Master thesis

Hub routing as an alternative to direct calls for third-party barge profiles

TIL5060: Thesis project

Naut Linssen

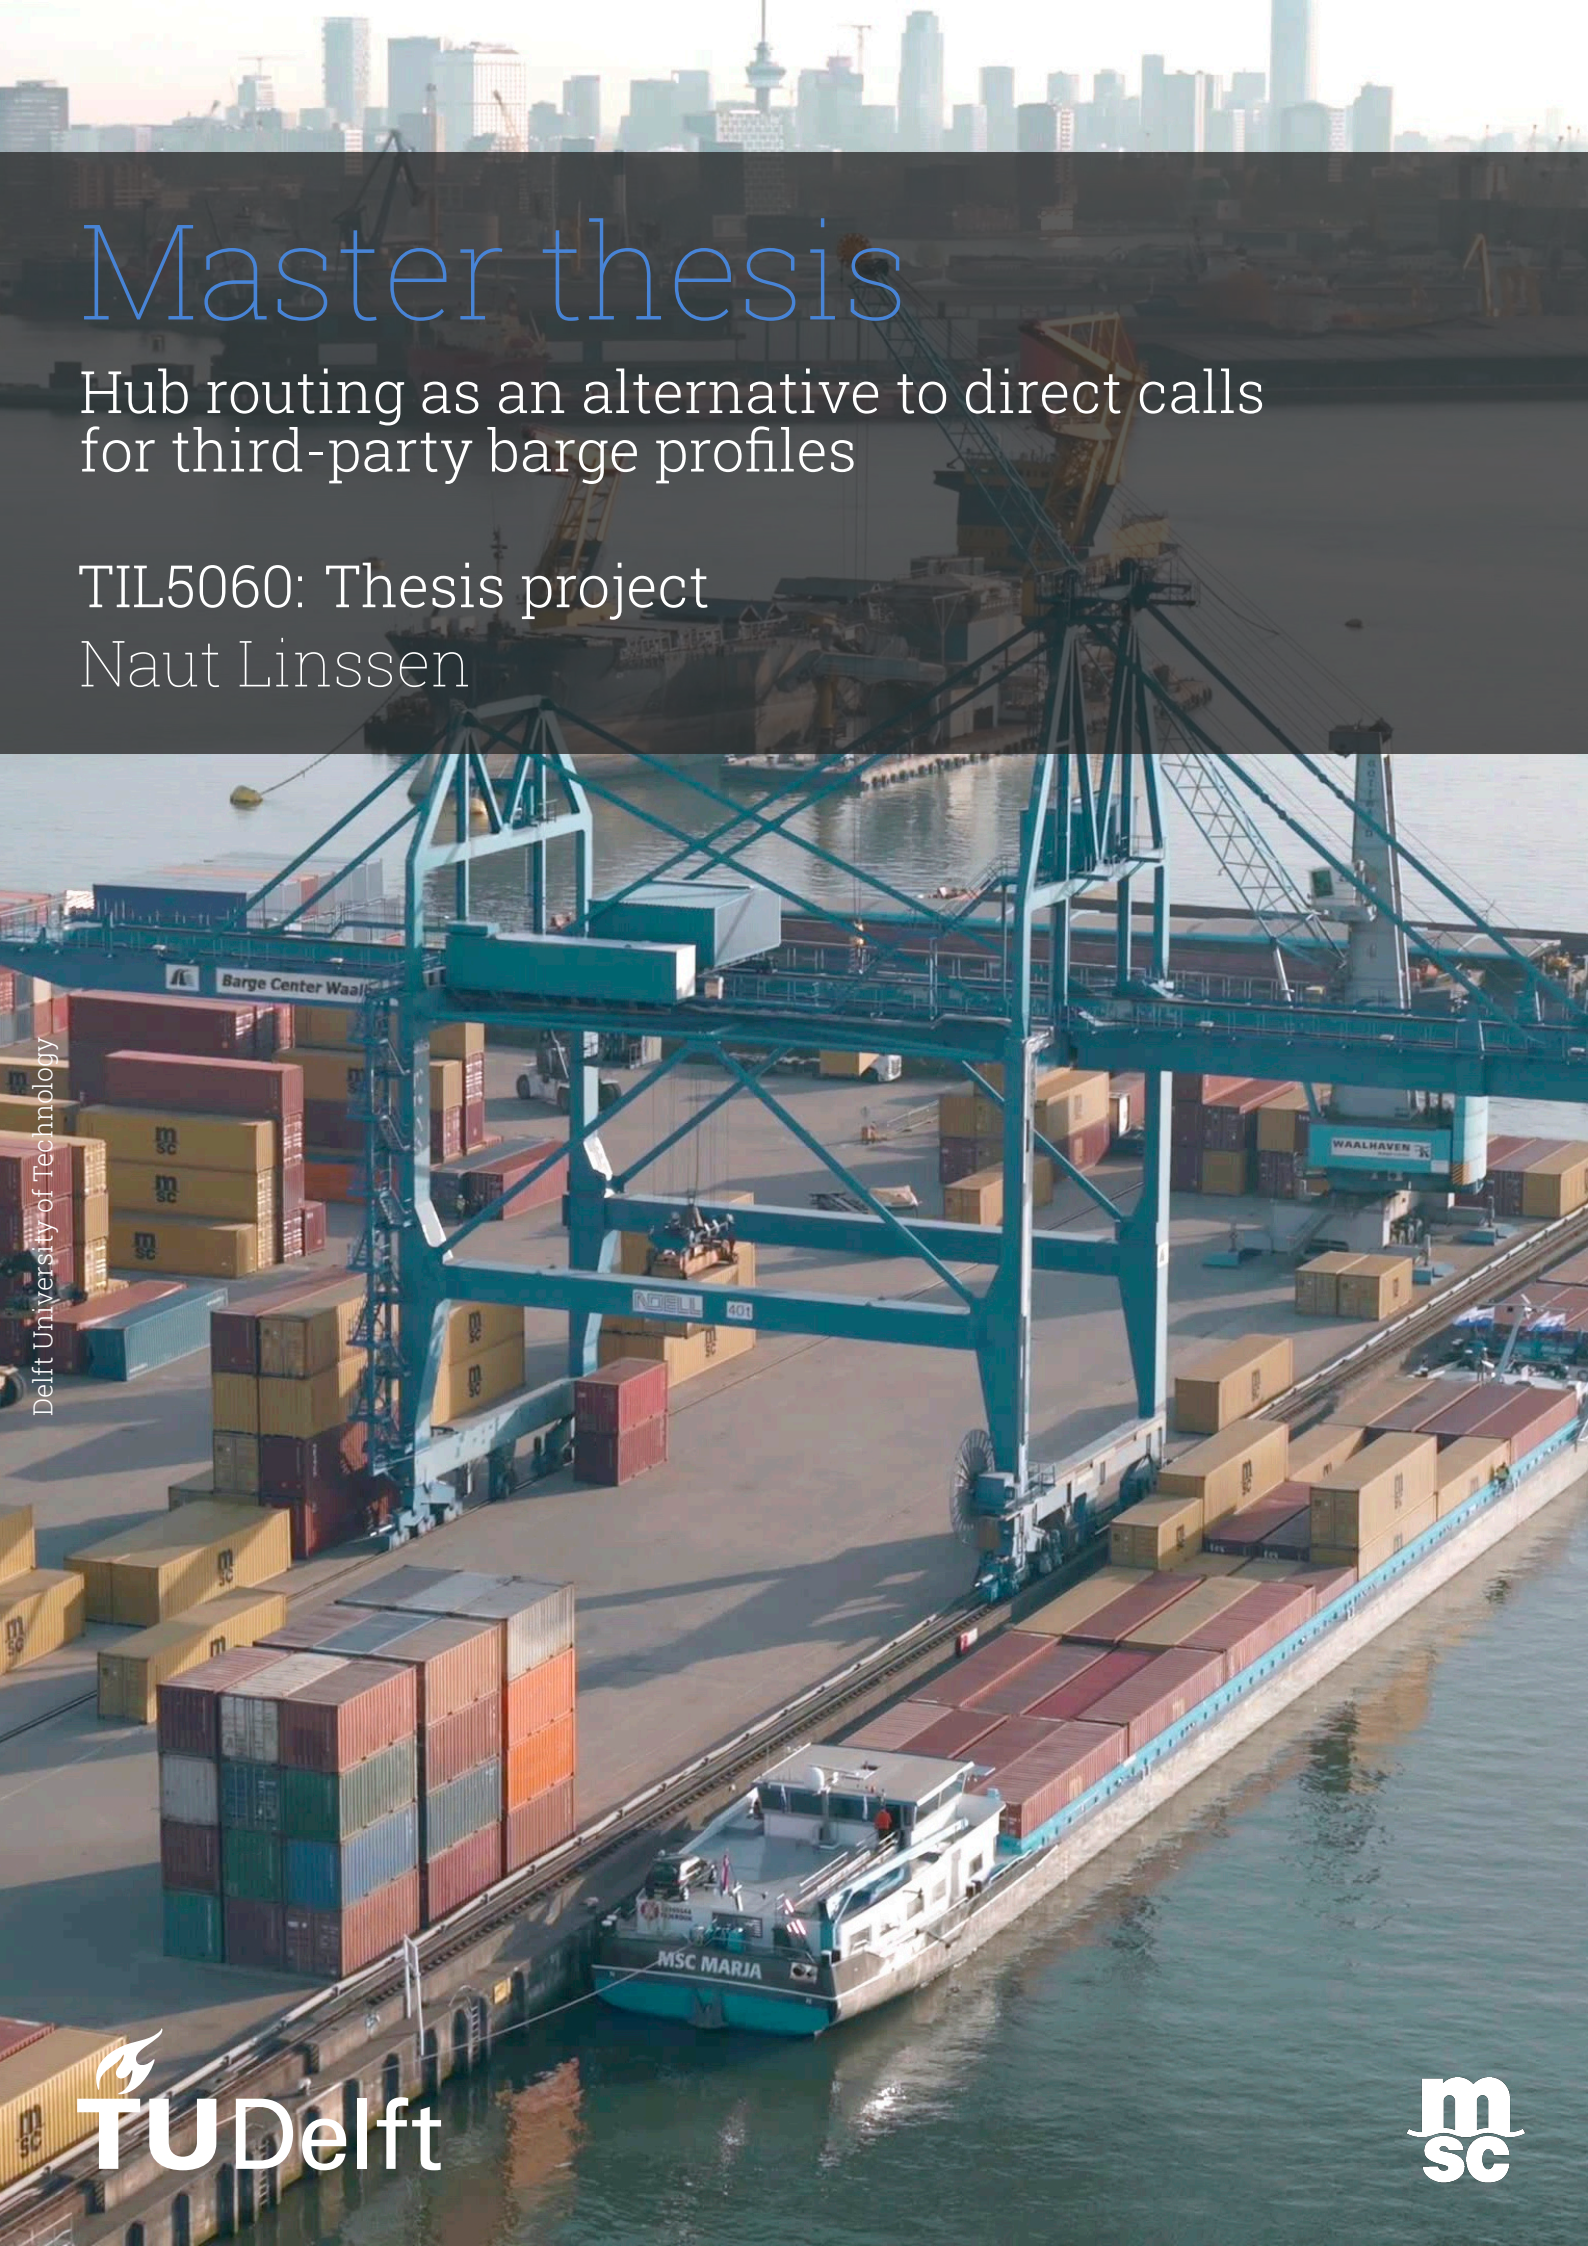

Delft University of Technology

TU Delft

msc

Master thesis

Hub routing as an alternative to direct calls for third-party barge profiles

by

Naut Linssen

Transport, Infrastructure and Logistics

Faculty of Civil Engineering and Geosciences

Student number: 5032563

Thesis committee:	Dr. S. (Stefano) Fazi	TU Delft, Chair and First supervisor
	Prof. dr. ir. E.B.H.J. (Edwin) van Hassel	TU Delft, Second supervisor
	C. (Casper) van der Woude	MSC, Company supervisor

Cover: Barge Center Waalhaven (Credits: MSC Nederland)

TU Delft

msc

Preface

This thesis represents the final part of my study period at the TU Delft. I have learned a lot in the past years, which have all been put to good use in the making of this thesis. The making of this thesis was a journey with ups and downs, but eventually it resulted in a thesis that I am proud of. I have learned a lot and I am happy to take that with me to the next chapter.

First, I want to thank my supervisors of the TU Delft. I would like to thank Stefano Fazi for being my first supervisor and chair. The clear and sometimes critical feedback from an academic point of view helped me write a thesis on the academic level. I also would like to thank my second supervisor Edwin van Hassel. You interested me for the maritime sector with the course 'Shipping management', which eventually brought me to MSC where I did my internship for this thesis. In addition to making me enthusiastic about the maritime sector, I thank you for the help I received from you with your feedback and especially the advice when I did not know what to do.

