flexibility. The model therefore highlights both the immediate potential of modular operations and the institutional prerequisites, distributed coordination, adaptive scheduling, and shared situational awareness, needed to realise them effectively.

Practitioner Validation and Implementation Perspective

To complement the model validation in Chapter 6 and Appendix E, the scenario findings were presented in a follow-up interview with the Business Manager Barge & Inland Terminals at the Port of Rotterdam. The discussion focused on how the simulated barge–module mixes relate to current barge planning practice, what levels of modular adoption are realistic, and how acceptable the assumed module-crane configuration is from a terminal perspective.

First, the expert linked the barge–module scenarios to the way inland work is currently organised. In his experience, roughly half of inland barge volume is handled through liner-type services with recurring fixed windows, while the remaining half consists of more irregular flows that are coordinated through the integrated barge-planning platform. Within this structure, modular intra-port traffic would naturally sit in the flexible, non-liner segment, where short calls and opportunistic slot filling add most value. This view supports interpreting the 50:50 barge–module scenario as more than an abstract split: it approximates a port in which conventional liner services continue to use fixed windows, while modules cover the non-liner share that is already coordinated via platform-based planning.

Second, the expert reflected on adoption levels and on how the scenarios map to plausible transition paths. While a 50:50 division of inland work between conventional barges and modular units was regarded as a credible long-term orientation, he judged it unlikely that ports would move directly to such

a high modular share. A more realistic near-term configuration would involve adding a modest volume of modular traffic on top of the existing system, on the order of ten percent of today's barge or road work, targeted at the part of the market that is already flexibly planned. This perspective suggests that the rapid-shift case with 100:10 barge–module shares is close to an implementable first step, whereas higher modular shares represent medium- to long-term or upper-bound cases rather than immediate policy options.

Third, the interview underscored an implementation constraint that is only stylised in the scenarios. The simulation assumes that modules are served by a pool of dedicated module cranes that can operate at terminal quays alongside terminal-owned ship-to-shore equipment. The expert indicated that container terminals are unlikely to accept external cranes driving onto their quays under current safety, liability, and automation arrangements. He considered alternative configurations more realistic, such as floating cranes positioned at barge quays or terminal-owned equipment explicitly dedicated to module handling. Under this interpretation, the module crane pool in the model is best read as an abstraction of additional, dedicated handling capacity for modules, whose ownership and integration would need to be negotiated in practice, rather than as a literal proposal to deploy barge-operated mobile cranes on deep-sea quays.

Taken together, the practitioner feedback supports the qualitative patterns observed in the results, in particular that a balanced barge–module mix can improve quay utilisation and barge turnaround without undermining deep-sea and feeder performance. At the same time, it emphasises that modular concepts are most naturally deployed in the platform-coordinated, non-liner segment of the barge market and that institutional acceptance of new handling arrangements is a critical condition. The modular scenarios should therefore be interpreted as evidence on the relative performance of alternative coordination concepts, conditional on creating suitable module-handling capacity, rather than as direct forecasts of short-term operational changes. This reading is consistent with the broader positioning of the model as an exploratory decision-support tool rather than a predictive planning system.

Limitations

While the simulation results offer valuable insights into the dynamics of port-wide vessel and terminal interactions, several limitations should be acknowledged. Access to detailed operational data from terminals and barge operators was restricted due to confidentiality constraints, which necessitated the use of representative averages for model calibration. Although these data provided realistic orders of magnitude, they may not capture the full operational variability observed in practice.

The model itself relies on several simplifying assumptions. Vessels were treated as homogeneous classes with deterministic handling rates and fixed service sequences, omitting stochastic variations in arrivals, crane performance, and terminal coordination. Environmental influences such as weather-related delays, maintenance downtime, or unplanned disruptions were excluded to maintain computational efficiency and reproducibility. As a result, the outcomes should be interpreted as indicative of relative performance differences rather than precise operational forecasts.

The experiments assume mixed-use quays and (effectively) uniform destination draws across terminals. In practice, traffic skews toward larger or preferred terminals and some ports operate dedicated barge quays. This simplification may understate localised window pressure and anchorage at high-volume terminals. An extension with terminal-specific demand weights and dedicated-quay logic would quantify distributional effects and fairness across concession areas.

The scope of the analysis was deliberately limited to vessel–terminal interactions within the Port of Rotterdam. Broader network effects, such as hinterland congestion, berth allocation conflicts between deep-sea terminals, and supply-chain propagation effects, were not simulated. Similarly, the arrival process was generated synthetically to ensure volume consistency, rather than based on published liner schedules. This is acceptable for congestion and capacity analysis but could be refined in future research using empirical schedule data.

The simulated time horizon covers representative four-week periods, which provide a realistic operational snapshot but do not account for seasonal variations or long-term scheduling dynamics. Moreover, sensitivity analyses were conducted only for selected parameters such as berth capacity and crane allocation. Other potentially influential operational factors, such as variability in vessel arrivals,

stochastic handling rates, berth-specific operating constraints, and heterogeneous scheduling policies across terminals, were not included in the current model formulation. Incorporating these dynamics in future work would allow for a more comprehensive assessment of the robustness of modular and autonomous vessel operations under realistic operating uncertainty.

Model validation posed an additional challenge, as large-scale deployment of autonomous or modular barge systems has not yet occurred in practice. Consequently, validation relied on expert interviews, secondary datasets, and plausibility checks rather than direct performance measurements. The follow-up practitioner assessment reported above strengthens confidence in the qualitative direction of the findings, but it does not remove the inherent uncertainty in extrapolating from simulated scenarios to real-world implementation. This introduces uncertainty in interpreting the absolute magnitude of the results, though the relative trends remain robust. Finally, the outcomes are specific to the layout and governance structure of the Port of Rotterdam and may not be directly transferable to other ports with different configurations or automation levels.

Finally, capacity findings are conditional on holding annual module TEU constant while scaling mission counts and module length. If future uptake increases total module work rather than redistributing it, the balance between call frequency, dwell, and berth-fit may shift; extending the analysis to non-constant TEU cases would test the robustness of these patterns.

These limitations do not undermine the validity of the study but rather highlight areas for future refinement. The model's transparent structure allows for straightforward integration of stochastic arrival processes, seasonal variations, and richer datasets as they become available, thereby improving representativeness and robustness in subsequent research.

Implications and Outlook

The findings demonstrate that discrete-event simulation is a powerful method for assessing future coordination strategies and technological transitions in port operations. From a managerial perspective, the model provides a decision-support tool to evaluate the trade-offs between operational efficiency, berth congestion, and resource utilisation under different fleet compositions. The evidence suggests that partial substitution of traditional inland barge services with modular autonomous units can increase system resilience, especially during peak demand or irregular arrival patterns. When selecting module capacity, the practical trade-off is between fewer, longer calls, which simplify scheduling and lower inland anchorage, and more, shorter calls, which may tighten berth packing in some geometries. Cranes remain secondary provided single-server effects are avoided.

A useful way to read these results is against the continuing trend toward larger sea-going ships, which concentrates exchanges into fewer, lumpy calls and makes mainline berths sensitive to micro-interruptions. Modular inland units move in the opposite direction at the unit level, yet they complement upsizing at the system level. By filling short berth opportunities with small, predictable exchanges, they smooth peaks around very large calls, protect crane intensity on ocean berths, and shorten the tail of berth occupancy after completion. In practice, modules should be planned as elastic capacity around the large-ship schedule, preferably on side or barge-dedicated quays sized so that one or two modules can complete within a typical short window, while primary deep-sea cranes remain focused on the main exchange.

Translating this into design, several choices at the vessel and quay interface unlock the operational gains observed. Module capacities should align with common spreader modes for 20- and 40-foot moves, allow one or two modules to complete within a short window at the case terminals, and keep convoy length compatible with bollard spacing and fender geometry so mooring times remain predictable. Manoeuvrability and approach control matter for reliability: fine-control propulsion, bow and stern thrusters, and stable station-keeping reduce lines-on to first-lift variance, and for autonomy, sensor and control stacks should be tuned to low-wash, port-speed approaches. On the quay, one assured crane per short call with surge to two caps dwell tails, quick tool changes avoid spreader penalties when alternating between 20- and 40-foot moves, and near-quay buffers sized to a module batch prevent micro-idleness. Where electric propulsion is assumed, duplicated charging pedestals positioned to avoid conflicts with crane envelopes support opportunity charging during handling. Mooring hardware should match module length increments, with quick-release or auto-mooring solutions that compress

berthing and unberthing intervals. Digitally, publishing berth-window catalogues with explicit minima that match chosen module sizes, together with lightweight reservation and timestamped commitment interfaces, enables autonomy stacks to plan approaches that reliably hit earliest-start times.

Where safety rules and local practice allow it, enabling parallel rather than strictly serial working of modules provides further benefit. Two modules worked in parallel, with one crane per module and clear crane zoning to avoid interference, compress completion time and cut rendezvous delay without raising peak crane pressure. To make this practical, bollard and fender spacing should allow both units to make fast without repositioning, near-quay buffers and charging points should be duplicated at workstation spacing, and line leads should support safe abreast mooring. Where abreast mooring is not permitted, adjacent short pockets with separate access achieve the same effect while preserving safety margins. These configurations remove single-server effects for modular traffic at the busiest quays and align with the more uniform utilisation observed in the scenarios.

From a research perspective, the modelling framework can be extended to incorporate uncertainty, dynamic scheduling, and energy metrics, allowing emissions and energy consumption to be assessed alongside operational efficiency. The approach is portable to other multi-terminal ports by adapting layout parameters and operating rules.

Although the simulation demonstrates the potential benefits of modular and autonomous vessel operations, real-world implementation depends on enabling conditions. A sufficient level of digital integration is essential to allow real-time exchange of berth and vessel information between terminals. In landlord-type ports, where operators plan independently within shared infrastructure, distributed coordination frameworks supported by shared data platforms are particularly well-suited. A high degree of quay automation is advantageous but not strictly required: highly automated terminals can integrate modular or autonomous vessels primarily through synchronised planning interfaces, whereas semi-automated or conventional terminals may first invest in digital gate procedures, vessel detection, and scheduling systems. Capacity choice should be paired with the number of module cranes at the terminals that carry the highest module load. Larger modules reduce insertions but lengthen per-visit dwell, therefore the preferred configuration is the smallest crane count that removes single-server effects for the chosen capacity at the busiest terminals.

The degree of autonomy on the vessel side also shapes coordination requirements. Autonomous navigation within the port does not imply autonomous berthing or loading. Modules that self-navigate but rely on assisted berthing would still require local guidance near the quay; higher autonomy levels reduce manual oversight but require robust communication with terminal control systems. The present findings are most relevant for intermediate autonomy, where modules operate autonomously in transit and interact with human-supervised or semi-automated terminal processes.

Beyond terminal coordination, the increasing number of modular vessels has implications for traffic management on port waterways. As the share of modules rises, total movements increase, and navigational interactions become more complex. While these additional sailing elements help smooth quay utilisation, they raise the importance of real-time situational awareness and collision-avoidance across access channels and turning basins. A distributed planning approach remains effective in this context, provided it is supported by shared traffic pictures and alerting. Configurations with a modest initial modular share already deliver measurable efficiency gains while keeping the number of moving units within current traffic-monitoring capacity, which suggests a feasible intermediate step toward broader automation.

As a next step, optimisation-based window assignment with look-ahead and multi-agent coordination mechanisms could be layered on top of the current rule set to test whether berth-fit and anchorage can be improved further under different module capacities and non-constant module TEU scenarios.

In conclusion, modularisation and autonomy are not only technological innovations, they are coordination challenges. Their successful adoption will depend on adaptive scheduling frameworks, transparent data sharing, and close alignment between terminal and port authorities. The next chapter synthesises these findings into overarching conclusions and recommendations for research and practice.

9

Conclusions and Recommendations

This study investigated how different coordination strategies and fleet compositions influence the operational efficiency of port-wide container transhipment. Using a discrete-event simulation model of the Port of Rotterdam, the research explored the potential impact of modular and autonomous barge concepts on key performance indicators such as throughput, berth utilisation, waiting time, and quay productivity. The model was designed to reflect the distributed coordination structure of the port, where terminals operate independently yet exchange scheduling information through shared digital platforms.

Summary of Key Findings

The simulation results demonstrate that hybrid fleets combining conventional and modular vessels can achieve significant efficiency gains without requiring additional infrastructure. The 50:50 configuration produced the most balanced outcomes, minimising average time in port and anchorage waiting while maintaining stable berth and crane utilisation. The 100:10 case revealed that benefits emerge even at low adoption levels, indicating that partial modularisation offers a practical and low-risk first step. In contrast, the higher module-share scenarios (25:75 and 0:100) further improve quay fit and reduce idling but lead to a substantial increase in the number of moving vessels, emphasising the need for enhanced traffic coordination and real-time situational awareness in port waterways. At fixed module volume, larger units mean fewer, longer calls. Module-mother visits lengthen mainly through extra terminal dwell, while barge, feeder and deep-sea times are essentially unchanged; inland anchorage drops because fewer insertions fragment the quay less. Effects on berth fit are terminal-specific, and although module-crane productivity ticks up with size, cranes are not the bottleneck. Read against the continuing upsizing of sea-going ships, the modular concept complements, rather than competes with, large calls: by filling short berth opportunities with small, predictable exchanges it protects crane intensity on mainline berths and shortens the tail of berth occupancy.

Overall, the study shows that modularisation strengthens operational resilience by distributing quay demand more evenly and reducing peak congestion effects. However, its success depends on adaptive scheduling mechanisms, information transparency between stakeholders, and synchronisation of berth allocation across terminals.

Answers to Research Questions

The findings presented above can be directly linked to the research questions introduced in Chapter 1. This section summarises how each question has been addressed through literature review, conceptual development, and simulation analysis.

SQ1: What are the current operational characteristics and bottlenecks of barge-based intra-port transhipment?

Current operations at Rotterdam's barge quays are shaped by a distributed planning environment, high variability in call size, and competition for berth access during deep-sea peaks. Waiting times mainly arise from protected window schedules and unsynchronised arrival patterns. These factors limit quay

efficiency and motivate the search for more flexible coordination schemes.

SQ2: Which autonomous or coordination-enhanced vessel strategies are applicable and how do they differ from current practice?
Autonomous modular barges introduce smaller, more frequent calls that can be dynamically dispatched between terminals. Unlike conventional barges operating under fixed schedules, these systems allow flexible coordination within a shared digital planning framework. The study shows that distributed coordination, supported by platforms such as Nextlogic, provides the necessary structure for such systems while preserving terminal autonomy.

SQ3: How can discrete-event simulation be used to model port-intra-port transhipment processes under varying demand and coordination scenarios?
The developed simulation framework integrates vessel movements, berth assignment, and crane utilisation into a unified decision environment. It allows controlled variation of demand patterns and fleet composition while maintaining representative port operations. The approach proved effective for capturing queuing behaviour, resource utilisation, and coordination effects, demonstrating its suitability as a decision-support tool for port planning.

SQ4: What are the operational effects of different coordination and fleet composition scenarios?
The simulation shows that introducing modular inland barges reduces average waiting times, shortens barge turnaround, and smooths quay utilisation, especially in mixed fleets such as the 25:75 and 50:50 barge–module ratios. Hybrid fleets provide the best overall balance between throughput, inland share, and resource use, while very high modular shares further improve quay fit but increase the number of vessel movements and the need for tighter traffic coordination on port waterways. Across all scenarios, efficiency gains are achieved without additional quay or crane infrastructure. Berth access and general crane capacity remain the governing resources, and crane duty factors stay moderate even at higher demand levels. Under constant annual module volume, changing module capacity mainly repackages work at the module interface, with capacities around 24–36 TEU minimising module–mother turnaround, anchorage use, and waiting, whereas very small or very large modules increase staging and delay without materially affecting deep–sea or feeder performance. Early-adoption futures with a modest modular share (100:10) already deliver lower barge waiting and higher inland throughput, indicating that ports can realise benefits through gradual hybridisation and improved berth allocation rather than abrupt full modularisation.

SQ5: What are the implications of adopting modular or autonomous vessel systems for port management and policy?
Operationally, ports adopting modularisation must ensure sufficient digital maturity to manage higher vessel traffic and more granular scheduling. Policy-wise, landlord port authorities should facilitate data-sharing standards and invest in coordination platforms that enable distributed planning across terminals. The transition should be accompanied by updated traffic management and safety frameworks to accommodate increasing vessel numbers in constrained waterways. Two practical rules follow; modules should be sized with local berth geometry and window scarcity in mind, using fewer insertions where short windows are scarce, and single-server effects should be avoided by providing at least two module cranes in module-heavy mixes, since returns beyond two were diminishing in the simulation experiments.

SQ6: What broader insights can be drawn for future port design and automation strategies?
The results support the view that modularisation and autonomy are not isolated technological changes but systemic shifts that primarily pose coordination challenges. Ports that combine distributed planning platforms, adaptive berth- and window-design, upgraded traffic management, and gradual hybrid adoption of modular units will be best positioned to benefit from automation. The findings suggest that Rotterdam's landlord governance model and existing digital infrastructure provide a strong foundation for such integrated strategies.

Recommendations for Operations and Policy

From an operational standpoint, hybrid adoption should be prioritised. A modest initial share of modular units, for example the 100:10 configuration moving toward a 50:50 mix, already delivers congestion relief while keeping navigational density manageable. Module capacity should be treated as a schedul-

ing lever: larger units reduce insertions and lower inland anchorage, though they lengthen per-visit dwell and can reduce berth fit where quay geometry is restrictive. Capacity choice should therefore be paired with berth-window design and look-ahead settings, rather than with indiscriminate addition of cranes. In practice, select module size and the minimum number of module cranes per terminal as a pair. For each candidate capacity, choose the smallest crane count that removes single-server effects at the busiest terminals. Two cranes provide a practical default for mixed traffic, with increases only where measured module queues persist.

Landlord authorities can accelerate this transition by strengthening distributed coordination. Publishing berth-window catalogues with explicit minima tuned to local module sizes allows arrivals to plan reliably for earliest start. Extending reservation and time-stamped commitment interfaces on shared platforms, such as Portbase and Nextlogic, supports predictable approaches and reduces anchorage. Traffic-management capability should be upgraded to provide real-time situational awareness for the higher number of movements, and safety frameworks should be updated where needed to allow parallel working of modules with clear zoning and mooring rules.