Furthermore, I would like to thank Giuseppe for the opportunity and the trust to start my thesis at MSC. I would like to thank my supervisor at MSC, Casper van der Woude, for his guidance and useful insights during the making of the thesis. In addition, the invitations to different meetings helped me to speak to the right people for my thesis and getting new insights I could use in the further development of the project. A special thanks to Mark who helped me with the crucial AIS data and Olivier for the fresh insights of the do's and don'ts during the making of a thesis. Furthermore, I also would like to thank Ilona and Duncan for sharing useful details about Intermodal or the ECT terminals whenever it was needed and for the beautiful photos, one of which is on the cover.

Finally, I would like to thank my family and friends for their support and positive words whenever I needed them. Not only during my thesis, but throughout my whole study period. Thank you for your patience and interest when I explained everything that I could not figure out, was proud of, or sometimes simply did not enjoy studying. The more I could explain it to you, the more I understood it myself or realised what needed to change.

Thank you all for being part of my journey, and I wish the reader as much joy while reading it as I had while finishing the thesis.

Naut Linssen
Delft, November 2026

Summary

Barge congestion at deepsea terminals is still a large issue. Deepsea vessels and barges with contractual agreements get priority, so the in-port time of other inland barges is increasing. The literature describes hub-and-spoke systems as a possible solution. In the Port of Rotterdam the Barge Center Waalhaven (BCW) is partly used as such a hub, with a shuttle service to the deepsea terminals at the Maasvlakte.

Hub-and-spoke systems have so far been evaluated as designs at network level. Whether a third-party operator has a cost reason to use a hub that already runs and charges a real tariff is not known. This thesis asks to what extent routing third-party inland barges through the hub changes their turnaround time and transport cost per TEU, and for which operator the hub is a promising alternative.

The direct route and hub route are both modelled with discrete-event simulation. AIS data of inland barges is used to set the parameters and to validate the direct-call simulation. Six scenarios are formed by three vessel classes (M6, M8 and M9) and two inland-leg profiles. Each scenario is run for thirty replications with common random numbers, so the two routings are compared as paired differences. A cost layer converts the time output into cost per TEU, using published cost rates per vessel class and the shuttle tariff that is charged in practice.

The direct call is the cheaper routing in all six scenarios. The hub costs 63.85 to 84.07 EUR per TEU more. The main cause is the shuttle fee of 67.50 EUR per TEU, while a barge without a fixed window pays nothing for a call at the deepsea terminal. On time the hub is faster only for the M8 class, by 2.29 to 2.68 hours per round trip. The throughput of the system is capped at about 540 containers per week by the shuttle and not by the hub, so the call size of the inland barges is a system outcome rather than a vessel property.

Four experiments were run on the M6 class. Three of them change the operation of the hub and each frees enough capacity to run one barge fewer. This lowers the gap to between 56.14 and 58.01 EUR per TEU. Without the fee the hub would then be the cheaper operation, but no fee exists that is low enough for the operator and high enough to cover the cost of the shuttle. The fourth experiment congests the deepsea terminal. The direct call then needs eight barges instead of seven on the long leg to move the same volume, and the queue at the terminal grows. The hub situation becomes faster in every case, by up to 39.08 hours per round trip on the long leg. The cost gap narrows to 30.23 to 32.02 EUR per TEU but does not close.

The hub-and-spoke system as it runs today is not profitable for third-party inland barge operators in the scenarios tested and the fee is the main reason. The size of the fleet is the main lever for both time and cost, because every barge added to a single-berth system joins the same queue. Under congestion an operator that values a short and predictable turnaround may accept paying about 30 EUR per TEU extra for a round trip that is more than a day shorter. The answer could change for profiles that were not tested, such as barges that make several deepsea calls per cycle, barges smaller than the M6 class and a cheaper shuttle operation.

Contents

Preface i

Summary ii

Nomenclature vii

1 Introduction 1

2 Literature review 4
2.1 Barge congestion and its causes . . . 4
2.2 Consolidation strategies . . . 4
2.3 Bundling and network types . . . 5
2.4 Hub-and-spoke . . . 6
2.5 Cost performance of the hub-and-spoke network . . . 7
2.6 Literature synthesis . . . 8

3 Problem description 10
3.1 Current state . . . 10
3.2 Problem statement . . . 12
3.3 Research questions . . . 14

4 Methodology 15
4.1 Research approach . . . 15
4.2 System description . . . 16
4.3 Conceptual model and simplifying assumptions . . . 17
4.4 Simulation model . . . 18
4.4.1 Container arrivals and deepsea terminals . . . 18
4.4.2 Shuttle service and fixed windows . . . 18
4.4.3 Hub operations and opening hours . . . 19
4.4.4 Inland barges and the inland leg . . . 19
4.4.5 Current-state simulation . . . 19
4.5 Input modelling . . . 20
4.5.1 From AIS data to visits . . . 20
4.5.2 The inland leg . . . 20
4.5.3 Distribution selection . . . 20
4.6 Cost layer . . . 21
4.7 Conclusion . . . 22

5 Data analysis 23
5.1 AIS data . . . 23
5.1.1 Deepsea terminal visits . . . 23
5.1.2 Inland-leg distribution . . . 24
5.2 Model parameters . . . 25

6 Verification and validation 28
6.1 Framework and scope . . . 28
6.2 Conceptual model validation . . . 28
6.3 Verification and white-box validation . . . 29
6.3.1 Traces . . . 29
6.3.2 Internal consistency checks . . . 29
6.3.3 Extreme condition tests . . . 29
6.3.4 White-box parameter checks . . . 31

6.3.5 Coverage of the verification . . . 31
6.4 Operational validation . . . 31

7 Numerical results **33**
7.1 Experimental framework and performance indicators . . . 33
7.2 Warm-up period . . . 34
7.3 Replications and statistical analysis . . . 35
7.4 Scenarios and fleet calibration . . . 35
7.4.1 Fleet-size calibration . . . 35
7.4.2 Scenario definition . . . 36
7.5 Baseline results . . . 36
7.5.1 Steady state performance of the hub-and-spoke . . . 36
7.5.2 Performance of the current state . . . 38
7.5.3 Decomposition of hub-side waiting time . . . 38
7.5.4 Current vs hub situation . . . 39
7.5.5 Paired differences hub and current situation . . . 40
7.5.6 Short versus long inland-leg . . . 41
7.5.7 Summary of findings . . . 41
7.6 Experiments on the M6 configuration . . . 42
7.6.1 Diagnosis of the berth queue . . . 43
7.6.2 Results of the experiments . . . 43
7.6.3 The congestion experiment . . . 44
7.6.4 Summary of the experiments . . . 45