Implementation should be staged and evidence-led. In the near term, improve data quality and window exposure, establish pilot corridors for modular calls, and monitor waiting, quay fit, and reliability. In the medium term, align window minima with the module sizes that emerge in practice, calibrate look-ahead in scheduling tools, and extend rules to cover parallel working at selected barge quays. A gradual, feedback-driven approach maintains operational reliability while scaling modular operations and provides the performance data required for policy refinement.

Recommendations for Design of Vessels and Quays

Design choices at the vessel–quay interface should complement the continuing upsizing of sea-going ships rather than compete with it. Small autonomous modules are best used as elastic capacity around very large calls, filling short berth opportunities, protecting crane intensity on mainline berths, and shortening the completion tail after ocean exchanges. Where local safety rules permit, terminals should be designed to enable parallel working of two modules, with independent moorings, one crane per module, and clear crane zoning. This compresses completion time and reduces rendezvous delay without raising peak crane pressure. If abreast mooring is not permitted, adjacent short pockets with separate access achieve the same effect while preserving safety margins.

On the vessel side, capacities should be chosen to complete one or two modules within typical short windows at the study terminals and to align with common spreader modes for 20- and 40-foot moves, so handling penalties are avoided. Manoeuvrability and approach control matter for reliability: bow and stern thrusters, fine-control propulsion, and stable station-keeping reduce variance between lines-on and first lift. For autonomous or supervised operation, sensors and control should be tuned to low-wash, port-speed approaches that are compatible with terminal safety envelopes.

On the quay side, one assured crane per short call with surge to two caps dwell tails, rapid tool changes minimise penalties when alternating between 20- and 40-foot moves, and a flexible mobile harbour crane can serve residual pockets of quay length. Mooring hardware should match the chosen module length increments, with continuous or closely spaced fenders and quick-release or auto-mooring solutions that shorten berthing and unberthing intervals. Near-quay buffers sized to a module batch prevent crane starvation during modular completion, and duplicated DC charging pedestals positioned outside crane envelopes enable opportunity charging during handling. Finally, digital exposure of window minima and lightweight reservation interfaces ensure that vessel control can plan approaches that reliably meet earliest start, translating the operational gains observed in the scenarios into day-to-day capability.

Future Research

The findings of this thesis open several directions for further research that can enhance both the analytical depth and practical applicability of the developed model.

One possibility for extension is to incorporate additional sources of operational uncertainty to evaluate system robustness. Extending the simulation to include stochastic vessel arrivals, variable handling

rates, and weather-related disruptions would provide a more realistic representation of day-to-day operations. This would allow assessing how modular and hybrid fleet configurations perform under irregular or high-variability conditions, an essential step toward reliability analysis. A direct extension is to relax the fixed-volume assumption in the capacity tests and study cases where module volume grows with size. This would clarify how fewer insertions, per-visit dwell, berth fit and anchorage rebalance when capacity changes also scale workload, and enable joint design of module size, berth-window policies and module-crane counts.

Another promising line of research concerns environmental and energy-related components. Integrating energy consumption and emissions estimation would enable assessment of the ecological footprint of different fleet compositions, offering valuable insights for carbon-reduction strategies and environmental policy evaluation.

A further refinement could involve embedding optimisation-based or agent-based coordination mechanisms. An optimisation layer could evaluate the theoretical upper bound of system performance, while a multi-agent framework would allow the simulation of autonomous decision-making and adaptive coordination among vessels and terminals in real time. This would bridge the current rule-based logic with more intelligent control architectures aligned with future automation levels.

Future studies could also examine structural variations in port layout and governance. In many ports, including Rotterdam, dedicated barge quays exist alongside multi-purpose deep-sea and feeder berths. Extending the model to simulate this configuration would allow exploration of how modular vessels distribute across quay types, whether they should compete for shared capacity, occupy available gaps, or operate under dedicated scheduling policies. Such analysis would clarify how berth segregation influences efficiency, fairness, and system-wide congestion.

Additional opportunities lie in expanding the simulation to larger-scale network models that capture multi-port or multimodal interactions, linking inland terminals, rail, and barge corridors. These extensions would support policy analysis on regional logistics integration and infrastructure investment planning.

Furthermore, relax the uniform terminal assignment and jointly size module capacity with module-crane inventory. In practice, traffic is skewed toward larger or preferred terminals and some ports operate dedicated barge quays; adding terminal-specific arrival weights and dedicated-quay logic would reveal distributional effects on berth fit, anchorage, and SLA fairness. In parallel, extend the capacity experiments to a grid over module size $\times$ module-crane counts and derive per-terminal sizing curves (e.g., iso-waiting or iso-SLA). A compact design-of-experiments or response-surface approach can map thresholds without exhaustive simulation and support terminal-specific configuration choices.

Finally, empirical validation remains a crucial step. Data from pilot projects or limited-scale deployments of modular or autonomous barges could be used to calibrate service times, manoeuvring behaviour, and coordination dynamics. These datasets would allow cross-verification of simulation outputs and strengthen confidence in the model as a decision-support tool for port authorities and planners.

Contributions to Research, Practice, and the MAGPIE Project

This study contributes to academic literature by bridging the gap between operational simulation and governance-oriented analysis of port automation. It demonstrates how discrete-event simulation can capture coordination dynamics under distributed planning, a perspective seldom addressed in the context of intra-port logistics. For practitioners, the work provides an analytical basis for evaluating modular vessel adoption and highlights that efficiency gains depend as much on information sharing and institutional design as on technology. The model thus serves as a foundation for both future research and strategic experimentation in ports transitioning toward autonomous and digitally coordinated operations.

In addition, the research forms part of the MAGPIE project (*sMArt Green Ports as Integrated Efficient multimodal hubs*), which aims to accelerate the transition towards sustainable, digitally connected, and automated port operations across Europe. Within this context, the developed simulation supports Work Package 5, Demo 6 on autonomous e-barge and transshipment, by providing a system-level assessment of operational effects under varying degrees of modular and autonomous deployment. The find-

ings offer a complementary perspective to the pilot-scale demonstrations, highlighting the coordination and scheduling conditions under which modular vessel concepts can enhance efficiency and resilience in port operations. The developed model may therefore serve as a foundation for future MAGPIE-related analyses that quantify large-scale deployment effects and evaluate digital coordination strategies across the multimodal network.

References

Abourraja, Mohamed, Sebastiaan Meijer, and Jaouad Boukachour (Sept. 2021). “A Model-Driven Design Approach For Ro-Ro And Container Terminals: From Requirements Analysis Down To Simulation Model Implementation”. In: pp. 9–20. DOI: 10.46354/i3m.2021.mas.002.

Alzate, Paola et al. (June 2024). “Operational efficiency and sustainability in smart ports: a comprehensive review”. In: *Marine Systems & Ocean Technology* 19.1-2, pp. 120–131. DOI: 10.1007/s40868-024-00142-z.

APM Terminals (n.d.). *Hoe goed konden we u afgelopen maand helpen?* URL: https://www.apmterminals.com/nl/maasvlakte/newsletter/kpis-mvmfk.

AUTOBarge (2025). *AUTOBarge Project – European Training and Research Network on autonomous barges for smart inland shipping*. URL: https://etn-autobarge.eu/project/.

Bierwirth, Christian and Frank Meisel (2010). “A survey of berth allocation and quay crane scheduling problems in container terminals”. In: *European Journal of Operational Research* 202.3, pp. 615–627. DOI: 10.1016/j.ejor.2009.05.031.

Central Commission for the Navigation of the Rhine (Apr. 2024). *DISRUPTION IN THE CONTAINER SUPPLY CHAIN*. Tech. rep. URL: https://ccr-zkr.org/files/documents/omanalysesthematiques/Report_Congestion_Final_en.pdf.

Cil, Ahmet Yunus, Dini Abdurahman, and Ibrahim Cil (May 2022). “Internet of Things enabled real time cold chain monitoring in a container port”. In: *Journal of Shipping and Trade* 7.1. DOI: 10.1186/s41072-022-00110-z.

Dai, Jam et al. (2008). “Berth Allocation Planning Optimization in Container Terminals”. In: *Supply Chain Analysis: A Handbook on the Interaction of Information, System and Optimization*. Springer US, pp. 69–104. ISBN: 978-0-387-75240-2. DOI: 10.1007/978-0-387-75240-2_4.

Douma, A.M. (Dec. 2008). “Aligning the Operations of Barges and Terminals through Distributed Planning”. PhD Thesis - Research UT, graduation UT. Netherlands: University of Twente. ISBN: 9789036527354. DOI: 10.3990/1.9789036527354.

Douma, Albert M., Peter C. Schuur, and J. Marco J. Schutten (Feb. 2011). “Aligning barge and terminal operations using service-time profiles”. In: *Flexible Services and Manufacturing Journal* 23.4, pp. 385–421. DOI: 10.1007/s10696-011-9080-9.

Dragovic, Branislav, Ernestos Tzannatos, and Nam Kyu Park (Mar. 2017). “Simulation modelling in ports and container terminals: Literature overview and analysis by research field, application area and tool”. In: *Flexible Services and Manufacturing Journal* 29. DOI: 10.1007/s10696-016-9239-5.

Duinkerken, Mark B. et al. (2007). “Comparing transportation systems for inter-terminal transport at the Maasvlakte container terminals”. In: *Container Terminals and Cargo Systems: Design, Operations Management, and Logistics Control Issues*. Springer Berlin Heidelberg, pp. 37–61. ISBN: 978-3-540-49550-5. DOI: 10.1007/978-3-540-49550-5_3.

European Commission Mobility and Transport and Huib van Essen (Dec. 2018). *Sustainable Transport Infrastructure Charging and Internalisation of Transport Externalities*. Tech. rep. URL: https://transport.ec.europa.eu/document/download/8f767f8a-772d-45b4-a2c7-a91b6ba0dbad_en?filename=2018-year-multimodality-external-costs-ce-delft-preliminary-results.pdf.

Flikkema, Maarten et al. (Nov. 2022). “MAGPIE: Towards the European Green port of the Future”. In.

Froeling, D.W. et al. (Dec. 2008). “Container Transferium Rotterdam”. In: pp. 1–6. DOI: 10.1109/INFRA.2008.5439632.

Frojan, Pablo et al. (2015). “The continuous Berth Allocation Problem in a container terminal with multiple quays”. In: *Expert Systems with Applications* 42.21, pp. 7356–7366. ISSN: 0957-4174. DOI: https://doi.org/10.1016/j.eswa.2015.05.018.

GHG Protocol Team (2011). *Corporate Value Chain (Scope 3) Accounting and Reporting Standard*. URL: https://ghgprotocol.org/sites/default/files/standards/Corporate-Value-Chain-Accounting-Reporing-Standard_041613_2.pdf.

Hagaseth, Marianne, Per Håkon Meland, and Egil Wille (Oct. 2023). “Structured description of autonomous inland waterway barge operations”. In: *Journal of Physics Conference Series* 2618.1, p. 012015. DOI: 10.1088/1742-6596/2618/1/012015.

Imai, Akio, Etsuko Nishimura, and Stratos Papadimitriou (2001). “The dynamic berth allocation problem for a container port”. In: *Transportation Research Part B: Methodological* 35.4, pp. 401–417. ISSN: 0191-2615. DOI: https://doi.org/10.1016/S0191-2615(99)00057-0.

IMO (2018). *Autonomous shipping*. URL: https://www.imo.org/en/mediacentre/hottopics/pages/autonomous-shipping.aspx?.

— (May 2021). *Autonomous ships: regulatory scoping exercise completed*. URL: https://www.imo.org/en/mediacentre/pressbriefings/pages/massrse2021.aspx?.

International Maritime Organization (1972). *Convention on the International Regulations for Preventing Collisions at Sea, 1972 (COLREGs)*.

Interreg (2021). *Interreg VB North Sea Region Programme*. URL: https://northsearegion.eu/avatar/about/.

Konings, Rob (2007). “Opportunities to improve container barge handling in the port of Rotterdam from a transport network perspective”. In: *Journal of Transport Geography* 15.6, pp. 443–454. ISSN: 0966-6923. DOI: https://doi.org/10.1016/j.jtrangeo.2007.01.009.

Konings, Rob et al. (2010). “Comparative Strategies for Developing Hinterland Transport by Container Barge: Analysis for Rotterdam and U.S. Ports”. In: *Transportation Research Record* 2166.1, pp. 82–89. DOI: 10.3141/2166-10.

Langen, Peter and Martijn Van der Horst (Mar. 2008). “Coordination in Hinterland Transport Chains: A Major Challenge for the Seaport Community”. In: *Maritime Economics and Logistics* 10, pp. 108–129. DOI: 10.1057/palgrave.mel.9100194.

Maritime Safety Committee (June 2019). *INTERIM GUIDELINES FOR MASS TRIALS*. Tech. rep. URL: https://www.register-iri.com/wp-content/uploads/MSC.1-Circ.1604.pdf.

Neagoe, Mihai et al. (2021). “Using discrete-event simulation to compare congestion management initiatives at a port terminal”. In: *Simulation Modelling Practice and Theory* 112, p. 102362. ISSN: 1569-190X. DOI: https://doi.org/10.1016/j.simpat.2021.102362.

Negenborn, Rudy et al. (Feb. 2023). “Autonomous ships are on the horizon: here’s what we need to know”. In: *Nature* 615, pp. 30–33. DOI: 10.1038/d41586-023-00557-5.

Neurohr, Christian et al. (Aug. 2025). *Towards Efficient Certification of Maritime Remote Operation Centers*. DOI: 10.48550/arXiv.2508.00543.

Nextlogic (2024). *Hoe het werkt*. URL: https://www.nextlogic.nl/zowerkthet/.

OECD (2012). *Strategic Transport Infrastructure Needs to 2030*. Tech. rep. URL: https://www.oecd.org/content/dam/oecd/en/publications/reports/2012/02/strategic-transport-infrastructure-needs-to-2030_g1g13d1c/9789264114425-en.pdf.

Ottjes, Jaap A. et al. (2007). “Simulation of a multiterminal system for container handling”. In: *Container Terminals and Cargo Systems: Design, Operations Management, and Logistics Control Issues*. Springer Berlin Heidelberg, pp. 15–36. ISBN: 978-3-540-49550-5. DOI: 10.1007/978-3-540-49550-5_2.

Pachakis, Dimitrios and Anne Kiremidjian (Sept. 2003). “Ship Traffic Modeling Methodology for Ports”. In: *Journal of Waterway Port Coastal and Ocean Engineering-asce - J WATERW PORT COAST OC-ASCE* 129. DOI: 10.1061/(ASCE)0733-950X(2003)129:5(193).

Pallis, Athanasios and Jean-Paul Rodrigue (2024). “Chapter 6.2 – Port Efficiency”. In: *Port Economics, Management and Policy*. URL: https://porteconomicsmanagement.org/pemp/contents/part6/port-efficiency/.

Port of Rotterdam (July 2023). *Port Call Process | PortOfRotterdam*. URL: https://www.youtube.com/watch?v=YSfVp_b1Yv4.

— (2025a). *Barge Performance Monitor*. URL: https://www.portofrotterdam.com/en/logistics/connections/intermodal-transportation/inland-shipping/barge-performance.

— (2025b). *Port Performance*. URL: https://www.portofrotterdam.com/en/logistics/port-performance.

— (n.d.[a]). *Containers - Enabling seamless supply chains*. URL: https://www.portofrotterdam.com/nl/logistiek/lading/containers.

Port of Rotterdam (n.d.[b]). *Optimising inland container shipping chain*. URL: https://www.portofrotterdam.com/en/logistics/connections/intermodal-transportation/inland-shipping/optimising-inland.

Port of Rotterdam et al. (n.d.). *CONTAINER TERMINALS AND DEPOTS IN THE ROTTERDAM PORT AREA*. Tech. rep. URL: https://www.portofrotterdam.com/sites/default/files/2021-06/container-terminals-and-depots-in-the-rotterdam-port-area.pdf.

Port of Rotterdam Authority (Nov. 2023). *The Port of Rotterdam Authority officially commissions Container Exchange Route*. URL: https://www.portofrotterdam.com/en/news-and-press-releases/the-port-of-rotterdam-authority-officially-commissions-container-exchange.

— (2025a). *Container Exchange Route (CER)*. URL: https://www.portofrotterdam.com/en/building-port/ongoing-projects/container-exchange-route-cer.

— (2025b). *Maasvlakte railway development*. URL: https://www.portofrotterdam.com/en/building-port/ongoing-projects/maasvlakte-railway-development.

— (2025c). *Rail transport — PortShuttle Rotterdam*. URL: https://www.portofrotterdam.com/en/logistics/connections/intermodal-transportation/rail-transport.

Portbase (2025). *Melding container Achterland*. URL: https://www.portbase.com/services/melding-container-achterland/.

Prasad, Dilip et al. (Feb. 2017). "Maritime situational awareness using adaptive multi-sensor management under hazy conditions". In: DOI: 10.48550/arXiv.1702.00754.

Puleri, Marzio et al. (2023). *A look at digital twins in port operations*. URL: https://www.ericsson.com/en/blog/2020/12/digital-twins-port-operations.

Register, Lloyd's (Dec. 2017). *Cyber-enabled ships: ShipRight procedure assignment for cyber descriptive notes for autonomous & remote access ships*. Tech. rep. URL: https://www.scribd.com/document/449320742/MO-Cyber-Enabled-Ships-ShipRight-Procedure-V2-0-201712?.

Rødseth, Ørnulf, Lars Andreas Wennersberg, and Håvard Nordahl (July 2022). "Levels of autonomy for ships". In: *Journal of Physics: Conference Series* 2311, p. 012018. DOI: 10.1088/1742-6596/2311/1/012018.

Rødseth, Ørnulf Jan and Lars Andreas Lien Wennersberg (2023). "A Criticism of Proposed Levels of Autonomy for MASS". In: *Proceedings of the 33rdEuropean Safety and Reliability Conference (ESREL2023)*. DOI: 10.3850/978-981-18-8071-1\{_}driver.

Rødseth, Ørnulf Jan, Lars Andreas Lien Wennersberg, and Håvard Nordahl (2023). "Improving safety of interactions between conventional and autonomous ships". In: *Ocean Engineering* 284, p. 115206. ISSN: 0029-8018. DOI: https://doi.org/10.1016/j.oceaneng.2023.115206.

Schiaretti, Matteo, Linying Chen, and Rudy Negenborn (2017). "Survey on Autonomous Surface Vessels: Part I - A New Detailed Definition of Autonomy Levels". In: *Lecture Notes in Computer Science* 10531, pp. 219–233. DOI: 10.1007/978-3-319-68496-3_15.

Shobayo, Peter and Edwin Van Hassel (June 2019). "Container barge congestion and handling in large seaports: a theoretical agent-based modeling approach". In: *Journal of Shipping and Trade* 4.1. DOI: 10.1186/s41072-019-0044-7.

Steenken, Dirk, Stefan Voß, and Robert Stahlbock (2004). "Container terminal operation and operations research—a classification and literature review". In: *OR Spectrum* 26.1, pp. 3–49. DOI: 10.1007/s00291-003-0157-z.

Truong, Nguyen et al. (2016). "A survey on trust computation in the Internet of Things". In: *Journal of Korean Institute of Communications and Information Sciences* 33, pp. 10–27. URL: https://www.researchgate.net/publication/316042146_A_Survey_on_Trust_Computation_in_the_Internet_of_Things.

UN Office for Sustainable Development (2015). *Sustainable Development Goals (SDGs)*. URL: https://unosd.un.org/content/sustainable-development-goals-sdgs.

Van der Horst, Martijn et al. (Jan. 2019). "Coordination problems in container barging in the port of Rotterdam: an institutional analysis". In: *Transportation Planning and Technology* 42, pp. 1–13. DOI: 10.1080/03081060.2019.1565164.

Veeke, Hans P.M., Gabriël Lodewijks, and Jaap A. Ottjes (Jan. 2008). *The Delft Systems approach*. DOI: 10.1007/978-1-84800-177-0. URL: http://doi.org/10.1007/978-1-84800-177-0.

Vidan, Pero et al. (Apr. 2019). "Autonomous Systems & Ships - Training and Education on Maritime Faculties". In: URL: https://www.researchgate.net/publication/354922331_Autonomous_Systems_Ships_-Training_and_Education_on_Maritime_Faculties.

World Bank (2007). *Port Reform Toolkit*. Tech. rep. Public-Private Infrastructure Advisory Facility (PPIAF). URL: https://ppp.worldbank.org/public-private-partnership/library/port-reform-toolkit-ppiaf-world-bank-2nd-edition.

Yu, Tai-Tuck, James Scanlan, and Gary Wills (Sept. 2007). "Agent-Based and Discrete-Event Modeling: A Quantitative Comparison". In: *Collection of Technical Papers - 7th AIAA Aviation Technology, Integration, and Operations Conference* 2. DOI: 10.2514/6.2007-7818.

A

Research Paper

The Impact of Autonomous Intra-Terminal Barge Concepts - A Case Study of the Port of Rotterdam

Maartje Janszen
MSc Mechanical Engineering
Delft University of Technology
Delft, The Netherlands

Dr. B. Atasoy
Faculty of Mechanical Engineering
Delft University of Technology
Delft, The Netherlands

Dr. J. Jovanova
Faculty of Mechanical Engineering
Delft University of Technology
Delft, The Netherlands

Prof. Dr. R. A. Zuidwijk
Rotterdam School of Management
Erasmus University Rotterdam
The Netherlands

Dr. ir. T. Verduijn
Port of Rotterdam – MAGPIE
Rotterdam, The Netherlands

***Abstract*— European hub ports rely on inland barges to connect deep-sea terminals with the hinterland. Fragmented local rules create short, irregular berth windows and strain reliability. Modular inland vessels can split into autonomous shuttles and exploit these windows more effectively. System-level effects of such concepts remain unclear. A discrete event simulation of a multi-terminal Rotterdam hub is developed. It includes rotation-aware barge itineraries and a rule-based berth admission scheme. Two crane pools are represented: general quay cranes and a dedicated module crane pool. A parameterised modular vessel model supports direct side-by-side tests. Experiments vary fleet mix, sea freight demand, crane count and speed, and module capacity. An early adoption case tests a 100:10 mix. Mixed fleets outperform the conventional baseline. Intermediate mixes, 50:50 and 25:75, shorten inland turnaround and raise total throughput. Deep-sea and feeder times remain close to baseline. At high modular shares, the module interface becomes decisive. At least two module cranes per terminal are needed to avoid single-server queues. Extra cranes mainly add headroom. With a fixed annual modular volume, modules of 24 to 36 TEU minimise staging and waiting. Handling speed is secondary at moderate uptake. Early adoption increases total throughput by about three to six percent. It also lifts inland share by two to four percent. Gains arise from better length–time packing at shared quays, not from higher sustained occupancy. Berth access and general quay cranes therefore remain governing resources. The results support a hybrid transition with distributed scheduling and a right-sized module interface.**

I. Introduction

Inland waterway transport is central to the efficiency and sustainability of European port systems [1]. Large multi-terminal hubs, such as the Port of Rotterdam, rely on barge services to connect deep-sea terminals with inland destinations. These services help reduce road congestion and support modal shift objectives [2], [3]. Yet barge operations are characterised by fragmented decision-making and heterogeneous local rules. Strong interdependencies between terminals compound these effects and challenge service reliability and port-wide coordination [2]. These challenges are acute when quay windows are short and irregular, since small timing mismatches compound across multi-call rotations [4].

A rich literature on port operations covers berth allocation, crane scheduling, and port-wide traffic generation. Simulation is widely used to represent stochastic processes that analytical models cannot easily capture [5]–[7]. However, most applications centre on sea-going vessels. Inland barges differ in two important ways. They often perform multiple calls within a port rotation, so delays at one terminal propagate to subsequent calls [4]. They also face short and irregular windows at the quay, which makes fine temporal coordination decisive for throughput and reliability [3], [8].

Four gaps motivate this study. First, there is little port-wide evidence on how barge coordination choices change waiting and service times, window adherence, and berth occupancy across multiple terminals. Second, traffic generation methods developed for deep-sea calls are not well suited to intra-port, multi-call barge rotations where delays propagate across the round [9]. Third, comparative side-by-side evaluations of coordination rules in multi-terminal landlord settings remain scarce [10]. Fourth, the operational consequences at port scale of emerging vessel concepts, including autonomous and modular barges that split into smaller units, are largely unquantified [11]–[15].

This paper addresses these gaps in four ways. It develops a rotation-aware traffic generator that constructs multi-call barge itineraries and propagates realised delays across subsequent calls [4], [9]. It introduces a rule-based berth-admission and resource-attachment scheme that reflects landlord-port practice, with explicit checks on effective quay length, terminal opening windows, and the availability of general and module crane pools. It links job size, attached cranes, service time, and occupied effective length into a unified key performance indicator (KPI) formulation that covers both time-based occupancy and length–time coverage [1]. It also provides a parameterised representation of autonomy-enabled modular vessels so that conventional and modular fleets can be analysed within a single framework [11], [13]–[15]. The framework is designed to integrate with existing port simulators. Where data are abundant, documented defaults and sensitivity bands allow analysts to assemble defensible traffic and control scenarios [9].

In the following sections, the background situates barge coordination in multi-terminal landlord ports and

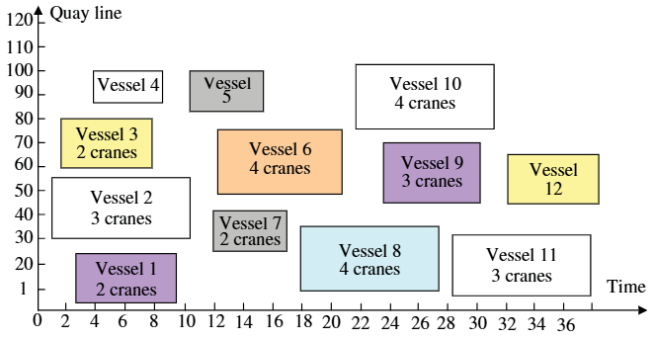

Fig. 1. Berth allocation with crane assignment in a length–time diagram [16]. Rectangles represent vessel calls, labels indicate assigned cranes; gaps reflect fragmented quay capacity over time.

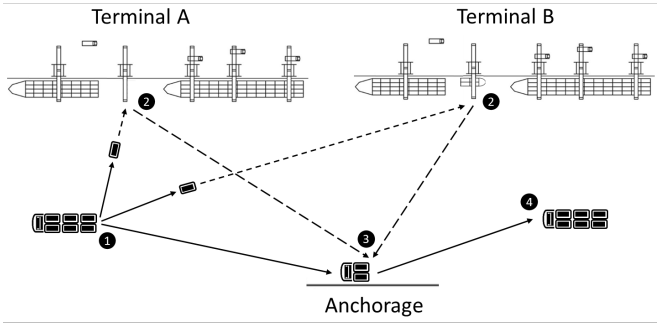

Fig. 2. Modular splitting in three phases. Phase 1: mother arrives and detaches modules. Phase 2: modules make short, parallel calls handled by a dedicated module-crane pool while the mother waits. Phase 3: modules rendezvous with the mother, and Phase 4: depart as a convoy.

reviews planning rules and digital appointment systems. The methodology describes the simulation framework, traffic generation, berth-admission logic, scenarios, and performance indicators. The results present system-level effects of modular adoption under varying equipment and demand. The discussion interprets mechanisms, implementation choices, and generalisability. The conclusion summarises implications for port and terminal decision makers and outlines directions for future work.

II. Background

Large multi-terminal ports allocate scarce quay length and quay-crane (QC) capacity across vessels that arrive over time. In landlord settings, decision rights are distributed across the port authority, terminal operators, and carriers, which makes coordination a system problem rather than a single-facility problem [1]. Inland barge rotations often include multiple terminal calls; delays at one call propagate to the next legs, affecting reliability at port scale [4], [9]. Appointment systems and integrated planning platforms (e.g., Portbase/Nextlogic in Rotterdam) were introduced to improve predictability and reduce waiting [2], [3], [8].

A length–time berth plan reveals how well calls pack along the quay. Fig. 1 shows a plan with explicit quay crane assignments. The key insight is fragmentation: the white spaces between rectangles are micro-intervals and length remainders that are hard to use when calls are short or irregular. Changing the number of quay cranes alters completion time, which in turn determines whether a call fits into residual windows. This links berth allocation directly to crane scheduling and motivates strategies that create short, schedulable calls.

Within this context, the paper analyses a coordination concept based on modular inland vessels that can split into autonomous shuttles (Fig. 2). A mother barge detaches one or more modules. These berth independently for brief exchanges handled by dedicated module cranes, then rendezvous before departure. By converting part of the inland workload into many short, schedulable calls, modular traffic targets residual pockets in the berth plan while preserving sea-going priorities. The central question is whether short, parallelisable calls improve waiting, berth fit (length–time coverage), and throughput at multi-terminal scale compared with conventional single-hull rotations. The analysis identifies the configurations under which these gains materialise and outlines implications for distributed planning and traffic management in landlord ports.

III. Methodology

This section sets out the simulation tools used to represent container handling in a multi-terminal port. A discrete-event approach is adopted to capture event timing, resource contention, and variability in arrivals and handling. The model comprises a high-level framework, a traffic generator that produces barge movements and terminal calls, and a berth-admission and service logic that governs access to quay length and cranes. The configuration is verified and validated against observed ranges before experiments are evaluated with KPIs.

A. Simulation Framework and Scope

The simulator represents a container hub with three vessel classes: deep-sea, feeder, and inland. The inland class is implemented in two configurations. The conventional configuration models single-hull barges. The modular configuration models a mother vessel that can detach one or more modules which berth independently and later rejoin the mother before departure. Terminals have finite quay length and opening windows, and two resource pools are distinguished: general quay cranes that handle sea-side and barge work, and module-handling cranes that are reserved for modules. Experiments use multiple replications of a four-week horizon, consistent with container-terminal simulation practice [5], [7].

B. System Configuration

The model represents a stylised Maasvlakte hub. It includes four vessel classes: deep-sea, feeder, inland barge and autonomous module. Quays are shared-use rather than class-dedicated, so any of the four classes may berth at the same quay. Access is governed by effective length, defined as hull length plus fore and aft buffers. Two

crane pools are represented: general quay cranes and a dedicated module-crane pool. The network comprises nautical legs between the entrance, anchorage and terminals. An anchorage buffers arrivals when berths are unavailable.

C. Arrivals, Rotations, and Modular Missions

Class-specific stochastic processes generate deep-sea, feeder, and inland arrivals at observed volumes. Inland services are created as rotations, namely sequences of terminal calls linked by stochastic sailing and manoeuvring times; realised start and completion at one call set the ready time for the next leg, so early and late starts propagate across the round, consistent with practice [4], [9]. A configured modular share converts a proportion of inland moves into module calls: a mother vessel anchors in the port area, child modules sail independently to their assigned terminals for short exchanges using the module-crane pool, then return to rendezvous before the mother departs (concept in Fig. 2). Per-call workload is sampled by class; vessel attributes such as length and speed determine effective quay use and sailing times.

D. Berth Admission, Sequencing, and Crane Attachment

Terminals operate mixed-use quays with a class-first admission order that reflects common practice [5], [6], [17]: deep-sea, then feeder, then inland (barge and modules share the inland tier). A call is admitted only when three gate checks all pass: (i) a non-overlapping berth segment of sufficient effective quay length is free, (ii) the terminal is within its opening window, and (iii) the crane rule for the class can be met. Conventional inland and feeder calls require at least one free general quay crane; modules require a module crane; deep-sea calls may berth without immediate crane attachment, after which general cranes attach when available. The decision flow is shown in Fig. 3.

Within a class, sequencing follows standard rules with simple tie-breakers. Deep-sea and feeder select the earliest ready call within the operational day, ties by first-come, first-served (FCFS), protecting contracted service levels. Inland selection uses weighted shortest processing time (WSPT), where the weight captures class or contractual priority and the processing time is the expected service time (moves divided by expected concurrent cranes); ties break by earliest ready time. Sensitivity runs with pure FCFS and pure SPT reproduce decentralised and congestion-driven practices and do not change the qualitative conclusions [5], [6].

Lower-priority inland traffic is also screened by a finish-sooner look-ahead: a lower-priority call is admitted only if it is predicted to complete before the next higher-priority call's earliest feasible start, up to a small safety buffer Δ that covers approach and set-up. Let $T_{\text{low}}^{\text{finish}}$ be the predicted completion time if admitted now and $T_{\text{high}}^{\text{start}}$ the earliest feasible start of the higher-priority call (accounting for sailing, berth, and crane constraints). Admission requires

$$T_{\text{low}}^{\text{finish}} + \Delta \leq T_{\text{high}}^{\text{start}}. \quad (1)$$

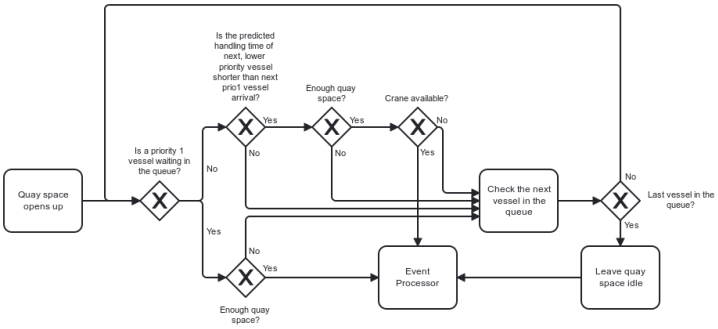

Fig. 3. Berth-admission decision flow used in the simulation.

If the inequality fails, the berth is held empty to protect the imminent higher-priority start, and the queue is re-evaluated at the next check.

On berthing, cranes attach from the relevant pool up to three limits: the class cap, free cranes in that pool, and spacing along the quay. Cranes can be rebalanced when neighbouring jobs complete, yet ongoing jobs are not pre-empted. After service, conventional inland vessels sail to the next terminal in the rotation; modules return to the mother's waiting point, and the mother departs once all modules have rejoined. This rendezvous logic influences the module-mother's port time but leaves terminal admission unchanged.

E. Scenarios

To examine mechanisms and limits, multiple experiments are designed.

1) Module share: the fraction of inland moves executed by modules rather than conventional barges, ranging from 100:0 (conventional:module) to 0:100.
2) Sea-freight demand band: a proportional scaling of sea-side load around the baseline to test robustness under lighter and tighter pressure on quay and cranes.
3) Module interface configuration: three variables at the module interface — (a) module cranes per terminal, (b) per-module capacity at fixed annual module TEU, and (c) module-crane handling rate.
4) Near-term implementation: an incremental adoption case with a small module share under current infrastructure to gauge early gains and practical requirements.

F. Performance Indicators

Performance is tracked at vessel, terminal, and port levels. Vessel-level metrics focus on turnaround times consisting of waiting, service, and sailing times by class. Terminal-level metrics include time-based berth occupancy, effective utilisation computed from length–time coverage of calls, and crane utilisation on average and at peak times. Port-level metrics include processed TEU, and anchorage shares by class to measure congestion.

G. Verification and Validation

Model code and event logic are verified with unit and consistency checks. Face validation is performed with subject-matter experts. Aggregate outputs are compared with observed ranges for arrivals, handling rates, and berth occupancy to check plausibility at system level.

IV. Results

This section reports outcomes of the four experiments. Unless noted otherwise, values are means over multiple independent replications of a four-week horizon.

A. Fleet Mix: Conventional Barges and Modules

Mixed fleets shorten inland turnaround and lift throughput while seagoing service remains broadly stable. The pattern in Table I follows from two mechanisms. First, converting part of the inland workload into many short module calls increases the chance that arrivals coincide with residual windows in the length–time plan, so inland work is admitted more often without displacing higher-priority deep-sea traffic. Barge time at the quay therefore changes little, but waiting collapses, and total throughput rises. Second, modules draw on a dedicated crane pool, so inland handling load shifts away from general quay cranes. With fewer attach lags and micro-interruptions on ocean calls, deep-sea and feeder turnaround falls slightly yet consistently. At terminal level this appears as lower average effective berth utilisation with peaks still near full use: short calls pack fragmented windows more efficiently, raising processed TEU per metre–hour while spending less time in high-occupancy states. When modules dominate, constraints move upstream of the quay: the mother's departure is tied to the slowest child and the module-crane interface saturates in scarce windows, shifting queues to anchorage and flattening the throughput gains relative to intermediate mixes. Overall, the 50:50 and 25:75 configurations capture most of the benefit by combining better length–time packing with a favourable redistribution of crane load, without triggering long rendezvous tails or sustained saturation at the module interface.