8 Discussion **46**
8.1 The shuttle fee decides the cost comparison . . . 46
8.2 Fleet size is the main lever for time and cost . . . 46
8.3 Operator profiles for which the hub could be attractive . . . 47
8.3.1 Barges with a long waiting time in port . . . 47
8.3.2 Barges with more than one call . . . 47
8.3.3 Smaller barges and smaller calls . . . 48
8.3.4 A cheaper shuttle service . . . 48
8.3.5 The value of reliability . . . 48
8.4 Limitations . . . 48
8.5 Further research . . . 49

9 Conclusion **50**
9.1 Answers to the sub-questions . . . 50
9.2 Answer to the main research question . . . 51
9.3 Recommendations . . . 52
9.4 Contribution . . . 52

References **54**

A Scientific paper **56**

B Assumptions and simplifications **71**

C Model graphs **73**
C.1 SPC warm-up control charts . . . 73
C.2 Hub stock . . . 74
C.3 Waiting-time . . . 75

D Verification and validation **77**
D.1 Trace verification of the model . . . 77

List of Figures

2.1 Line bundling network. Source: Janic et al. (1999) 5
2.2 Trunk line with collection and distribution network. Source: Janic et al. (1999) 5
2.3 Hub-and-spoke design. Source: Konings et al. (2013) 6

3.1 Visualisation of the current state of the barge network 11
3.2 Visualisation of the current state zoom-in on the Maasvlakte 12
3.3 Visualisation of the future hub-and-spoke system with BCW as the hub 13

4.1 Flowchart of the current barge process 16
4.2 Flowchart of the hub-and-spoke process 17
4.3 Conceptual model of the hub-and-spoke network 18

5.1 Weibull fits for the sailing and dwell components of the short and long inland-leg profiles, with bias-corrected KS p-values. 25

6.1 Hub stock over a full run on the M6 long profile 30
6.2 Container flow per replication under the baseline and under two extreme settings 30

7.1 Warm-up period against mean barge turnaround time for scenario M6 with a long inland leg 34
7.2 Mean number of barges in the hub berth queue per hour of the week, M6 long leg, 30 replications. 43

C.1 Warm-up period of the mean barge turnaround time for all six scenarios 73
C.2 Stock of the hub per class on the short and long leg 74
C.3 Waiting time at the hub per class on the short and long leg 75

D.1 Output print statements hub-and-spoke simulation 78
D.2 Output print statements current state simulation 79

List of Tables

5.1 Candidate distributions for the inland-leg components. AIC per family (lower is better), fitted Weibull parameters and bias-corrected KS p-values. 24
5.2 Parameter values of the case study 26

6.1 Container flow identities asserted at the end of every replication, over the 30 replications of the M6 long baseline. 29
6.2 Extreme condition test on the current state simulation: no container arrivals at the deepsea terminal 31
6.3 Operational validation of the current state simulation 32

7.1 Warm-up period per scenario, obtained with the SPC method of Robinson (2007) on the inland barge turnaround time. 34
7.2 Fleet-size calibration per scenario at the hub. 35
7.3 Hub performance of the baseline scenarios 36
7.4 Call size and vessel capacity utilisation per baseline scenario. The minimum frequency is the visit rate at which the weekly supply would be moved with fully loaded vessels. . . 37
7.5 Direct-call performance of the baseline scenarios 38
7.6 Decomposition of hub-side waiting time per barge visit 39
7.7 Decomposition of the time-cost advantage per deepsea call 39
7.8 Paired differences between the hub and the direct-call situation 40
7.9 Paired differences in time and cost (long minus short per vessel class) 41
7.10 Decomposition of the cost gap into sailing, waiting and fees. Waiting includes the waiting at the hub or the deepsea terminal, the loading time and the dwell at the inland terminal, all at the waiting rate. 41
7.11 Definition of the experiments 42
7.12 The operational experiments against the direct-call situation 43
7.13 Results of the congestion experiment 44

B.1 Assumption register for the conceptual model 71

Nomenclature

Abbreviations

Abbreviation	Definition
AIC	Akaike information criterion
AIS	Automatic Identification System
BCW	Barge Center Waalhaven
CC	Cargo consolidation
CER	Container Exchange Route
CFS	Container freight station
CI	Confidence interval
DDE	Delta Dedicated East (ECT terminal)
DDN	Delta Dedicated North (ECT terminal)
DES	Discrete-event simulation
DST1, DST2	Deepsea terminal 1 and 2
ECT	Europe Container Terminals
FCL	Full container load
FEU	Forty-foot equivalent unit
ITT	Inter-terminal transport
IWT	Inland waterway transport
KPI	Key performance indicator
KS	Kolmogorov–Smirnov
LCL	Less than container load
MMT	Modular Mobile Terminal
MSC	Mediterranean Shipping Company
SPC	Statistical process control
TEU	Twenty-foot equivalent unit
TT	Turnaround time

1

Introduction

Nowadays, the container is at the core of our transport system and the container throughput of ports is very large (Shenkar & Rosen, 2018). For example, the container throughput of the port of Rotterdam was 133.4 million tonnes and 13.8 million TEU in 2024 (Port of Rotterdam, 2025) and that is not the only port. The container transport continues to rise globally, with container vessels that have room for more than 24,000 TEU. The average call size on such vessels in Rotterdam is 8,300 TEU (Roux et al., 2023). This is the result of pursuing economies of scale (Zheng & Yang, 2016), making the ports expand the container terminals' capacity. Due to the expansion of the container terminals, the hinterland capacity and performance becomes more relevant as well. The hinterland transportation connects countries and regions without a connection to a sea or seaport after all. From the seaport containers are transported within the hinterland, by trucks, trains, and barges. The transportation mode can differ by the destination of the containers and sometimes a combination of transport modes is used. For example, the container is transported from the deepsea terminal per barge and then dropped at an inland terminal where a truck can take the container to the final destination.

Over the years, Inland Waterway Transport (IWT) has become an important way of transporting containers into the hinterland. Due to the cost-efficient and sustainable transport of barges, this transport mode generates higher economies of scale than other modes of transport (Nicolet et al., 2023), such as rail or truck transport. In Rotterdam 3 million TEU were transported by barge in 2020, which is equal to 32% of the market share. The Port of Antwerp shipped 2.5 million TEU by barge with a market share of 35% (Roux et al., 2023). Due to the sustainable future of the IWT, The European Commission aims to increase transport by inland waterways and short sea shipping by 25% by 2030 and by 50% by 2050 (European Commission, 2021). The infrastructure costs, truck congestion at the ports, and the sustainable costs confirm the urge to move to IWT. The infrastructure cost of barge transportation is 12.60 EUR per thousand ton-kilometers. For the train and truck this is respectively 45.21 and 48.42 EUR (Bu & Nachtmann, 2023). On the environmental cost, barge transport scores good as well compared to the other transportation modes. Transporting containers by barge is 0.27 cents per ton per kilometer compared to 2.01 cents by train and 0.80 cents by truck (Bu & Nachtmann, 2023). The last argument to use barge transport for the hinterland transportation is the truck congestion at the port of Rotterdam. 40% of the trucks experience heavy delays that can be reduced by the switch of container transport from truck to barge (Bu & Nachtmann, 2023).

Nevertheless, the barge sector faces challenges. In Europe these challenges are mainly focused on the largest ports such as the port of Rotterdam and the port of Antwerp (Shobayo & van Hassel, 2019). One of the main challenges is the persistent barge congestion at deepsea terminals. This congestion has multiple underlying causes, including the extended time barges spend in the port. A key contributor to these longer port stays is the relatively small volume of containers handled per barge call. Therefore, barges need to make calls (6 to 8) across various deepsea terminals to be fully loaded or unloaded. When a delay appears at one terminal it can easily lead to missing a time window at the following terminals. This can cause a domino-effect that leads up to long delays (Roux et al., 2023). Another reason is the prioritisation of larger vessels. The berthing and handling capacity at seaports are still

limited for which ocean-going vessels, feeder vessels and smaller barges need to compete at the same time. During peak times, peak loads appear at the terminals, which create bottlenecks in the handling capacity. The result is priority of ocean-going and feeders' vessels over smaller inland vessels, because larger vessels have higher operating costs (Friedhoff et al., 2023). A third factor that plays a role in the small space available for barges is the fact that there are no contractual relations between barge operators and deepsea terminals according to the interviewees in the research of Oganesian et al. (2021). Because deepsea vessels do have a contract with the deepsea terminals, these vessels are prioritised over barges. Taking into account that barges also need to visit different deepsea terminals to load and unload containers, the result is a long waiting line of barges at the deepsea terminals.

According to a recent study from Vermaat (2026) the average time barges from the upper Rhine spend in ports can increase to 52 hours. The research of Nicolet et al. (2023) raises the bar even further and mentions a waiting time of 120 hours in the port of Rotterdam. Reducing the barge in-port times can lead to shorter transport times, and a higher chance of more barge round trips, which can eventually lead to lower transportation costs (Van Es, 2019). In the general literature the hub-and-spoke model is introduced. With the hub-and-spoke model barge congestion at deepsea terminals can be reduced, due to the bundling of container flows and the use of a consolidation hub (Konings et al., 2013). In the current situation small-size barges need to plan many direct-calls (6 to 8) at different deepsea terminals, leading to a lower overall terminal productivity (Roux et al., 2023). When these small flows are combined in the hinterland terminals, the call sizes at deepsea terminals can increase. With the use of shuttle barges sailing between one or two deepsea terminals and a consolidation hub, container flows can be combined, which has a high impact on the efficiency in the port area (Otten et al., 2020). Besides the advantages of this system, such as faster handling at the deepsea terminal, higher reliability, kilometer reduction, lower costs due to dwell time in the port, and increased call sizes of barges, there are also a few disadvantages. The consolidation hub is an extra node in the network, which means an extra handling. This results in extra transhipment costs and possible extra travel time (Otten et al., 2020). However, barge congestion at deepsea terminals is large, which means that the benefits could exceed the costs.

The hub-and-spoke concept has been studied, but not from the position of the operator that has to decide whether to use it. The recent study of Vermaat (2026) assessed the feasibility of a consolidation hub with a qualitative framework and recommended that quantitative data should be added. Nicolet et al. (2023) modelled a floating hub on regional container flows and recommended a vessel-level analysis, so that waiting times at terminals and hubs can be represented in more detail. The research of Hu et al. (2019) reviewed the inter-terminal transport part and concluded that quantitative methods are needed that reflect the financial interests of the different actors. Roux et al. (2023) state the trade-off directly, namely that the benefit of bundling at the hub has to be larger than the cost of the extra transhipment step, but they do not quantify it. The studies mentioned treat the cost of that extra handling step as a modelled cost. In an operating system it is a tariff that the operator pays. A third-party barge operator pays a fee per container for the shuttle leg. In exchange it avoids the sailing to the Maasvlakte and the waiting at the deepsea terminal, where an operator without a fixed window pays no fee for a call. Whether that exchange is favourable depends on how many containers the operator moves per call, on the cost rate of the inland barge and on how much time the hub routing actually saves. None of these is a property of the hub alone. They are properties of the operator and they differ between operators that use the same hub.

This thesis addresses the part of the operator properties and how they differ between operators that use the same hub in order to find the objective of the thesis. The objective is to quantify the operator cost and turnaround time of routing third-party inland barges through a hub-and-spoke system relative to the direct deepsea calls, and to establish whether and for which inland barge profiles this reduces cost and turnaround times, given the inland-leg profile and the hub system's handling capacity. With the use of AIS data and discrete-event simulation, the objective of this thesis can be realised. Two routings are realised in this thesis. The hub route uses a shuttle barge that sails between the deepsea terminals and the hub, Barge Center Waalhaven (BCW), using a fixed window. The inland barges sail between the hub and the hinterland. In the direct route, inland barges sail directly to the deepsea terminal without a fixed window instead of using the hub. The two routings are compared at the level of a single fleet cycle and with the tariffs that are charged in practice. The comparison is made for three inland barge classes and two different inland legs, so the outcome is reported per operator profile rather than as one

figure for the system as a whole.

The remainder of this report is structured as follows. Chapter 2 reviews the literature on barge congestion, consolidation strategies and hub-and-spoke networks, and closes with the research gap. Chapter 3 describes the case and states the problem with the research questions. Chapter 4 sets out the methodology, consisting of the system description, the conceptual model, the simulation model and the cost layers. Chapter 5 describes the AIS data, the fitted distributions and the case parameters. Chapter 6 reports the verification and validation of the model. Chapter 7 presents the experimental framework and the numerical results. Chapter 8 interprets the results and states the limitations. Finally, chapter 9 answers the research questions and gives the recommendations.