B. Effects of Different Demands

Sea-freight intensity was scaled to 80%, 100%, and 120% while keeping call mix, priorities, berth windows, and crane inventories fixed. Two mixed fleets (50:50 and 25:75) are compared with the conventional baseline (100–0) to test robustness under lighter and tighter sea-side pressure. Table II summarises the demand sensitivity of key indicators. Throughput elasticity is essentially unchanged across mixes: a 20% sea-freight increase raises total TEU by about +13–15% in every case. The performance gap appears in delay sensitivity. Barge turnaround in 100–0 rises by +16.1 h between 80% and 120% demand, but only +8.3 h in 50–50 and +3.2 h in 25–75; at 120% demand, absolute barge times in both mixed systems remain well below the conventional system even at 80%. Deep-sea and feeder times drift up only slightly (consistent with priority and ship-to-shore capacity).

TABLE I

Summary of relative changes in KPIs compared with the baseline (100:0). Positive values denote increases, negative values reductions.

KPI (% vs 100:0)	75:25	50:50	25:75	0:100
Vessel-level				
Barge turnaround time	−29.1	−53.7	−68.8	N/A
Deep-sea turnaround	−2.5	−6.7	−8.8	−10.2
Feeder turnaround	−3.9	−11.0	−14.9	−16.2
Barge waiting time	−38.5	−71.5	−92.2	N/A
Terminal-level				
Avg. eff. berth util.	−9.4	−19.8	−34.2	−51.5
Avg. cranes used (total)	−7.6	−15.6	−26.5	−39.8
Port-wide				
Total throughput (TEU)	+7.1	+12.3	+16.7	+17.7
Inland throughput (TEU)	+32.2	+58.6	+80.0	+86.8
Inland share (%)	+4.6	+8.1	+10.6	+11.6

On the terminal side, average effective berth utilisation increases with demand yet remains below saturation in all cases, and the mixed fleets achieve the same sea-side load with lower overall crane duty at Terminal B. In short, modular uptake does not change how volumes scale with demand, but it markedly dampens how delays grow when the system is tight; 50–50 captures most of this benefit with a moderate module footprint, while 25–75 offers the strongest protection of inland performance.

TABLE II

KPIs under ±20% sea-freight demand. Values at 80% and 120% show change vs the 100% case in parentheses.

Demand	Total throughput ($\times 10^6$ TEU)	Inland share (%)	Barge turnaround (h)	Average berth util. (%)	Avg. crane duty factor[1] (%)
			100:0		
80%	0.95 (−15.2%)	24.4	24.5 (−9.2)	76.5 (−3.8)	67.0 (−5.0)
100%	1.12	19.7	33.7	80.3	67.0
120%	1.29 (+15.2%)	16.3	40.6 (+6.9)	81.6 (+1.3)	68.0 (+1.0)
			50:50		
80%	1.07 (−15.1%)	32.7	12.1 (−3.5)	57.8 (−6.5)	50.0 (−6.0)
100%	1.26	27.8	15.6	64.3	56.0
120%	1.43 (+13.5%)	24.2	20.4 (+4.8)	70.3 (+6.0)	62.0 (+6.0)
			25:75		
80%	1.12 (−14.5%)	35.4	9.3 (−1.2)	45.9 (−6.9)	43.0 (−6.0)
100%	1.31	30.3	10.5	52.8	49.0
120%	1.48 (+13.0%)	26.7	12.5 (+2.0)	59.8 (+7.0)	55.0 (+6.0)

C. Module Interface Configuration

This section varies three elements of the module interface: crane count per terminal, handling speed, and per-module capacity. Effects are shown for the 50:50 and 25:75 mixes.

[1]Average cranes in use divided by cranes installed.

1) Case: Crane Configuration: Two results dominate. First, there is a clear threshold in module handling. With only one module crane per terminal, module waiting explodes and mother calls inherit very long turnarounds, which suppresses inland and total throughput. Adding a second module crane collapses module queues and restores throughput; a third mainly adds headroom, and an effectively unbounded module pool does not raise total TEU beyond the +2 or +3 cases, which shows the ceiling lies elsewhere. Second, quay access and general cranes remain governing. Deep-sea and feeder turnarounds move only slightly across configurations, and Terminal B operates in a high but non-saturated regime within each mix. Removing general cranes chiefly hurts barges; adding module cranes shifts inland work away from the general pool without increasing peak general-crane load. Anchorage patterns mirror this: in 25-75 the modules absorb staging when the module pool is undersized, while in 50-50 barge anchorage is higher overall but remains modest once two module cranes are available.

TABLE III
25:75 MIX: AVERAGE MODULE WAITING AND TOTAL THROUGHPUT BY CRANE CONFIGURATION.

Configuration	Module waiting (h)	Total throughput ($\times 10^6$ TEU)
-2 cranes	0.68	1.30
-1 crane	143.70	1.18
+1 m crane	143.20	1.18
+2 m cranes	0.60	1.31
+3 m cranes	0.19	1.31
∞ m cranes	0.10	1.30
Baseline (100:0)	-	1.12

TABLE IV
50:50 MIX: AVERAGE MODULE WAITING AND TOTAL THROUGHPUT BY CRANE CONFIGURATION.

Configuration	Module waiting (h)	Total throughput ($\times 10^6$ TEU)
-2 cranes	0.39	1.25
-1 crane	26.59	1.24
+1 m crane	26.50	1.25
+2 m cranes	0.20	1.26
+3 m cranes	0.12	1.26
∞ m cranes	0.10	1.26
Baseline (100:0)	-	1.12

In practice, at least two module cranes per terminal are needed for the high-module 25-75 configuration to avoid a single-server bottleneck. In the balanced 50-50 mix, one to two module cranes already keep module waiting low under the tested demand; further cranes mainly provide robustness rather than higher throughput.

2) Case: Handling Speed: Slower module handling leaves barges, deep-sea and feeders broadly unchanged, while it drives a sharp rise in module staging and waiting once the dedicated pool becomes heavily used. In the high-module 25–75 mix, module waiting per child increases from well under an hour at 1.5 min/TEU to more than 140 h at 3.5 min/TEU, and total throughput erodes from about 1.31×10^6 to 1.18×10^6 TEU (Table V). The balanced 50–50 mix is far more forgiving: waiting grows, but port-wide throughput stays near 1.25×10^6 TEU across the range (Table VI). Terminal B operates in a high-utilisation yet non-saturated regime for all speeds, with load shifting from general to module cranes as handling slows. The mechanism is therefore not longer service at the quay, but queueing upstream of the quay once the module interface saturates. Moderate modular uptake preserves the port-wide gains even with slower module cranes, whereas high modular shares require a fast, well-dimensioned module pool to avoid severe waiting and loss of inland throughput.

TABLE V
25-75 MIX: MODULE WAITING PER CHILD AND TOTAL THROUGHPUT BY MODULE-HANDLING SPEED.

	sp 1.5	sp 2.0	sp 2.5	sp 3.0	sp 3.5
Module waiting per child (h)	0.60	7.23	65.02	109.98	144.16
Total throughput ($\times 10^6$ TEU)	1.31	1.30	1.25	1.21	1.18

TABLE VI
50-50 MIX: MODULE WAITING PER CHILD AND TOTAL THROUGHPUT BY MODULE-HANDLING SPEED.

	sp 1.5	sp 2.0	sp 2.5	sp 3.0	sp 3.5
Module waiting per child (h)	0.20	0.48	0.85	2.87	27.57
Total throughput ($\times 10^6$ TEU)	1.26	1.25	1.25	1.26	1.24

3) Case: Module Capacity: Changing module size repackages a given inland workload into many short visits or fewer, heavier pulses. With annual module TEU held constant, intermediate capacities of 24–36 TEU minimise module staging and waiting in both mixes, while 12 TEU creates many intrusive arrivals that must be staged and 48 TEU creates longer visits that need larger windows. In the 25–75 mix, module waiting falls from about 4.9 h at 12 TEU to around 0.5 h at 24–36 TEU, then rises again to about 2.7 h at 48 TEU (Table VII). In the 50–50 mix the same pattern appears at lower levels (Table VIII). Berth-fit and general-crane loading at Terminal B remain on a stable plateau across capacities, so the effect is driven by staging and packing rather than by sustained local saturation. Port-wide volumes are similar for 12–36 TEU within each mix; the 48 TEU cases lift throughput because they also carry more module TEU. In practice, capacities around 24–36 TEU offer a robust compromise: they keep module turnaround short, avoid excessive anchorage, and preserve stable service for barges and sea-going traffic, while very small or very large modules increase staging pressure without relieving quay constraints.

TABLE VII

25-75 MIX: MODULE WAITING PER CHILD AND TOTAL THROUGHPUT BY MODULE CAPACITY.

	12 TEU	24 TEU	36 TEU	48 TEU
Module waiting per child (h)	4.92	0.57	0.46	2.71
Total throughput ($\times 10^6$ TEU)	1.30	1.31	1.30	1.41

TABLE VIII

50-50 MIX: MODULE WAITING PER CHILD AND TOTAL THROUGHPUT BY MODULE CAPACITY.

	12 TEU	24 TEU	36 TEU	48 TEU
Module waiting per child (h)	0.26	0.20	0.25	0.38
Total throughput ($\times 10^6$ TEU)	1.26	1.26	1.26	1.33

4) Synthesis Configuration: Taken together, the configuration cases point to clear thresholds and diminishing returns. At high modular shares, at least two dedicated module cranes per terminal are required to avoid single-server queues; in a balanced 50:50 mix, one to two cranes already keep waiting low. Handling speed is secondary at moderate uptake but becomes critical once the module pool approaches saturation. With annual modular TEU held constant, mid-range capacities around 24–36 TEU minimise staging and waiting, while very small or very large units increase anchorage. Beyond these thresholds, extra cranes or faster cycles mainly add robustness rather than higher port-wide throughput, because quay access and the general crane pool remain governing. The common mechanism is improved temporal fit at the quay: short, schedulable calls that avoid single-server effects reduce anchorage and raise effective length–time coverage. This sets up the early-adoption results, where a modest modular share yields gains provided the module interface is dimensioned to these thresholds.

D. Early Modular Adoption Scenario

Introducing a modest modular share (100:10) raises total throughput by about 3 to 6% and lifts the inland share by roughly 2.4 to 3.8 percentage points, without materially affecting seagoing service levels (Table IX). Deep-sea and feeder port times stay within a narrow band around the baseline (Table X). Their average waiting times are around an hour or less in all settings, with one slow-speed sharing case slightly above one hour (the 0QC, 3.0 min/container run in Table XI). This indicates that priority and ship-to-shore capacity continue to protect ocean calls.

Deep-sea and feeder port times stay within a narrow band around the baseline (Table X). Their average waiting times are around an hour or less in all settings, with one slow-speed sharing case slightly above one hour (the 0QC, 3.0 min/container run in Table XI). This indicates that priority and ship-to-shore capacity continue to protect ocean calls.

TABLE IX

PORT-WIDE THROUGHPUT AND INLAND SHARE FOR THE 100:0 BASELINE AND EARLY MODULAR ADOPTION FUTURES (100:10). TOTAL THROUGHPUT IN $\times 10^6$ TEU; INLAND THROUGHPUT (BARGE + MODULE) IN $\times 10^5$ TEU.

Scenario	Total ($\times 10^6$)	Change vs 100:0	Inland ($\times 10^5$)	Inland share (%)
100:0 baseline	1.12	–	2.20	19.7
100:10, dedicated module cranes (1.5 min/container)				
+1MC	1.17	+4.4%	2.66	22.8
+2MC	1.17	+4.4%	2.66	22.7
100:10, modules on quay cranes (1.5 min/container)				
0QC	1.16	+3.2%	2.55	22.1
+1QC	1.17	+4.5%	2.68	22.9
+2QC	1.19	+5.8%	2.79	23.5
100:10, modules on quay cranes (3.0 min/container)				
0QC	1.16	+3.2%	2.55	22.1
+1QC	1.17	+4.9%	2.67	22.7
+2QC	1.18	+5.4%	2.76	23.4

TABLE X

AVERAGE PORT TIME (H) BY VESSEL TYPE FOR THE 100:0 BASELINE AND EARLY MODULAR ADOPTION FUTURES (100:10). "+1MC"/"+2MC": DEDICATED MODULE CRANES AT 1.5 MIN/CONTAINER; "0QC/+1QC/+2QC": ADDITIONAL QUAY CRANES WHEN MODULES ARE HANDLED ON THE GENERAL CRANE POOL AT 1.5 OR 3.0 MIN/CONTAINER.

Scenario	Barge (h)	Deep-sea (h)	Feeder (h)	Module mother (h)
100:0 baseline	33.65	28.51	15.39	–
100:10, dedicated module cranes (1.5 min/container)				
+1MC	33.89	28.48	15.36	1.24
+2MC	33.40	28.72	15.69	1.46
100:10, modules on quay cranes (1.5 min/container)				
0QC	36.77	29.11	15.78	26.24
+1QC	29.14	28.22	15.08	18.72
+2QC	23.64	27.48	14.49	14.02
100:10, modules on quay cranes (3.0 min/container)				
0QC	33.16	28.91	15.76	23.79
+1QC	29.30	28.03	14.95	20.08
+2QC	25.08	27.36	14.26	15.79

The main effects sit in inland traffic and depend on how cranes are organised. If modules must share the general quay-crane pool without extra capacity (0QC), crane constraints at inland terminals become frequent, barge waiting increases relative to the baseline, and module children incur long queues before service (Table XI). Dedicating module cranes (+1MC, +2MC) keeps barge waiting close to the baseline and reduces module waiting to near zero. Adding general cranes when modules run on the quay-crane pool (+1QC, +2QC) brings barge waiting down from about 25–26 h to roughly 20 h and then 15 h, while module waiting falls to single-digit hours.

Matching module handling speed to barge speed

TABLE XI

AVERAGE WAITING TIME (H) AT TERMINALS BY VESSEL TYPE FOR THE 100:0 BASELINE AND EARLY MODULAR ADOPTION FUTURES (100:10). CONVENTIONS AS IN TABLE X.

Scenario	Barge (h)	Deep-sea (h)	Feeder (h)	Module child (h)
100:0 baseline	23.62	0.74	0.36	–
100:10, dedicated module cranes (1.5 min/container)				
+1MC	24.15	0.73	0.30	0.13
+2MC	24.64	0.89	0.40	0.21
100:10, modules on quay cranes (1.5 min/container)				
0QC	25.59	0.92	0.47	9.52
+1QC	19.80	0.89	0.44	7.07
+2QC	14.64	0.71	0.39	5.14
100:10, modules on quay cranes (3.0 min/container)				
0QC	23.29	1.08	0.56	8.96
+1QC	19.19	0.71	0.39	7.18
+2QC	16.00	0.65	0.23	5.63

(3.0 min/container) leaves the port-wide gains essentially intact. The total-throughput increases remain within the same 3 to 6% range across +1QC and +2QC, compared with their 1.5 min/container counterparts (Table IX). This shows that availability and organisation of the crane interface matter more than faster per-container cycles at these adoption levels.

At terminal scale, operations tighten but do not saturate. At the representative inland terminal, average effective berth utilisation remains near the baseline plateau (about 50 to 52% on average, with short peaks near full use), which indicates that early adoption mainly fills residual windows rather than driving sustained saturation. Crane loading shifts rather than explodes: compared with a baseline average-to-installed duty near two thirds, the 0QC sharing case lifts duty substantially, while +1QC and +2QC restore slack toward baseline levels. Bottleneck counters reflect the same transition: crane-limited time rises in the 0QC case and recedes as extra capacity is provided, while quay-limited periods remain negligible.

Port-wide, all 100:10 variants add inland moves on top of stable ocean volumes. With dedicated module cranes (+2MC), total TEU increases by about 4.4% and the inland share reaches about 22.7%. With modules on the quay-crane pool, +1QC and +2QC yield increases of about 4.5 to 5.8% and inland shares around 22.9 to 23.5% (Table IX).

At terminal scale the picture is of tighter but non-saturated operations. At the representative inland Terminal B, average effective berth utilisation stays close to the baseline plateau (about 50–52% on average with short peaks near full use), indicating that early adoption mainly fills residual windows rather than pushing berth-metre-hours towards sustained saturation. Crane loading shifts rather than explodes: compared with the baseline average-to-installed duty of 66.7%, gen0 lifts the duty to about 75% by funnelling modular work through general cranes, while gen1 and gen2 restore slack to roughly 72% and 69% respectively. The bottleneck counters reflect the same transition, with crane-limited time rising sharply only when modules share the general pool and receding as extra cranes are added, while quay-limited periods remain negligible.

Port-wide, all 100-10 variants add inland moves on top of stable ocean volumes. The default set-up delivers around +4.4% total TEU with an inland share near 22.7%; the more capacity-intensive gen2 reaches about +5.8% and 23.5% (Table **??**). Taken together, these results suggest that for early adoption the key design choice is how the crane interface for modular traffic is organised. Dedicated or incrementally expanded inland crane capacity converts the added calls into shorter barge queues and low module staging, while leaving seagoing performance broadly unchanged.

V. DISCUSSION

This section reflects on the findings of the simulation study and discusses their implications for operational efficiency and coordination across port-wide vessel–terminal operations. It places the results in context, outlines limitations, and identifies directions for future development.

A. Interpretation of Results

The composition of the vessel fleet and the configuration of terminal resources have a pronounced influence on port-wide performance. A balanced mix of conventional barges and modular or autonomous vessels delivers the highest overall efficiency: the intermediate 50:50 scenario yields the lowest average time in port and the most balanced quay utilisation, which suggests that hybrid operations are a pragmatic transition pathway when introducing modular concepts alongside existing barge services. A fully modular system further reduces waiting time but increases the number of calls, which can fragment berth use if scheduling coordination is insufficient.

When annual modular volume (TEU) is held constant, increasing module capacity substitutes fewer, longer calls for more, shorter ones. Module-mother visits therefore lengthen mainly through additional terminal dwell, while barge, feeder and seagoing times remain broadly unchanged. Fewer insertions reduce short-window fragmentation and lower inland anchorage. Effects on berth fit are terminal specific: where short calls previously created unusable gaps, fewer and larger calls improve packing; where geometry and long blocks dominate, longer visits can leave remainder gaps that reduce effective fit. In the capacity experiments, crane-driven congestion is negligible; module-crane productivity rises slightly with capacity as changeovers fall, but cranes are not the binding constraint.