2

Literature review

2.1. Barge congestion and its causes

Barge handling at deepsea terminals carries several structural disadvantages (Van Koningsveld et al., 2021). The first is the small size of the barges. The cranes at the deepsea terminals are dimensioned for ocean-going vessels, so their productivity drops when they handle a much smaller barge. The second is the priority of deepsea vessels, which regularly interrupts the handling of a barge that holds no fixed window (Van Koningsveld et al., 2021). The third is the number of terminals a barge must visit to collect a full load, which becomes even more time-consuming when the barge has to yield to deepsea vessels at each of them. The research of Shobayo and van Hassel (2019) used an agent-based model to find the optimal scenario with respect to the handling of containers in large seaports with connection to the hinterland and barge congestion at deepsea terminals. The conclusion of this research and the optimal scenario to prevent barge congestion at deepsea terminals, was that there simply needs to be more barge infrastructure. Therefore, the container terminals need to expand their infrastructure which brings investments, as an extra quay wall and new cranes, to increase the handling capacity. However, the question remains what the costs would be for the terminal operators and if the benefits can outweigh the costs.

2.2. Consolidation strategies

Another strategy to reduce barge congestion is investigated by Bedoya-Maya et al. (2025). The research focuses on minimizing the total number of containers and trips through the consolidation of cargo (CC) in a container freight station (CFS). The study aims to assess whether consolidating “Less than Container Load” (LCL) shipments into “Full Container Load” (FCL) shipments at specific hubs is economically viable. The goal is to maximize loading factors, thereby reducing the total number of containers and trips required. The study shows that the average costs and emission are reduced per container. However, the study also found that the implementation of consolidation directly at seaports or at individual inland terminals resulted in negative cost savings (losses) in all scenarios due to high handling and logistics costs. Other aspects that make the CC strategy less appealing are for example the critical water level episodes, due to climate change. During these episodes, the capacity of barges is reduced, automatically making cargo consolidation less attractive (Shobayo et al., 2026). Recommendations for future research are therefore to assess exogenous disruptions in the model such as the water level fluctuations.

The responsibility of decreasing the barge in-port times is partly in the hands of the barge operators. Deepsea terminals have mentioned several times that the small call size of the barges is one of the reasons the barge performance is quite low. The berthing of a barge costs approximately 5 minutes. This is the time in which a crane is not moving containers (Van Es, 2019). With a larger call size, the ratio of call sizes and barge berthing will increase. In the research of Roux et al. (2023) a system with higher call sizes for barges is introduced. Due to bundling of container flows between seaports and hinterland ports, the system could mitigate the barge congestion. Bundling container flows could mean

a reduction in the number of visits of barge operators at deepsea terminals and an increase in the call size. The port of Rotterdam has applied different versions of bundling, all with positive feedback. Although these actions have tackled a certain amount of inefficiency, there is still work to do to further mitigate inefficiency and congestion.

2.3. Bundling and network types

The research of Hu et al. (2019) reviews the studies of inter-terminal transport (ITT) in port areas and the hinterland seeking to identify research gaps. ITT is linked to the consolidation idea and allows customers to visit decentralized transhipment points where they can leave the further haulage to the destination terminal to a service provider in the port ecosystem. This prevents calls for inland vessels at deepsea terminals. An example is the Container Exchange Route (CER) in the Port of Rotterdam. This route bundles containers for rail and barge on a dedicated road of 17 kilometers between Maasvlakte terminals, depots, distribution centres and the State Inspection Terminal on the Maasvlakte (Port of Rotterdam, n.d.). CER is mostly used to exchange containers, that would otherwise form small call sizes (Van Es, 2019). However, not many of these services have been realised yet, due to the uncertainty about the cost and benefits of each stakeholder involved (Friedhoff et al., 2023).

In the research of Hu et al. (2019) different ITT networks are mentioned. ITT network could be a line network, a hub-and-spoke network, a trunk line with collection and distribution network or a mix of these networks. The hub-and-spoke network is the most frequently analysed. Caris et al. (2011) used a discrete-event simulation to analyse four hub-and-spoke scenarios in Antwerp, simulated against the current situation to examine the waiting times and utilisation at the hubs and the sea terminals. Another literature research confirms that hub-and-spoke networks can improve the cost performance of container barge transport, specifically in the Rotterdam hinterland (Hu et al., 2019). For further research the study of Hu et al. (2019) mentions quantitative methodologies are needed in the methodology perspective, in order to reflect the financial interests of different actors.

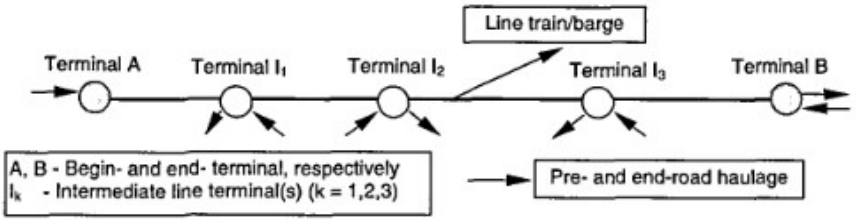

Figure 2.1: Line bundling network. Source: Janic et al. (1999)

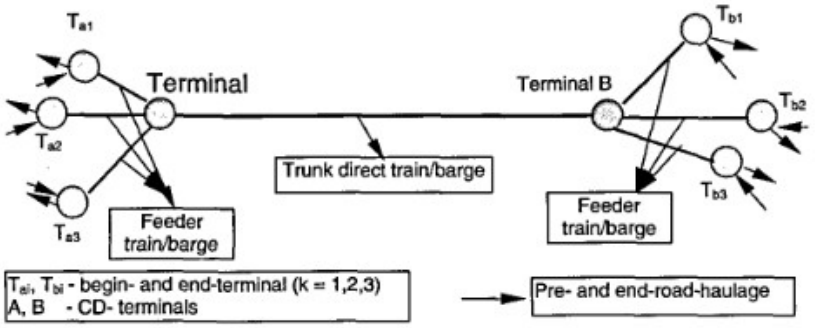

Figure 2.2: Trunk line with collection and distribution network. Source: Janic et al. (1999)