Benefits can be realised quickly, provided the crane interface is organised for modular work. In the rapid-shift 100:10 scenario, dedicated or added inland crane capacity reduces inland waiting and preserves low berth loading; if modules must share the general quay crane pool, barge waits can rise.

Operationally, meaningful efficiency gains are achievable without major infrastructure expansion, provided that berth allocation and crane scheduling are adapted to handle smaller, more frequent arrivals, so that short modular

calls fill residual windows without displacing higher-priority seagoing work. Modular units distribute quay demand more evenly and sustain continuous crane activity by reducing idle intervals between vessel exchanges. Although average crane productivity remains relatively stable across scenarios, utilisation becomes more uniform under mixed and modular conditions, with fewer interruptions due to vessel waiting. This implies a shift from peak-intensive deployment to steady, balanced utilisation, supporting more predictable labour and energy use. Capacity experiments reinforce this reading: changing module capacity primarily rebalances call frequency and per-call dwell without creating persistent crane pressure. Quay access and berth-window design continue to govern system performance.

These implications are particularly relevant for Rotterdam's landlord governance model, where the port authority manages shared infrastructure and private terminal operators hold concessions. The outcomes highlight the importance of coordination mechanisms that transcend terminal boundaries so that aggregate berth demand stays synchronised with overall port capacity. Integrated barge-planning platforms (for example, Nextlogic) exemplify this need; synchronised scheduling reduces anchorage and enhances berth-fit efficiency without adding quay length. With constant annual modular TEU, larger per-call capacity reduces insertions and short-window fragmentation, lowering inland anchorage even as per-visit dwell increases, which clarifies why tools that smooth the timing of short calls are effective levers alongside physical capacity.

At Maasvlakte 2, where terminals employ high levels of automation, modular or autonomous vessels could be integrated with limited quayside adaptation. The principal challenge is operational synchronisation rather than infrastructure readiness. A modular arrival structure aligns with continuous handling cycles designed for predictable, high-frequency service. For the port authority, the transition should be accompanied by increased data sharing, common scheduling frameworks, and adaptive berth allocation policies to balance utilisation across the waterway system.

More broadly, modular and autonomous vessel concepts strengthen resilience and adaptability. Rotterdam's distributed planning approach is advantageous under uncertainty, allowing terminals to adjust locally while preserving system-wide coherence. The transition towards modularisation is not only technological but organisational, requiring digital integration and procedural flexibility. The model highlights both immediate operational potential and the institutional prerequisites—distributed coordination, adaptive scheduling, and shared situational awareness—needed to realise it. Deep-sea anchorage remains low and stable across scenarios (about 8–12% depending on mix), with average waiting below one hour by design.

B. Limitations

Confidentiality constraints limited access to detailed operational data, necessitating representative averages for calibration. While realistic in order of magnitude, these may not capture full operational variability. The model treats vessels as homogeneous classes with deterministic handling rates and fixed sequences, omitting stochastic variation in arrivals, crane performance, and coordination. Environmental influences such as weather, maintenance downtime, and disruptions are excluded to preserve computational tractability. Outcomes should therefore be read as indicative of relative differences rather than precise forecasts.

Experiments assume mixed-use quays and effectively uniform destination draws across terminals. In practice, traffic skews towards larger or preferred terminals and some ports operate dedicated barge quays, which may increase local window pressure and anchorage at high-volume sites. Extending the model with terminal-specific demand weights and dedicated-quay logic would quantify distributional effects and fairness across concessions.

Scope is limited to vessel–terminal interactions within Rotterdam. Broader network effects such as hinterland congestion, deep-sea berth conflicts, and schedule propagation are outside the model, as are empirical liner schedules; synthetic arrivals ensure volume consistency. The four-week horizon offers a realistic snapshot but omits seasonal dynamics. Sensitivities focus on berth capacity and crane allocation; other factors such as arrival variability, stochastic handling, berth-specific constraints, and heterogeneous scheduling policies are not yet included.

Validation is necessarily indirect, given the absence of large-scale modular deployments. Confidence stems from expert interviews, secondary datasets, and plausibility checks; the follow-up practitioner assessment supports qualitative trends but cannot remove uncertainty in absolute magnitudes. Finally, capacity results hold annual modular TEU constant; if uptake increases total modular work, the balance between call frequency, dwell, and berth fit may shift.

These limitations help position the results alongside prior work. The focus on length–time packing and the coupling of berth allocation with crane assignment follows the analytical view in [16], [18], extended here to inland rotations and distributed decision rights. The rotation-aware treatment of barges echoes delay propagation mechanisms noted in [4], [9]. Relative to sea-going focussed simulation surveys [5]–[7], the scope shifts to port-internal coordination where thresholds and single server effects are decisive. The patterns are consistent with landlord-port practice [1] and platform-based coordination [3], [8]. The mixed use quay assumption and the abstracted module interface may overstate pooling benefits compared with studies that assume dedicated quays or centralised control [10], which motivates the scenario extensions outlined above.

C. Implications and Outlook

Discrete-event simulation proves effective for assessing coordination strategies and technological transitions in port operations. Managerially, the model supports evaluating trade-offs between efficiency, berth congestion, and resource utilisation under different fleet compositions. Partial substitution with modular units increases resilience, especially during peaks or irregular arrivals. With fixed modular volume, larger units trade fewer insertions for longer per-call dwell; cranes remain secondary so long as single-server effects are avoided.

Read against the upsizing of seagoing ships, modular inland units complement system performance: by filling short berth opportunities with small, predictable exchanges they protect crane intensity on mainline berths and shorten the completion tail. Practically, modules should be planned as elastic capacity around large-ship schedules, preferably on side or barge-dedicated quays sized so that one or two modules complete within a typical short window, while primary seagoing cranes focus on the main exchange.

Design choices at the vessel–quay interface unlock these gains: select module capacities aligned with common spreader modes for 20- and 40-foot moves; keep convoy length compatible with bollard spacing and fender geometry for predictable mooring; ensure fine-control propulsion and station-keeping to reduce variance between lines-on and first lift. On the quay, assure one crane per short call with surge to two; minimise penalties when alternating between 20- and 40-foot moves; size near-quay buffers to a module batch; position duplicated charging pedestals outside crane envelopes; and, where permitted, enable safe parallel working with clear zoning and mooring rules.

D. Recommendations for operations and policy

Prioritise hybrid adoption. Begin with a modest modular share, for example 100:10, and move toward 50:50 where performance and traffic-management capacity justify it. Treat module capacity as a scheduling lever, pairing capacity choice with berth-window design and look-ahead settings rather than adding cranes indiscriminately. Select module size and the minimum number of module cranes per terminal as a pair, choosing the smallest count that removes single-server effects at the busiest terminals; two cranes are a practical default in module-heavy mixes.

Strengthen distributed coordination. Publish berth-window catalogues with explicit minima tuned to emerging module sizes, extend reservation and time-stamped commitment interfaces on shared platforms, and maintain real-time situational awareness for higher movement counts. Update safety frameworks where needed to allow parallel working with clear zoning and mooring rules.

Stage implementation. In the near term, improve data quality and window exposure, establish pilot corridors for modular calls, and monitor waiting, berth fit, and reliability. In the medium term, align window minima with observed module sizes, calibrate scheduling look-ahead, and extend parallel-working rules at selected barge quays.

E. Recommendations for the design of vessels and quays

Use small autonomous modules as elastic capacity around large-ship exchanges, filling short berth opportunities and protecting crane intensity on mainline berths. Where rules permit, enable parallel working of two modules with independent moorings and one crane per module; if abreast mooring is not permitted, provide adjacent short pockets with separate access.

Choose capacities that complete one or two modules within typical short windows and align with common spreader modes for 20- and 40-foot moves. Ensure manoeuvrability and approach control to reduce variance between lines-on and first lift. On the quay, assure one crane per short call with surge to two, minimise tool-change penalties, size near-quay buffers to a module batch, position duplicated charging pedestals outside crane envelopes, and expose window minima digitally so autonomy stacks can plan to earliest start.

F. Future research

Extend the model to incorporate uncertainty in arrivals, handling rates, and weather-related disruptions to assess robustness. Relax the fixed-volume assumption in capacity tests to study cases where modular volume scales with size, enabling joint design of module size, window policies, and crane counts. Embed optimisation-based window assignment and multi-agent coordination to bound performance and simulate adaptive decision-making. Explore dedicated barge-quay configurations and terminal-specific arrival weights to reveal distributional effects on berth fit, anchorage, and service fairness. Combine a grid over module size and crane counts with compact design-of-experiments or response surfaces to derive per-terminal sizing curves. Pursue empirical validation using data from pilots to calibrate service times, manoeuvring, and coordination dynamics.

G. Contributions to research, practice, and MAGPIE

The study links operational simulation with governance-oriented analysis of port automation, showing how discrete-event simulation can capture coordination dynamics under distributed planning for intra-port logistics. For practitioners, it provides an analytical basis for evaluating modular adoption and underlines that efficiency gains depend as much on information sharing and institutional design as on technology. Within MAGPIE (Work Package 5, Demo 6 on autonomous e-barge and transhipment), the model offers a system-level assessment of operational effects under varying degrees of modular and autonomous deployment, complementing pilot demonstrations and supporting future analyses of large-scale deployment and digital coordination across the multimodal network.

VI. Conclusions

This paper examined whether short, parallelisable module calls improve waiting, berth fit and throughput at multi-terminal scale, and identified the configurations that enable these gains in landlord ports. The simulation shows

that hybrid fleets, especially 50:50 and 25:75, shorten inland turnaround and raise inland and total throughput while keeping seagoing service stable. Early adoption is feasible: a 100:10 mix increases total throughput by about 3 to 6 % and lifts the inland share by roughly 2 to 4 % without harming deep-sea or feeder times.

Governing resources are berth access and general quay cranes. In high module shares, at least two module cranes per terminal are required to avoid single-server effects; further cranes mainly add headroom. With annual module volume held constant, capacities around 24 to 36 TEU minimise staging and waiting. Module-handling speed is secondary at moderate modular shares but becomes critical when the dedicated pool approaches saturation. These results indicate that modularisation can unlock latent quay capacity through better temporal fit, provided that admission and crane interfaces are dimensioned and scheduled accordingly.

REFERENCES

[1] A. Pallis and J.-P. Rodrigue, "Chapter 6.2 – port efficiency," in *Port Economics, Management and Policy*, 2024. [Online]. Available: https://porteconomicsmanagement.org/pemp/contents/part6/port-efficiency/

[2] Port of Rotterdam, "Barge performance monitor," 2019. [Online]. Available: https://www.portofrotterdam.com/en/logistics/connections/intermodal-transportation/inland-shipping/barge-performance

[3] Portbase, "Melding container achterland," 2025. [Online]. Available: https://www.portbase.com/services/melding-container-achterland/

[4] P. Shobayo and E. Van Hassel, "Container barge congestion and handling in large seaports: a theoretical agent-based modeling approach," *Journal of Shipping and Trade*, vol. 4, no. 1, 2019. [Online]. Available: https://doi.org/10.1186/s41072-019-0044-7

[5] D. Steenken, S. Voß, and R. Stahlbock, "Container terminal operation and operations research—a classification and literature review," *OR Spectrum*, vol. 26, no. 1, pp. 3–49, 2004. [Online]. Available: https://doi.org/10.1007/s00291-003-0157-z

[6] C. Bierwirth and F. Meisel, "A survey of berth allocation and quay crane scheduling problems in container terminals," *European Journal of Operational Research*, vol. 202, no. 3, pp. 615–627, 2010. [Online]. Available: https://doi.org/10.1016/j.ejor.2009.05.031

[7] B. Dragovic, E. Tzannatos, and N. K. Park, "Simulation modelling in ports and container terminals: Literature overview and analysis by research field, application area and tool," *Flexible Services and Manufacturing Journal*, vol. 29, 2017.

[8] Nextlogic, "Hoe het werkt," 2024. [Online]. Available: https://www.nextlogic.nl/zowerkthet/

[9] D. Pachakis and A. Kiremidjian, "Ship traffic modeling methodology for ports," *Journal of Waterway, Port, Coastal, and Ocean Engineering*, vol. 129, no. 5, pp. 193–202, 2003.

[10] M. N. Abourraja, A. El Hilali Alaoui, L. Benabbou, and A. El Mhamedi, "Simulation approaches in seaport container terminals: A literature review," *Simulation Modelling Practice and Theory*, vol. 116, p. 102482, 2022.

[11] International Maritime Organization, "Autonomous shipping," 2018. [Online]. Available: https://www.imo.org/en/mediacentre/hottopics/pages/autonomous-shipping.aspx

[12] ——, "Autonomous ships: regulatory scoping exercise completed," 5 2021. [Online]. Available: https://www.imo.org/en/mediacentre/pressbriefings/pages/massrse2021.aspx

[13] Lloyd's Register, "Cyber-enabled ships shipright procedure – autonomous ships," Tech. Rep., 2016.

[14] M. Schiaretti, L. Chen, and R. Negenborn, "Survey on autonomous surface vessels: Part i - a new detailed definition of autonomy levels," *Lecture Notes in Computer Science*, vol. 10531, pp. 219–233, 2017. [Online]. Available: https://doi.org/10.1007/978-3-319-68496-3_15

[15] Ø. J. Rødseth, L. A. L. Wennersberg, and H. Nordahl, "Improving safety of interactions between conventional and autonomous ships," *Ocean Engineering*, vol. 284, p. 115208, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0029801823015901

[16] Q. Zeng, Z. Yang, and X. Hu, "Disruption recovery model for berth and quay crane scheduling in container terminals," *Engineering Optimization*, vol. 43, no. 9, pp. 967–983, 2 2011. [Online]. Available: https://doi.org/10.1080/0305215x.2010.528411

[17] A. Imai, E. Nishimura, and S. Papadimitriou, "Berth allocation problem," *Transportation Research Part E: Logistics and Transportation Review*, vol. 37, no. 5, pp. 329–348, 2001. [Online]. Available: https://doi.org/10.1016/S1366-5545(00)00023-2

[18] J. Dai, W. Lin, R. Moorthy, and C.-P. Teo, *Berth Allocation Planning Optimization in Container Terminals*. Springer US, 2008, pp. 69–104. [Online]. Available: https://doi.org/10.1007/978-0-387-75240-2_4

B

Port Call Process

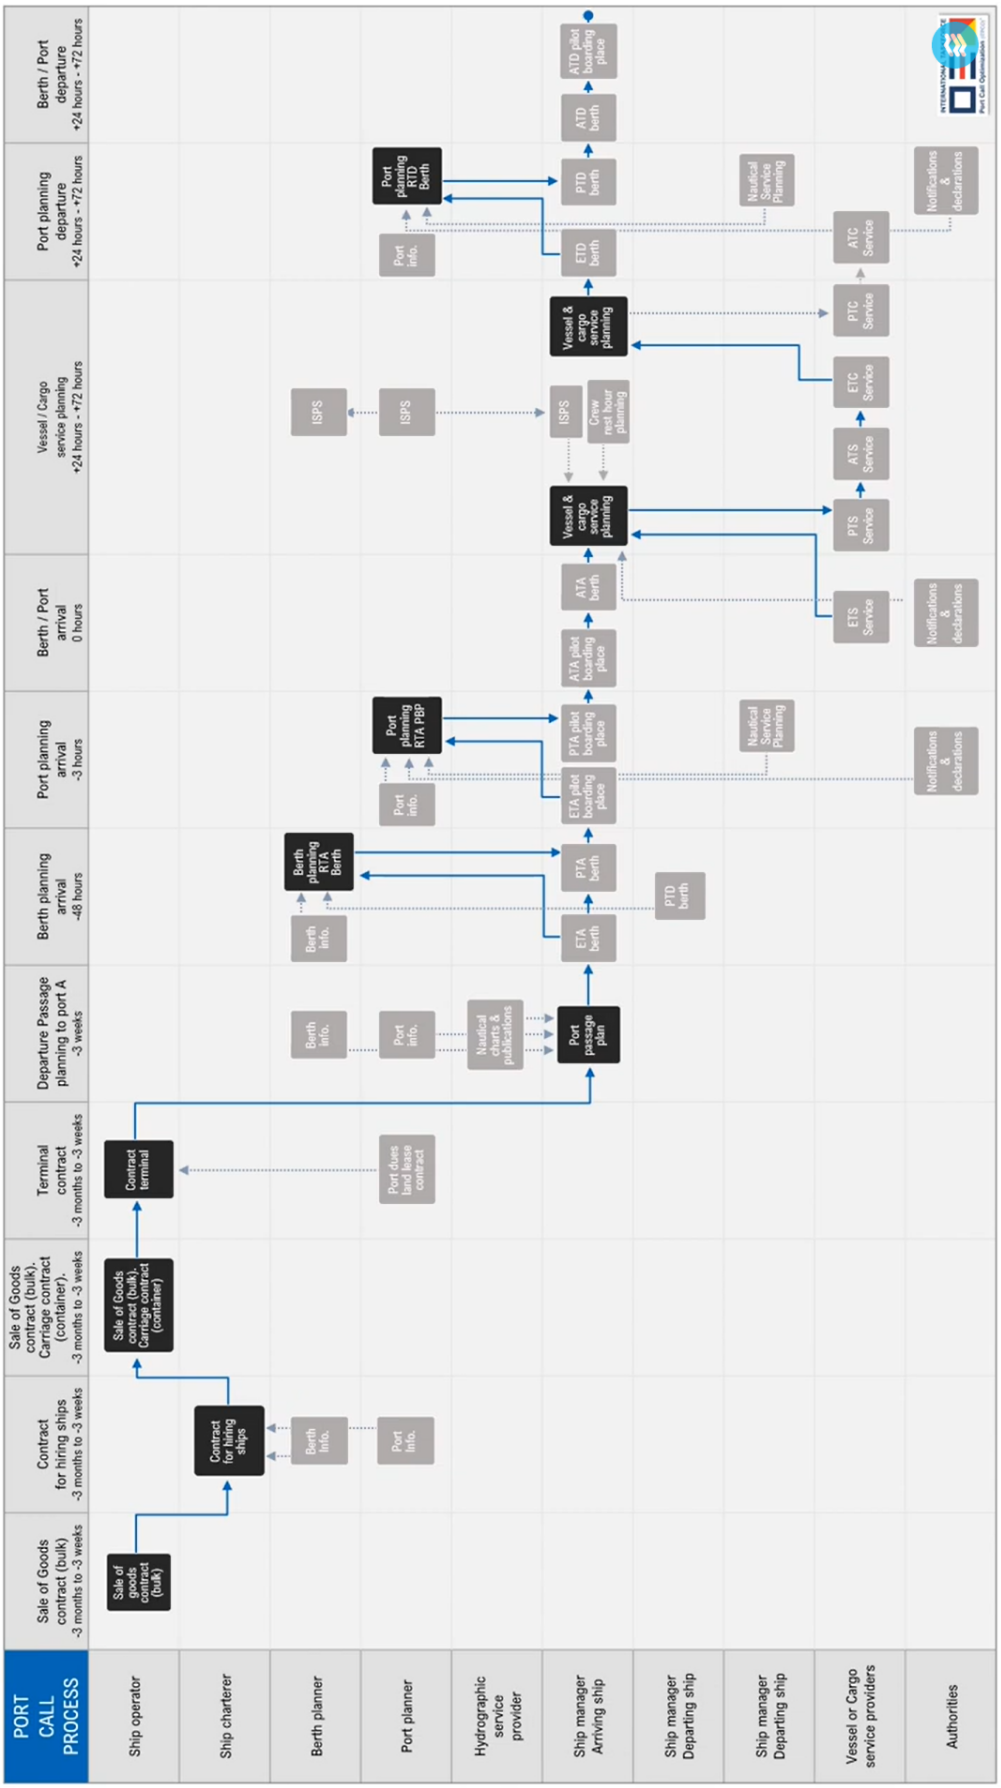

Figure B.1: Port call process (Port of Rotterdam, 2023)

C

Additional Concept Details and Special Cases

This appendix provides background narratives, diagrams, and out-of-scope variants. It supports the high-level overview in Chapter 3. The simulation experiments implement only the baseline and modular splitting concepts; all binding rules for implemented scenarios are specified in Chapter 5. Roundtrip service and Modular interchange are not modelled. Figures C.1 and C.2 are included to visualise the process logic and decision points discussed below.

Roundtrip Service

A fixed loop visits a predefined set of terminals on a repeating circuit. Predictable rotations can help synchronise with time windows and reduce ad hoc repositioning. Exposure remains to missed windows and the risk that small calls fall below minimum-move expectations.

Operational picture: The vessel completes a sequence such as A $\rightarrow$ B $\rightarrow$ C, then returns to A and repeats. Moves for each terminal are collected into the next pass of the loop. The concept suits steady flows and stable windows. As shown in Figure C.1, the fixed route and repeat loop are explicit, and there is no in-port resequencing within a rotation.

When it works well: Stable demand per terminal, reliable window regimes, and plans that protect the loop from cascading delays.

Risks: Delay at one stop can cascade to later stops in the rotation. Thin flows may create sub-threshold calls that either wait or underuse berth time.

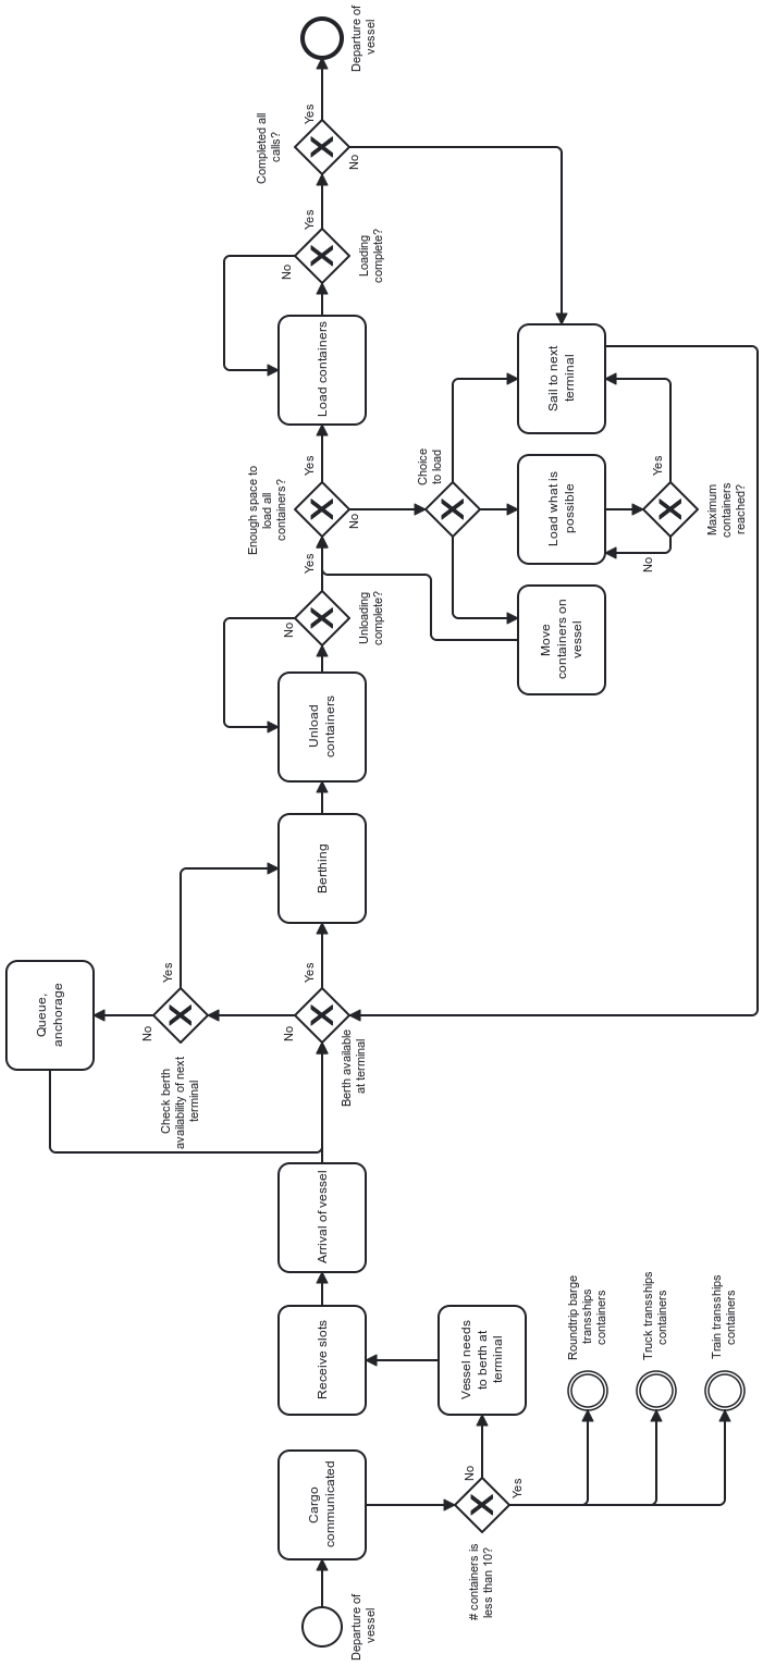

Figure C.1: Roundtrip service: simplified process logic (BPMN) showing the fixed loop and absence of in-port resequencing.

Illustrative Modelling of Truck-based ITT

The current simulation experiments do not model truck-based ITT. This subsection documents how such behaviour could be represented in a future extension, based on practice in Rotterdam and the literature.

Not all ITT flows in Rotterdam are executed by barge. When volumes between two terminals are small, it is often more efficient to use trucks rather than schedule an additional barge call. This practice is confirmed in literature and by practitioners; Van der Horst and De Langen (2008) note that truck-based ITT remains the dominant mode for small container flows because of its flexibility, while Konings (2010) and industry documentation indicate that barge calls are typically scheduled once a minimum volume is reached. In practice this threshold is context-dependent and negotiated bilaterally, and expert interviews indicate that it is commonly applied in Rotterdam.

How it could be modelled

A minimal and defensible representation is a *minimum call-size threshold* θ [TEU] that classifies small ITT flows to truck:

- Classification rule: For each prospective inter-terminal move in a voyage plan, if the expected load is $< \theta$ then handle by truck; if $\geq \theta$ then generate a barge (or module) call.
- DES integration: Truck-handled flows do not create vessel or module events, berth allocations, crane use, or anchorage waiting. They are counted only in throughput totals to preserve mass balance.
- Sensitivity: Vary θ to reflect different consolidation practices or policy regimes (e.g., $\theta \in \{0, 10, 50, 100\}$ TEU).

Implications for indicators

With the threshold representation above, truck ITT is exogenous to port-resource dynamics. It is excluded from berth, crane, and waiting statistics by construction. If an emissions analysis is performed that includes road transport, a separate accounting pathway would be needed; in studies that isolate the operational effects of barge coordination, truck ITT can be omitted from emissions per TEU.

Notes for future extensions

More detailed variants could assign a nominal truck travel time, gate capacity, or a simple queue for truck ITT between specific terminal pairs. These additions are not required for barge-focused coordination studies but may be useful for whole-system road–water trade-off analyses.

Case simplification (illustrative only): Inter-terminal flows below a minimum call-size θ [TEU] are handled by truck (or rail) rather than generating a barge or module call. This rule is not used in the present experiments; it is recorded here to support future work.

Modular Interchange

Modular interchange extends modular operations by allowing modules from different mothers or convoys to be exchanged mid-voyage. The intention is to align modules with terminals that are ready to handle them, without tying each module to its original mother throughout the port stay.

Operational picture: Two or more mothers arrive, split into modules, then exchange modules according to terminal readiness. Exchanged modules sail directly to their next terminal and later recombine with either their original mother or a designated partner at a rendezvous point. Figure C.2 highlights the decision gateway where modules may be exchanged across convoys, and the downstream recombination point.

When it works well: High synchronisation maturity, reliable short windows, and shared situational awareness across operators and terminals.

Risks: Additional coordination burden and handover complexity, with stringent information requirements on timing, permissions, and module status.

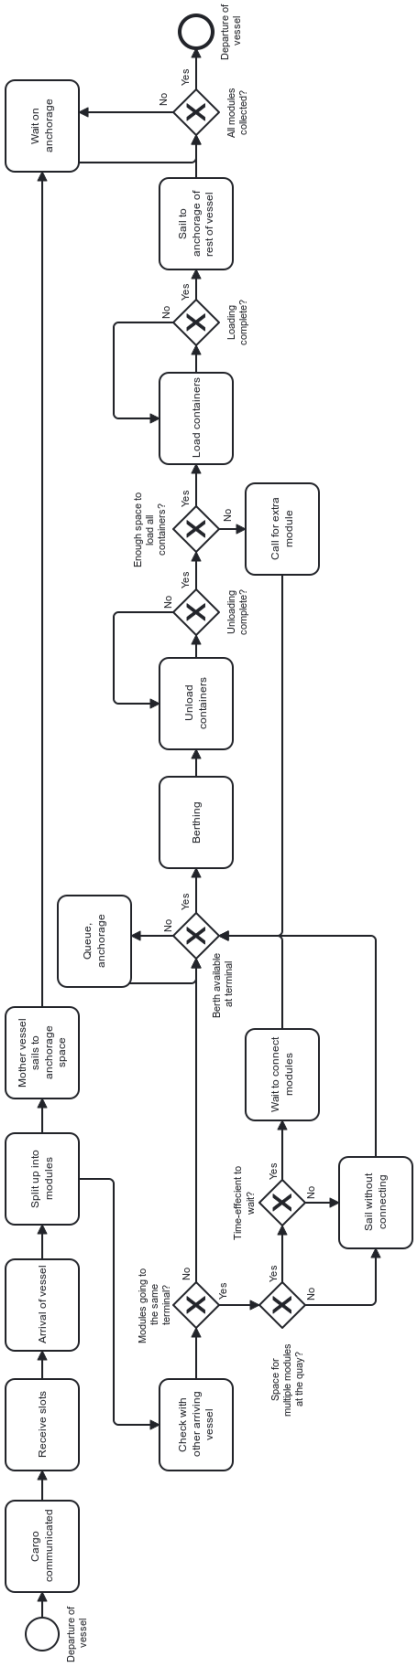

Figure C.2: Modular interchange: simplified process logic (BPMN) showing the exchange gateway and subsequent recombination.

Special Cases for Modular Operations

The cases below are adaptations of modular splitting and capture operational edge conditions that may be useful for future experiments. They are described to clarify trade-offs and to provide context for the split–rendezvous–recombination logic shown in Figure C.2. They are not implemented unless explicitly referenced in Chapter 5.

Excess Load Beyond Planned Module Capacity

The planned load for a terminal exceeds the capacity of its assigned module. The following policies can be instated:

1. Call an extra module: Dispatch an additional module dedicated to the overflow. Improves compliance with minimum moves, increases berth and crane demand.
2. Backfill on a module with spare capacity: Allow a second module to call sequentially at the same terminal to pick up the remainder. Reduces extra split cost, lengthens in-port sailing for that module.
3. Staggered pick-up: Leave the overflow for the next voyage when service promises allow. Eliminates immediate pressure, increases dwell.

Very Small Calls Below Minimum Move

A terminal's work is too small to justify a dedicated module call under minimum-move expectations. The following policies can be instated:

1. Two-stop module: Allow one module to visit two nearby terminals sequentially, combining small calls into one tour.
2. Consolidate upstream: Reassign containers to a different terminal when permitted, so the module meets the minimum at a single stop.
3. Truck pick-up: Use truck ITT for the remainder, consistent with the θ rule described illustratively in the section about truck-based inter-terminal transport.

Load-only Calls at a Terminal

A terminal has no unloading for the voyage, only outbound containers to load. The following policies can be instated:

1. Earliest-free module with capacity: After finishing its primary terminal, the earliest-free module that has spare capacity sails to perform the load-only call.
2. Pre-positioned empty module: Allocate an empty module from the start for the load-only call, to protect schedules at other terminals.
3. Defer to next voyage: If windows or capacity are tight, move the load-only call to the next rotation.

Departure Rules

Not all modules return to the rendezvous in time, or the plan foresees a different number of modules at departure. The following policies can be instated:

1. Strict recombination: Mother departs only when all assigned modules arrive. Maximises consolidation, increases risk of departure delay.
2. Flexible recombination: Mother may depart with a subset when a time limit is reached; remaining modules rejoin a later mother or exit independently. Protects punctuality, adds coordination overhead.

Chapter 3 points to this appendix for fuller narratives and edge-case context. All operative rules used in the experiments are specified in Chapter 5.

D

Model Settings and Inputs

This appendix is the single source of numerical values and configuration used by the simulation. It is organised to make experiments reproducible and to keep Chapter 5 free of repeated numbers. The simulation chapter explains the methodology, policies, and event flow; the present appendix records the concrete settings they use. These are divided in the following parts: Run control (horizon, warm-up, seeds), port layout and resources, vessel classes and handling parameters, arrival-process parameters for the internal generator, berth-admission, planning rules as implemented, and experiment scenarios.

D.1. Global Run Control

tocsectionD.1 Global Run Control

Table D.1: Run control and reporting.

Setting	Value / rule
Total simulated days	32 (2 warm-up, 28 analysis, 2 cool-down)
Event scan / policy tick	5 minutes
Precision target	Replicate until 95% CI half-width of primary KPIs $\leq$ 5% of the mean
Random seeds	Master seed per scenario; independent streams for arrivals (per class), call sizes, and handling-rate noise; substream = replication index
Reporting windows	24 h, 7 d, and analysis window (28 d)

D.2. Port Layout, Distances, and Sailing Times

tocsectionD.2 Port Layout, Distances, and Sailing Times

The model instantiates a stylised Maasvlakte with five container terminals, an anchorage, and entrance/departure nodes. The coordinates below are used for plotting the layout (1 map unit $\approx$ 0.2 km); operational distances are taken directly from the fixed matrix in Table D.3 rather than computed from coordinates. Sailing times equal distance divided by class speed (Table D.4).

Table D.2: Terminal geometry and installed resources (baseline).

Name	X [unit]	Y [unit]	Quay length [m]	Quay cranes [#]
Terminal A (APMT)	30	10	2000	13
Terminal B (RWG)	0	0	1700	16
Terminal C (ECT Euromax)	150	150	1500	16
Terminal D (ECT Delta N)	200	50	1800	17
Terminal E (ECT Delta S)	205	0	1800	17
Anchorage	320	150	–	–
Entrance	300	200	–	–
Departure	310	190	–	–

Table D.3: Intra-port distance matrix used by the simulation (km).

	Entrance	Terminal A	Terminal B	Terminal C	Terminal D	Terminal E	Anchorage	Departure
Entrance	0.0	10.0	10.0	5.5	3.0	4.0	0.5	0.0
Terminal A	10.0	0.0	0.5	5.0	8.0	11.0	9.5	10.0
Terminal B	10.0	0.5	0.0	5.0	8.0	11.0	9.5	10.0
Terminal C	5.5	5.0	5.0	0.0	3.5	6.0	4.5	5.5
Terminal D	3.0	8.0	8.0	3.5	0.0	4.0	2.5	3.0
Terminal E	4.0	11.0	11.0	6.0	4.0	0.0	3.5	4.0
Anchorage	0.5	9.5	9.5	4.5	2.5	3.5	0.0	0.5
Departure	0.0	10.0	10.0	5.5	3.0	4.0	0.5	0.0

Note: Entrance and Departure are colocated in the current set-up; both appear because they are distinct nodes in the event logic. If not needed in figures, the Departure row/column can be omitted without changing the simulation. Crane pools: Each terminal has a general quay-crane pool (counts above). A module-crane pool serves module calls and is varied by scenario (Section D.6). Berths: A vessel may occupy a quay segment only if effective length (hull + fore/aft buffers) fits; spacing by crane is enforced along the quay. There are no terminal opening hours in the model.

D.3. Vessel Classes and Handling

tocsectionD.3 Vessel Classes and Handling

Table D.4: Vessel classes and handling parameters.