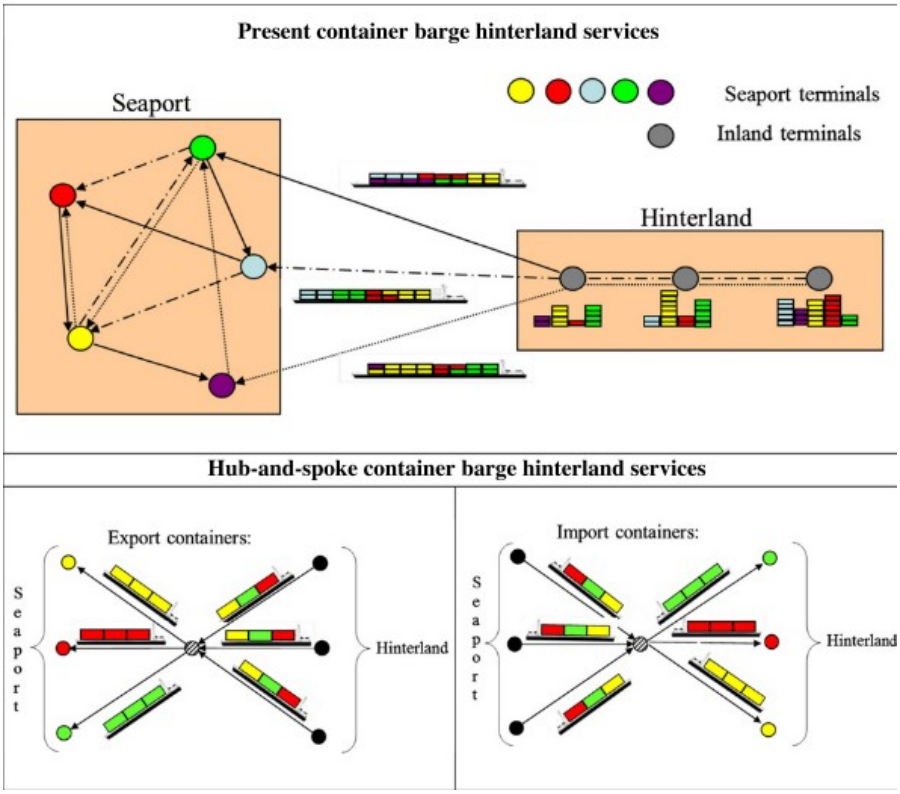

Figure 2.3: Hub-and-spoke design. Source: Konings et al. (2013)

2.4. Hub-and-spoke

The research of Konings et al. (2013) developed a hub-and-spoke model to benefit from the economies of scale of container transport and to improve container barge transport in the hinterland of Rotterdam. A system that uses the bundling of container flows between seaports and hinterland terminals. The hub-and-spoke system transforms the current barge transport structure to significantly reduce terminal visits through a centralised sorting approach. Containers from the seaport terminal with different hinterland destinations are transported together to a consolidation hub and then regrouped and bundled with containers from other seaport terminals. The use of a consolidation hub forms an additional stop and handling steps in the hinterland, resulting in additional costs (Vermaat, 2026). However, by bundling containers, economies of density and economies of scale can be achieved (Konings, 2006). The decoupling of barges from the deepsea terminals should reduce the in-port stay time and improve the reliability of container transportation (Konings et al., 2013). Especially for third-party barge operators without fixed windows the decoupling has a positive effect on the in-port time, because they only need to sail to the hub instead of the deepsea port, which prevents them from calling at the deepsea terminals and joining the waiting queue at the terminals. This could mean more round trips for the barge operators in the same period of time, which would spread the fixed transportation costs over more transport services and therefore decreases the cost per load unit (Konings, 2003). Nevertheless, the economic advantage of bundling at the hub must be large enough to compensate its main disadvantage, which is the cost of the additional transhipment operation in the hub (Roux et al., 2023).

The research of Nicolet et al. (2023) proposed a floating consolidation hub known as the Modular Mobile Terminal (MMT) to address the congestion and long waiting times for container barges in the ports of Rotterdam and Antwerp. This MMT is a variation on the hub terminal in the hub-and-spoke system. Instead of visiting multiple deepsea terminals, inland barges drop and collect their containers at a floating MMT. Dedicated shuttle barges then transport the consolidated cargo to and from specific deepsea terminals. The results of the research reveal that MMTs are most suitable for barges transporting small cargo volumes below 60 TEU. The benefits of the MMT are spatially dependent. The longer-distance hinterland connections capture the greatest cost savings compared to the shorter routes. The capital

cost of the terminal modules is identified as the most critical factor for profitability (Nicolet et al., 2023). Therefore, an examination of which party is going to invest in this infrastructure will be interesting to consider in further research. The practical operations of the shuttle barges to deepsea terminals would depend on this matter because of the impact on the relationship between the MMT and the deepsea terminal. Furthermore, this research relied on the regional container flows, but recommends to use a vessel-level analysis to model the terminals, shuttles and other vessels in more detail. With a vessel-level analysis, waiting times of vessels at the terminals and hubs can be modelled in more detail. Other recommendations for further research mentioned in Nicolet et al. (2023) regarding the demand are the use of an uneven split of containers between deepsea terminals and different types of barges used in the hinterland.

The hub-and-spoke model has requirements that need to be implemented in order to make it generally feasible. One of the most important requirements is fixed windows to offer congestion avoidance and an increased reliability (Vermaat, 2026). Fixed windows are needed for the implementation of the barge service between the seaport and the hub. This service can consist of shuttle barges sailing between the hub and one deepsea terminal or a line service that sails between a few deepsea terminals and the hub. Shuttle barge service works perfectly with larger call sizes and line barge systems are preferred with smaller call sizes, so calls can be made at several deepsea terminals. Furthermore, the use of push boat push barge formation is strongly suggested when implementing a hub-and-spoke network (Konings et al., 2013). Push barges can be used as floating stacks, so that direct transhipment can take place in the hub. This causes a decrease in the handling cost of exchanging containers which is a critical factor in the feasibility of the hub-and-spoke network (Konings et al., 2013). In addition, with the use of push barges, a push barge convoy can be suggested. A push barge convoy can consist of a push barge and a container barge, where the push barge is being pushed by the container barge. The push barge can be released at an inland terminal to be loaded or unloaded, after which the container barge can continue its journey. This concept is already used between Rotterdam and Moerdijk and caused a reduction in the cost (Otten et al., 2020). Another idea is also mentioned in the report of Otten et al. (2020) and is a combination of bundling containers and decoupling the loading and unloading process at the deepsea terminal, so the utilisation rate of the deepsea terminal can be optimised. This system is called the ‘lego-barge system’. The system is based on the use of unmanned push barges laying in the port, that can be called for loading/unloading whenever needed. This could be an interesting idea for the interpretation of shuttle barges.

For a hub-and-spoke network, the location of the hub is an important factor considering the cost performance of the network. One of the reasons is the influence that location has on the turnaround times of the vessels at the seaport link as well as the hinterland link. If the operational time of a vessel is close to the sailing schedule, the utilisation rate of the vessel will be higher. Therefore, a hub location that enables a perfect turnaround time for vessels creates a better performance of the hub-and-spoke network (Konings et al., 2013). Another argument about the hub location is also mentioned in the research of Konings et al. (2013) and states that the distance to the seaport matters a lot considering the flow of containers. The greater the distance, the more a hub within the hub-and-spoke network becomes a corridor-specific hub. When the distance to the seaport is smaller, more hinterland directions will fall in the catchment area of the hub. This sounds like a positive thing, but if the process capacity of the hub cannot deliver to the demand of the inland terminals within the catchment area, the big catchment area of the hub becomes less interesting. In this thesis the hub location is taken as given. The distance to the seaport is small which means the hub has a large catchment area.