Class	Length [m]	Speed [kn]	Max cranes [#]	Handling time [min/move/crane]	Fore–aft buffer [m]
Deep-sea	366	6–10	4	2.0	20
Feeder	200	5–8	2–3	2.5	10
Barge	110	4–7	1	3.0	8
Module	20	3–5	1 (module crane)	1.5	6

Productivity Variability - Per-crane effective productivity is a bounded multiplicative deviation of the nominal rate:

$$p_c \;=\; \bar{p}_{\text{class}}\,(1+\varepsilon_c), \qquad \varepsilon_c \sim \mathcal{N}(0,\sigma_p^2) \text{ truncated to bounds.}$$

Table D.5: Variability parameters by equipment.

Equipment	Mean [moves/h]	σ_p	Truncation [factor]	Notes
Quay crane (sea)	30	0.12	[0.6, 1.4]	Deep-sea & feeder work
Quay crane (barge)	20	0.12	[0.6, 1.4]	Lower mean via handling time
Module crane	40	0.10	[0.7, 1.3]	Dedicated to modules

Call Sizes - Call sizes (load and discharge TEU) are drawn from class distributions used by the generator; fitted parameters are listed here for reproducibility.

- Deep-sea: Normal $(\mu{=}4{,}476,\ \sigma{=}224)$, min 500, max $20{,}000$.
- Feeder: Normal $(\mu{=}1{,}000,\ \sigma{=}250)$, min 100, max $4{,}000$.
- Barge: Normal $(\mu{=}240,\ \sigma{=}60)$, min 40, max 600.
- Module mother: derived from module capacity; baseline ≈ 63 TEU per mother mission (Section D.9).

D.4. Arrival Processes (Generator Inputs)

tocsectionD.4 Arrival Processes (Generator Inputs)

All classes are generated stochastically (no liner timetables or appointments). For each class:

- an arrival process is specified (NHPP with daily profile or homogeneous Poisson with hour-of-day placement);
- per-call work is sampled from the class distribution in Section D.3;
- terminal assignment follows class–terminal shares (kept fixed across scenarios unless stated).

D.5. Planning and Admission Rules (as Implemented)

tocsectionD.5 Planning and Admission Rules (as Implemented)

Table D.6: Berth-admission and planning parameters.

Rule	Implementation
Priority order (baseline)	Deep-sea $>$ Feeder $>$ Barge $=$ Module; FCFS within class
Admission gates	(i) effective-quay space fits; (ii) class-specific crane rule satisfied. No terminal time windows are modelled
Deep-sea/feeder attachment	May berth without immediate crane attachment; cranes can attach with delay
Look-ahead buffer Δ	If admitting a lower-priority call would pre-empt a higher-priority start within Δ, keep berth empty (scenario value in Table D.7)
Crane assignment	Allocate up to class cap and available pool; allow mid-service rebalance subject to quay spacing
Anchorage	Infinite capacity; queue re-evaluated on each ServiceCompletion and periodic check

Table D.7: Priority vectors and look-ahead used across scenarios.

Scenario	Priority vector	Δ (min)
Baseline	[Deep-sea, Feeder, Barge=Module]	20
Barge-friendly	[Deep-sea, Barge=Module, Feeder]	20
Strict sea-first	[Deep-sea, Feeder, Barge, Module]	20

D.6. Barge−module Mix Ratios and Counts

tocsectionD.6 Barge–module Mix Ratios and Counts

Deep-sea and feeder call totals are held fixed at $N_{\mathrm{DS}}^{\mathrm{ref}} = 2{,}073$ and $N_{\mathrm{FE}}^{\mathrm{ref}} = 1{,}088$ (see Section D.7). Inland demand is varied by the barge–module mix while keeping total inland TEU approximately constant across mixes. Per–call averages are $t_B = 240$ TEU per barge call and $t_M = 63$ TEU per module mother mission.

Table D.8: Barge–module mixes with fixed sea-call totals and nearly constant inland TEU.

Mix (B%–M%)	Deep-sea	Feeder	Barge calls	Module missions	Barge TEU	Module TEU	Total inland TEU
100–0	2,073	1,088	13,000	0	3,120,000	0	3,120,000
75–25	2,073	1,088	9,750	12,381	2,340,000	780,003	3,120,003
50–50	2,073	1,088	6,500	24,762	1,560,000	1,560,006	3,120,006
25–75	2,073	1,088	3,250	37,143	780,000	2,340,009	3,120,009
0–100	2,073	1,088	0	49,524	0	3,120,012	3,120,012

Notes. Small deviations (up to 12 TEU) from the 3,120,000 TEU target arise from rounding call counts to integers. Deep-sea and feeder calls are listed for completeness and remain unchanged across mixes.

D.7. Sea-freight Load Variants

tocsectionD.7 Sea-freight Load Variants

Load stress is applied by changing the *number of deep-sea and feeder calls* generated by the model while keeping their call-size distributions and all other parameters fixed. For a given scenario with scale

factor $s \in \{0.8, 1.2\}$, the generator targets

$$N'_{\mathrm{DS}} = s\, N^{\mathrm{ref}}_{\mathrm{DS}}, \qquad N'_{\mathrm{FE}} = s\, N^{\mathrm{ref}}_{\mathrm{FE}},$$

and then samples exactly N'_{DS} and N'_{FE} arrival timestamps using the class-specific daily profile (Section D.4). Barge and module processes are unchanged.

Table D.9: Target numbers of sea calls used in the load-variation scenarios.

Vessel type	80% load	120% load
Deep-sea	1 658	2 488
Feeder	870	1 306

The baseline (reference) totals $N^{\mathrm{ref}}_{\mathrm{DS}}$ and $N^{\mathrm{ref}}_{\mathrm{FE}}$ are reported with the traffic summary in Section D.6.

D.8. Crane-allocation Scenarios

tocsectionD.8 Crane-allocation Scenarios
Additive (quay cranes held constant): Global module-crane pool size $= 1, 2, 3$ (baseline $= 2$). Substitution (net-neutral): Replace one quay crane per terminal with one module crane, and separately two with two.

	Run 1		Run 2		Run 3		Run 4		Run 5	
Terminals	Q	Md	Q	Md	Q	Md	Q	Md	Q	Md
Terminal A	13	2	13	1	13	3	11	2	12	1
Terminal B	16	2	16	1	16	3	14	2	15	1
Terminal C	16	2	16	1	16	3	14	2	15	1
Terminal D	17	2	17	1	17	3	15	2	16	1
Terminal E	17	2	17	1	17	3	15	2	16	1

Figure D.1: Scenarios for crane availability and substitution

D.9. Module Capacity Parameters

tocsectionD.9 Module Capacity Parameters
Modules are composed of 40 ft containers, counted as 2 TEU each. Capacity levels are obtained by varying the number of 40 ft bays in the lengthwise direction while keeping width and height fixed at two across and three high, i.e. layouts $1{\times}2{\times}3$, $2{\times}2{\times}3$, and $3{\times}2{\times}3$ containers for 12, 24, and 36 TEU respectively. Only the *effective length* influences berth fit in the model; beam and draft are abstracted.

To account for clearances and onboard systems (hull, propulsion, batteries, etc.), a fixed allowance of $L_{\mathrm{fix}} = 8$ m is added. The per-bay pitch is taken as $L_{40} = 6$ m, which rounds the nominal 40 ft length (5.896 m) to include spacing. The effective module length is therefore

$$L_{\mathrm{module}}(b) \;=\; b\, L_{40} + L_{\mathrm{fix}}, \qquad L_{40} = 6\,\mathrm{m}, \quad L_{\mathrm{fix}} = 8\,\mathrm{m}, \quad b \in \{1, 2, 3\}.$$

Table D.10: Module layouts, capacities, and effective lengths used in the simulation.

Capacity (TEU)	Layout (L×W×H)	40 ft bays b	TEU	Effective length L [m]
12	$1{\times}2{\times}3$	1	12	$1 \cdot 6 + 8 = 14$
24	$2{\times}2{\times}3$	2	24	$2 \cdot 6 + 8 = 20$
36	$3{\times}2{\times}3$	3	36	$3 \cdot 6 + 8 = 26$

Note. The rounding to $L_{40} = 6$ m and the fixed allowance $L_{\mathrm{fix}} = 8$ m are modelling choices that fold spacing and equipment into a single length parameter for berth-fit calculations.

To compare capacity levels on an equal-demand basis, the annual number of module missions must be scaled so that total annual module TEU is constant across scenarios. Let M_{ref} be the number of module missions in a reference case with per-module capacity C_{ref} (TEU), and let C_{new} be the capacity in a new case. The missions for the new case follow

$$M_{\text{new}} \;=\; M_{\text{ref}} \times \frac{C_{\text{ref}}}{C_{\text{new}}}.$$

In the model, `module_mission` counts mothers, while children are spawned in the simulation, therefore the scaling is applied to `ANNUAL_ARRIVALS["module_mission"]`.

For the reference case with a barge-module ratio 25-75, the capacity of one module $C_{\text{ref}} = 24$ TEU with $M_{\text{ref}} = 37{,}143$ module missions:

$$\begin{aligned}
\text{12 TEU:} \quad M_{12} &= 37{,}143 \times \frac{24}{12} = 74{,}286,\\
\text{24 TEU:} \quad M_{24} &= 37{,}143 \times \frac{24}{24} = 37{,}143,\\
\text{36 TEU:} \quad M_{36} &= 37{,}143 \times \frac{24}{36} = 24{,}762.
\end{aligned}$$

Similar calculations can be done for the module-barge mix 50-50. This results in the amount of modules displayed in Table D.11

Table D.11: Annual module missions for equal total demand, normalised to the 24 TEU reference

Capacity (TEU)	mix 25-75	mix 50-50
12	74,286	49,524
24	37,143	24,762
36	24,762	16,508

E

Validation Checks

This appendix supplements Chapter 6 by providing additional material on the validation of the simulation model. It consolidates insights from expert interviews and documents comparisons with publicly available indicators.

E.1. Expert Interviews

The interviews were selected to capture complementary perspectives that directly shape the model set-up: (i) the organisational/planning viewpoint of the Port of Rotterdam, and (ii) the vessel-operations viewpoint of a barge operator.

Interview: Business Manager Barge & Inland Terminals

Interviewee: Business Manager Barge & Inland Terminals, Port of Rotterdam
Purpose: Validate modelling assumptions and collect operational guidance for barge and modular operations.

Scope and topics

The interview covered: (i) vessel classes and priority rules, (ii) terminals and quay representation, (iii) indicative bands for call sizes and crane use, (iv) metric boundaries for time in port, and (v) the position of a modular mother-with-modules concept in current planning practice.

Consolidated findings and modelling implications

Planning priorities and service discipline - Deep-sea receives precedence, feeder services are scheduled by terminals, and barges are coordinated via an integrated platform. A pragmatic hierarchy is *deep-sea* $\rightarrow$ *feeder* $\rightarrow$ *barge*; a modular concept tied to deep-sea-critical flows could be positioned together with barges. This hierarchy is consistent with an FCFS-within-class scheme and a class-priority rule.

Terminals and quay realism - For a Rotterdam-like case, representing five deep-sea terminals on the Maasvlakte improves plausibility. Barge quays can differ from deep-sea quays within the same terminal; distinguishing these improves berth-access realism and berth-fit logic. This is an option for future work as there is a mix between mixed and seperated quays.

Call sizes and crane allocation (indicative bands) - Typical patterns include one quay crane for barges and multiple parallel cranes for deep-sea. Indicative productivity: barges $\sim$20 moves/h per crane; deep-sea $\sim$30–40 moves/h per crane. Barge call-size thresholds used operationally include a minimum around 15 moves and fixed-window thresholds in the $\sim$175–300 range. These bands are suitable for scenario design and sensitivity analysis.

Staging of modules - When the mother vessel is off-berth, a practical staging location for modules is along Maasvlakte barge quays; modelling a dedicated waiting node supports analysis of recombination effects.

Alignment with the model - The current model implements class priority (deep-sea over barge), FCFS within class, and deterministic single-case checks for verification. These elements align with the advised service discipline.

Post-interview Correspondence: Port of Rotterdam (Email Summary)

Type: Follow-up email after the interview
Purpose: Validate assumptions and provide aggregate reference figures used for plausibility checks and scenario bands.

Summary of Figures from Correspondence:

- Deepsea: average call size $\approx 4{,}500$ TEU, average time at berth ≈ 40 h.
- Feeder (shortsea): average call size $\approx 1{,}000$ TEU, average time at berth ≈ 12 h.
- Deepsea vessels $> 18{,}000$ TEU: average call size $\approx 11{,}500$ TEU, average time at berth ≈ 70 h.

Conversion: 1 container $= 1.7$TEU.
Port 2024 container facts and figures: total container vessel calls $\approx 8{,}701$ of which deepsea $\approx 2{,}073$ with average call size $4{,}476$ TEU, feeder $\approx 1{,}088$ with average call size $1{,}000$ TEU, shortsea $\approx 5{,}540$ (primarily at RST, excluded from scope).
Barge: average call size per rotation ≈ 240 TEU, that is ≈ 140 moves.

Interview Expert HTS

Purpose: Elicit vessel-side operational practices affecting call formation, sequencing, and realised service.

The discussion focused on: practical call formation and split across terminals, arrival and waiting patterns at Maasvlakte, typical crane assignment and realised productivity ranges at barge quays, operational responses to protected windows and deep-sea movements, and implications for time-in-port measurement from a vessel perspective.

Operators tend to consolidate moves at a subset of terminals given current appointment and availability patterns, which amplifies day-to-day variability in barge call sizes. In addition, many terminals apply a minimum call-size policy for barge appointments; sub-threshold rotations are typically asked to combine moves or are deprioritised in scheduling. Observed waiting is concentrated around peaks induced by deep-sea operations and protected windows; FCFS within class with “fit-before-priority” captures the dominant mechanism, provided that quayside opening-hour constraints are respected. One-crane service is the norm at barge quays, with realised rates in the high teens to low twenties moves/h, depending on set-up losses and short interruptions. Using a band of 18–22 moves/h is representative when no downtime is explicitly modelled.

E.2. Comparison with Publicly Available Data

The model is compared with publicly available indicators from the Port of Rotterdam and APM Terminals. The comparison focuses on indicators that are less sensitive to scope and scale differences, uses like-for-like filters when possible, and records definitional differences explicitly. When differences are observed, a plausible explanation is provided and the effect on conclusions is assessed.

Anchorage

Figure E.1 shows the percentage of deepsea and feeder vessels proceeding directly to berth versus anchorage for the Port of Rotterdam. Figure E.2 shows the same indicator in the simulation. The simulated shares fall within the range of the observed system and exhibit a similar variability, approximately eight percentage points. The alignment is considered satisfactory for a plausibility check.

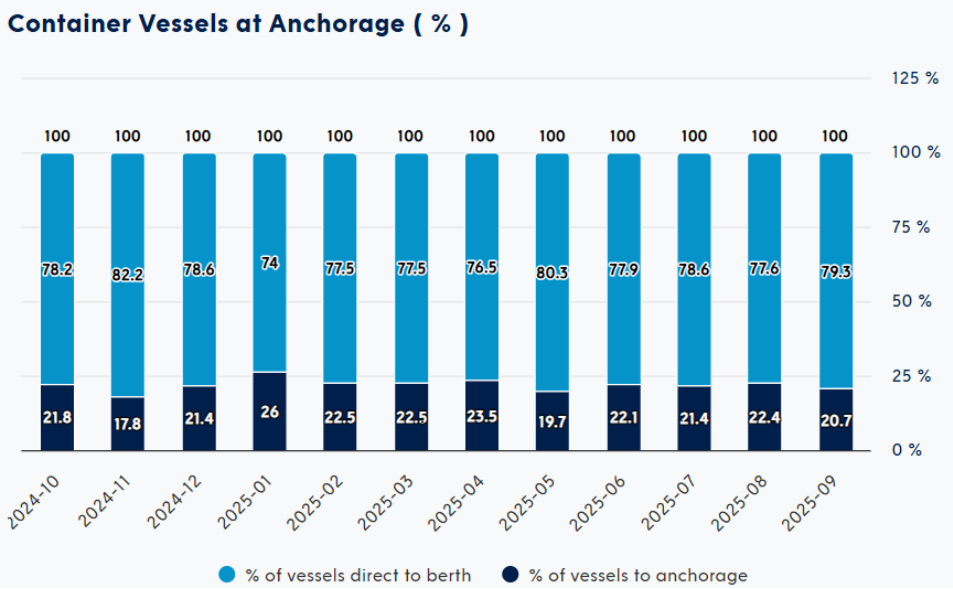

Figure E.1: Percentage of deepsea and feeder ships going directly to berth versus to anchorage in the Port of Rotterdam (Port of Rotterdam, 2025b).

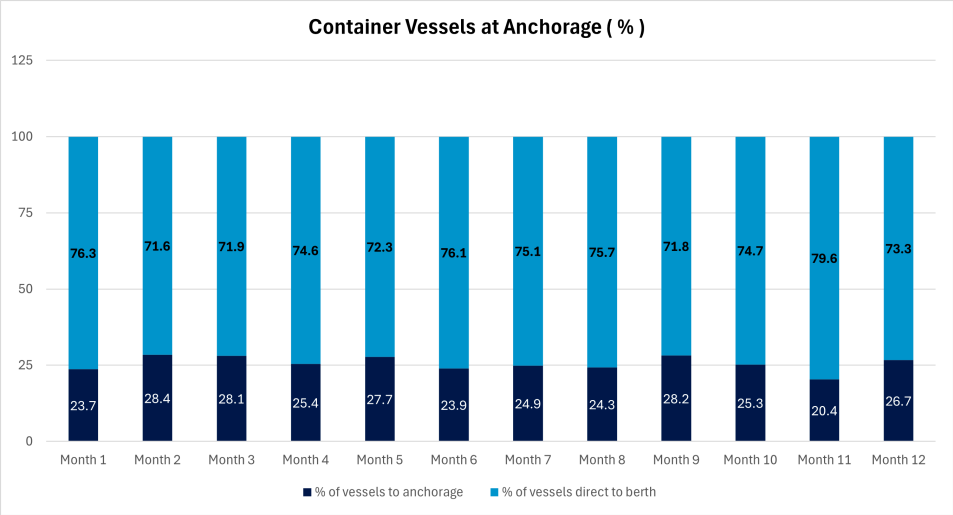

Figure E.2: Percentage of deepsea and feeder ships going directly to berth versus to anchorage in the simulation

Turnaround time

For each vessel class, Port of Rotterdam indicators are compared with simulation results. In the Port of Rotterdam statistics, time in port comprises sailing from sea through the Nieuwe Waterweg to one or more berths and back out past the Maasvlakte, time at berth measures the berth interval per call. Waiting is reported per waiting event. In the simulation, the same definitions are used, except for the sailing boundary, which is measured between the entrance and departure nodes located at the opening of Maasvlakte 2. Sailing time in the simulation is therefore shorter. For fair comparison, an additive adjustment to sailing time is noted when interpreting differences.

Deepsea The model's deepsea call-size distribution averages ∼4,467 TEU, consistent with the Port of Rotterdam context. Public indicators for time at berth and time in port are shown in Figures E.3a–E.3b. The scenarios assume *no downtime* during handling (no planned/unplanned crane outages, weather stops, or shift-change effects beyond the nominal rate), so simulated berth times constitute an optimistic lower bound relative to public figures that include such effects. The simulation yields a mean at-berth time of 25.22 h, with waiting 0.51 h and internal sailing 0.45 h (Table E.1). Because the model's sailing boundary is at Maasvlakte 2, external approach/departure along the Nieuwe Waterweg is excluded; once this boundary difference is acknowledged, the non-berth portion aligns in order of magnitude. The remaining gap to public berth-time averages can be explained by operational overheads (mooring/unmooring, lashing, hatch work, crane changeovers) and effective productivity below the nominal rate. Applying a plausible effective crane efficiency (e.g., ∼80%) plus modest fixed berth overheads brings the implied service time into the public range.

	Time at berth (h)	Waiting (h)	Sailing within boundary (h)
Deepsea (simulation)	25.22	0.51	0.45

Table E.1: Deepsea simulation summary (scenario mean values). Sailing excludes the external approach/departure beyond Maasvlakte 2.

Statistics Port of Rotterdam 2024 (Port of Rotterdam, 2025b):

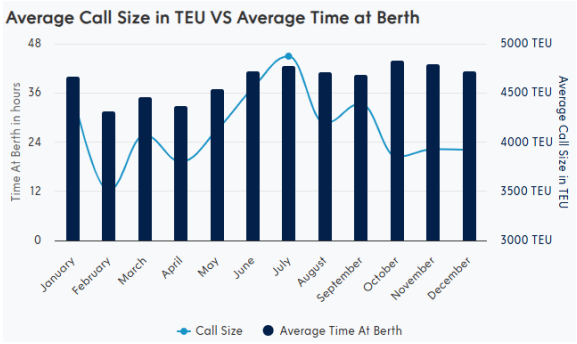

(a) Deepsea time at berth

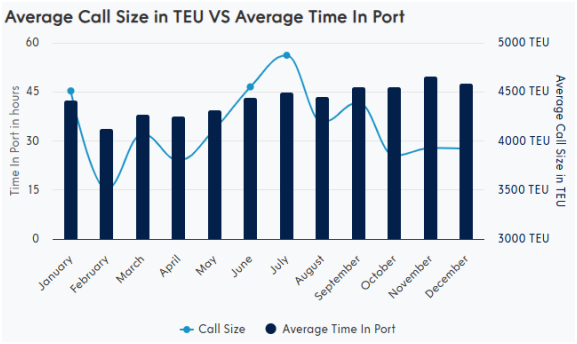

(b) Deepsea time at port

Feeder Public indicators report an average time at berth of about 12.3 h for feeders, with time in port around 41.4 h. The simulation reproduces the berth component closely for comparable call sizes ($\sim 1{,}000$ TEU), but yields a much shorter time in port because scope and routing are simplified. In the scenarios, feeders make a single terminal call on Maasvlakte 2, sailing is measured only inside the Maasvlakte boundary, and handling assumes no downtime. By contrast, the public indicator covers the full approach/departure via the Nieuwe Waterweg, includes pre-arrival and inter-call queues, and reflects that feeders often visit multiple terminals per rotation. These factors add substantial sailing and waiting outside the modelled boundary and between calls.

Over a ten-week window the simulation averages 12.33 h at berth, 0.17 h waiting, and 0.56 h internal sailing, for a model time in port of 12.84 h (Table E.2). The residual gap to the public 41.4 h figure is therefore attributed to (i) external approach/departure that is out of scope in the model, (ii) multi-terminal rotations with inter-berth sailing and queuing, and (iii) short operational overheads and micro-interruptions that are not explicitly modelled.

	Call size (TEU)	At berth (h)	Waiting (h)	Sailing (in-boundary, h)
Feeder (simulation mean)	1001.14	12.33	0.17	0.56

Table E.2: Feeder simulation summary (ten-week average). Time in port (model) $=$ 12.84 h; external approach/departure is excluded by design.

Interpretation: Given matched call sizes and aligned berth times, the port-time discrepancy primarily reflects scope differences (geographical boundary and number of calls per visit) and the no-downtime assumption in the simulation. For feeders, the most like-for-like validation target is therefore *time at berth*, while *time in port* is used qualitatively with these differences noted.

Statistics Port of Rotterdam 2024 (Port of Rotterdam, 2025b):

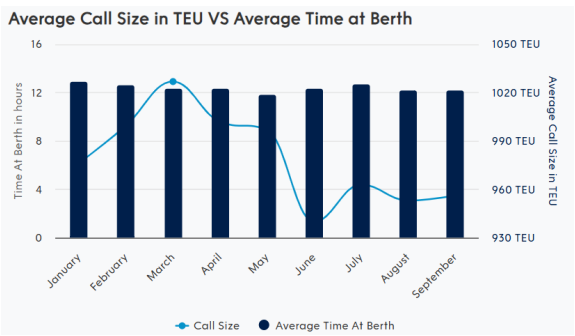

(a) Feeder time at berth

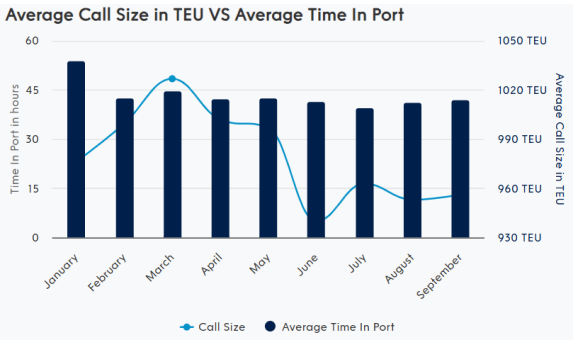

(b) Feeder time at port

Results simulation:

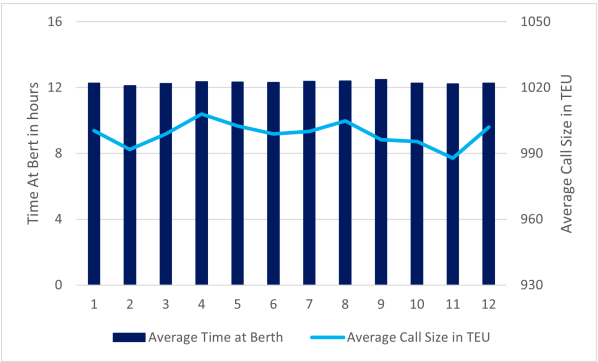

Figure E.5: Time at berth for feeders in the simulation

Barge Statistics Port of Rotterdam barge performance monitor(Port of Rotterdam, 2025a):

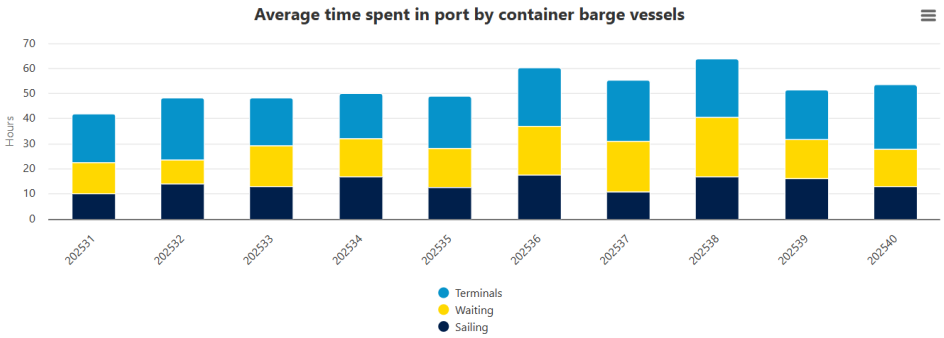

Figure E.6: Average time spent in the port by container barge vessels

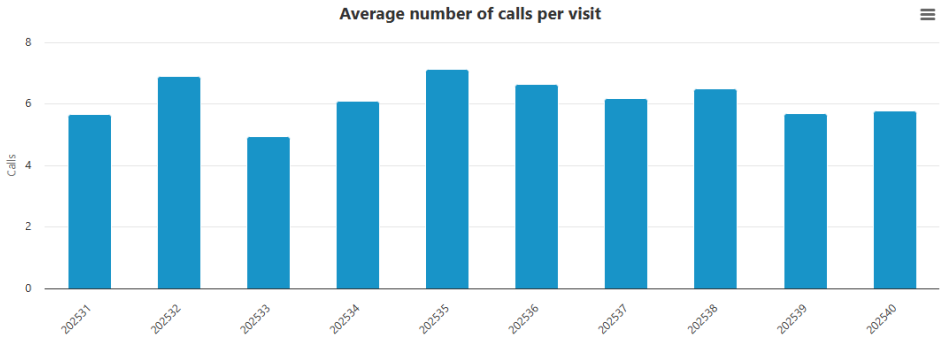

Figure E.7: Average number of calls per visit in the Port of Rotterdam

The simulation represents five terminals, whereas the Port of Rotterdam comprises a larger set. As a result, average calls per visit and time at terminals in the Port indicators span a wider range than in the model. Figure E.6 reports the Port of Rotterdam barge performance monitor for average time in port, and Figure E.7 shows the corresponding average number of calls per visit. The simulated values in Table E.3 are of the same order once scope differences are considered. Two comparability notes apply. First, sailing time in the Port indicators covers the full approach and departure via the Nieuwe Waterweg, while the simulation measures sailing within a boundary at Maasvlakte 2, therefore sailing components are not directly comparable and are excluded from quantitative comparison. Second, the scenarios assume no downtime during handling, so simulated service times represent an optimistic lower bound relative to public indicators that include outages and set-up losses.

In the Port indicators the average number of calls per barge visit is approximately six (Figure E.7), whereas the simulation yields a mean of 2.87 calls. Given that time at terminals scales with the number and dispersion of calls per rotation, shorter simulated terminal times are expected. Table E.3 summarises the simulated barge turnaround components used for comparison.

	Time at terminals (h)	Waiting (h)	Avg. calls per visit
Barge	11.9	9.4	2.87

Table E.3: Simulation: average barge turnaround components and calls per visit. Sailing is excluded from comparison due to boundary differences.

Terminal Productivity

The model's crane productivity is compared with the weekly performance snapshot reported by APM Terminals Maasvlakte 2. The public dashboard reports deepsea productivity, barge productivity, rail productivity, and truck turn-around time (APM Terminals, n.d.). For comparability, the model reports crane moves per hour for conventional quay cranes, which are used by both barges and deepsea vessels.

Figure E.8: APMT performance monitors: (a) barge and (b) deepsea.

Table E.4 lists the simulated weekly averages for crane moves per hour. The values are consistent with the public snapshot ranges, which supports the plausibility of the configured crane rates.

Week	Crane moves per hour
Week 1	25.94
Week 2	26.24
Week 3	26.41
Week 4	26.01

Table E.4: Simulation: weekly average crane moves per hour (conventional quay cranes).

Comparability notes - (i) Public snapshots aggregate across wider scopes and time windows than the model scenarios, (ii) the simulation uses a reduced terminal set, (iii) the sailing boundary differs, which shifts time-in-port components. These differences are recorded and considered when interpreting any residual gaps.

Synthesis - Across anchorage share, time-in-port components, and crane productivity, simulated values fall within plausible ranges of public indicators after accounting for scope and definition differences. Residual discrepancies are explainable by known simplifications in the terminal set, the sailing boundary, and scenario-specific call-size distributions. The validation is therefore considered sufficient for the comparative analysis of coordination strategies.

F

Results

F.1. Experiment: Different Barge and Module Ratios

Table F.1: Average time at terminals and waiting by vessel class and barge-module mix (hours).

Mix	Vessel class	Average. time at terminals	Average time waiting
100-0	Barge	7.61	23.62
	Deep-sea	28.32	0.74
	Feeder	15.51	0.36
	Module (child)	0.00	0.00
75-25	Barge	7.88	14.52
	Deep-sea	27.22	0.65
	Feeder	14.32	0.44
	Module (child)	0.55	0.14
50-50	Barge	8.07	6.72
	Deep-sea	26.17	0.38
	Feeder	13.34	0.21
	Module (child)	0.55	0.20
25-75	Barge	8.20	1.84
	Deep-sea	25.48	0.28
	Feeder	12.69	0.10
	Module (child)	0.55	0.60
0-100	Barge	0.00	0.00
	Deep-sea	25.10	0.27
	Feeder	12.46	0.10
	Module (child)	0.55	25.27

F.2. Experiment: Demand

Table F.2 summarises how anchorage behaviour changes with sea freight demand for the 100-0, 50-50 and 25-75 mixes. For each demand level, it reports the average number of arrivals routed via anchorage or directly to berth, together with the average waiting time in anchorage per call for each vessel class.

Table F.2: Anchorage behaviour by demand level, mix and vessel class. “Direct” and “Anch.” give the average number of arrivals of that class routed directly to berth or via anchorage; shares are percentages of total arrivals of that vessel class; “Wait” is the average anchorage waiting time per call (hours).

Mix	Demand	Vessel type	Direct	Direct share (%)	Anch.	Anch. share (%)	Wait (h)
				100:0			
100:0	80%	Deep-sea	119.9	91.4	11.3	8.6	0.40
100:0	100%	Deep-sea	139.6	85.0	24.7	15.0	0.74
100:0	120%	Deep-sea	147.3	75.0	49.0	25.0	1.69
100:0	80%	Feeder	65.0	96.7	2.2	3.3	0.19
100:0	100%	Feeder	79.4	93.9	5.2	6.1	0.36
100:0	120%	Feeder	89.7	89.0	11.1	11.0	1.04
100:0	80%	Barge	1989.8	74.6	676.0	25.4	5.35
100:0	100%	Barge	1725.8	68.5	792.1	31.5	8.36
100:0	120%	Barge	1445.9	61.0	925.9	39.0	10.88
				50:50			
50:50	80%	Deep-sea	123.7	94.0	7.9	6.0	0.15
50:50	100%	Deep-sea	147.3	89.3	17.6	10.7	0.36
50:50	120%	Deep-sea	161.4	81.7	36.2	18.3	0.83
50:50	80%	Feeder	65.8	97.6	1.6	2.4	0.05
50:50	100%	Feeder	81.0	95.2	4.1	4.8	0.19
50:50	120%	Feeder	92.5	91.0	9.2	9.0	0.43
50:50	80%	Barge	1380.4	91.4	130.8	8.6	1.07
50:50	100%	Barge	1274.2	85.8	211.4	14.2	2.14
50:50	120%	Barge	1150.9	79.0	307.0	21.0	3.67
50:50	80%	Module	3924.4	73.6	1411.1	26.4	0.11
50:50	100%	Module	3854.6	72.3	1477.7	27.7	0.20
50:50	120%	Module	3709.8	69.7	1613.8	30.3	0.51
				25:75			
25:75	80%	Deep-sea	126.2	95.8	5.6	4.2	0.11
25:75	100%	Deep-sea	151.6	91.9	13.3	8.1	0.27
25:75	120%	Deep-sea	169.2	85.5	28.8	14.5	0.63
25:75	80%	Feeder	66.2	98.1	1.3	1.9	0.04
25:75	100%	Feeder	81.5	95.9	3.5	4.1	0.10
25:75	120%	Feeder	94.3	92.6	7.5	7.4	0.33
25:75	80%	Barge	739.6	97.0	23.1	3.0	0.18
25:75	100%	Barge	710.6	93.6	48.2	6.4	0.54
25:75	120%	Barge	672.3	88.4	88.4	11.6	1.21
25:75	80%	Module	3006.1	37.6	4979.0	62.4	0.38
25:75	100%	Module	2919.9	36.5	5071.2	63.5	0.57
25:75	120%	Module	2735.4	34.3	5247.4	65.7	1.10

Across all three mixes, higher sea freight demand increases both the number of calls that pass through anchorage and the average waiting time. The conventional 100-0 case shows the strongest growth: barge calls see a marked rise in anchorage use and waiting, from moderate staging at 80 % demand to prolonged off-berth times at 120 % demand. Deep-sea and feeder calls also experience longer anchorage waits as demand rises, but their absolute waiting times remain much smaller than for barges.

In the mixed systems, the additional staging effort is redistributed. For the 50-50 and 25-75 mixes, deep-sea and feeder anchorage waiting times remain short across the full 80–120 % band, and barge anchorage is substantially lower than in the 100-0 case at the same demand levels. Instead, most

anchorage activity is concentrated in the module (child) fleet. Module calls already exhibit frequent, short anchorage episodes at 80 % demand, and their average waiting time increases with demand, particularly in the 25-75 mix. This pattern shows that the mixed concepts use modules as a flexible buffer to absorb timing mismatches and berth-window constraints, while keeping anchorage for deep-sea and feeders limited and reducing staging requirements for conventional barges compared with the purely barge-based system.

Overall, the demand experiment indicates that within the tested range the port remains below a congestion regime. Rising sea freight intensity increases anchorage and waiting in all configurations, but in the mixed cases most of this additional staging is taken up by modular inland traffic rather than by sea-going vessels.

F.3. Experiment: Early Modular Adoption Scenario

Table F.3: Barge anchorage behaviour for baseline and early modular adoption futures. The anchored share is the fraction of barge calls that visit anchorage at least once.

Scenario	Anchored share (%)	Average anchorage wait (h)
Baseline	31.5	8.36
100:10, dedicated module cranes (1.5 min/container)		
+1MC	33.0	8.43
+2MC	33.2	8.45
100:10, modules on quay cranes (1.5 min/container)		
0QC	33.7	9.43
+1QC	30.1	6.72
+2QC	25.0	4.93
100:10, modules on quay cranes (3.0 min/container)		
0QC	34.0	8.23
+1QC	31.3	6.92
+2QC	26.1	5.40