2.5. Cost performance of the hub-and-spoke network

As mentioned earlier, the handling costs and location of the hub are important factors in the feasibility of the hub-and-spoke network (Konings et al., 2013). The cost model presented in the research of Konings et al. (2013) compares the present hinterland services against a set of hub-and-spoke settings and reports the outcome as a margin for handling in the hub. The margin is the largest amount the exchange may cost before the network stops being cheaper than direct calls. What the study delivers is therefore a ceiling on the price of the extra node rather than a saving that bundling produces.

How large that ceiling is, decides whether the concept works. One of the hub locations examined is the Eem-Waalhaven area at 40 km from the Maasvlakte, which is the relation that the shuttle service studied

in this thesis operates on today. At that location the margin lies between roughly 1 and 19 EUR per TEU, and it narrows as the hinterland vessel grows larger. The sailing savings that produce the margin lie between roughly 19 and 28 EUR per TEU. Moving the hub further from the seaport widens it, because a longer seaport leg leaves more room for economies of scale, and the widest margin found anywhere in the study reaches 36 EUR per TEU. Section 2.4 introduced direct board to board transhipment with push barges acting as floating stacks for exactly this reason. It is the handling process that has to keep the exchange under the ceiling.

The ceiling rests on quantities that the model supplies rather than produces. The present situation in the research of Konings et al. (2013) is represented by nine terminal calls per port visit with an average of one hour of waiting at each. The hub situation includes a single call with one hour of waiting. Waiting times at the seaport terminals are considered to be small because dedicated barge quays with agreed service levels are assumed. The saving therefore scales with the number of calls removed. The waiting time that the hub avoids and the waiting time it introduces are both parameters. This is a reasonable construction for a concept that did not yet exist when the model was built, but it means the ceiling has never been confronted with waiting times generated by vessels competing for a real berth.

Later work has not closed the condition. The assessment of the floating variant discussed in Section 2.4 works from regional container flows and recommends a vessel-level analysis, so that waiting at terminals and hubs can be represented in more detail (Nicolet et al., 2023). Simulation has been used to examine waiting time and capacity utilisation at candidate hubs, without addressing the cost effectiveness that introducing a hub produces for the barge services (Caris et al., 2011; Konings et al., 2013). A separate way treats the network as something still to be designed and asks where the hub should go and which vessels should serve it (Ypsilantis & Zuidwijk, 2019; Zhang et al., 2020; Zhou et al., 2023). Times and demand are deterministic there, and cost is the technical cost of running the system rather than a price anyone is charged (Zheng & Yang, 2016). That suits a configuration which is still open. It is also the reason the hub location is a decision variable in that literature and a given here.

2.6. Literature synthesis

In the literature review, key topics of barge congestion at deepsea terminals have been discussed. The findings of the studies reviewed are combined and interpreted to find new insights and understanding of the research topic. These findings are explained in this section.

In the hub-and-spoke model the inland terminal functions as a consolidation hub from which third-party barge operators can collect and deliver containers. This prevents barges from sailing directly to the deepsea terminal and having a long in-port stay time due to the barge congestion. Between the deepsea terminals and the consolidation hub, shuttle barges transport bundled containers. This system works and is implemented in the Port of Rotterdam as well as the Port of Antwerp (Roux et al., 2023).

The quantitative research found in the literature is mostly deterministic optimisation on aggregate flows and does not produce waiting time as a resulting variable of the model. The research of Ypsilantis and Zuidwijk (2019) constructed a model to optimise the network of dry ports within the hinterland and showed how the network designs adapt to congestion. Caris et al. (2011) used simulation, but to compare hub configurations on waiting time and utilisation and not whether an operator should pay for the use of the hub. The research of Nicolet et al. (2023) mentions an important note about the financing of the system. The shuttle system and the hub require investments and ownership to be able for a hub-and-spoke system to work. It is therefore important that these investors are known. In the same research a recommendation is mentioned about the use of vessel-level data, since this creates detailed insights in vessel queueing, number of containers per call and the turnaround time of inland barges. Furthermore, the reviewed studies also evaluate consolidation at a given call size or as an operators choice, while the call size is itself a variable that both bundling and water levels move. Hub-and-spoke barge networks have so far been evaluated as designs, comparing hypothetical configurations at network level using modelled handling costs, which leaves open whether a third-party barge operator has a cost reason to use a hub that is already used as an extension of the deepsea terminal and charging a real tariff.

This research takes the network design as given and evaluates the routing choice facing a fleet run by a third-party barge operator. The research follows the recommendation to combine simulation with a case study for container-on-barge systems (Bu & Nachtmann, 2023). Discrete-event simulation of both routings, the shuttle service between the Maasvlakte and the Barge Center Waalhaven and the direct call at the deepsea terminal, produces the waiting times. AIS data of individual inland barges sets the parameters and validates the direct-call simulation, which delivers the vessel-level analysis recommended by Nicolet et al. (2023).

3

Problem description

3.1. Current state

Consolidation hubs are potential solutions to the barge congestion at deepsea terminals. In this research the function of the Barge Center Waalhaven (BCW) as consolidation hub is being reviewed and alternative hub settings are being considered. With a shuttle barge system between Maasvlakte and BCW (implemented by MSC), BCW is used as a dedicated hub of MSC. In the current situation clients of MSC have the option of using the BCW terminal to drop and collect containers with their trucks, instead of trucking all the way to Maasvlakte. The advantage of this system is the savings in travel time because of the distance and the truck congestion at the deepsea terminal.

However, barge congestion at the deepsea terminals remains a frequently occurring problem. The barge congestion is caused by a combination of factors. In the current situation, there are two types of barges; barges with fixed windows and barges without fixed windows. Barges can collect containers at one terminal or visit different terminals per cycle. The barge congestion does not have a large negative effect on fixed window barges compared to barges without fixed windows. Deepsea vessels are prioritised over barges without fixed windows due to the absence of a contractual relationship between the deepsea terminal and the barge operator (van der Horst, 2016).
The result is a waiting line of barges that is being created at the deepsea terminal. The waiting line increases even more when the barge operator needs to make many calls at different deepsea terminals without owning fixed windows. This causes a risk of missing the planned calls at the different terminals and is one of the reasons why delays can lead up to 52 hours on average for barge operators (Vermaat, 2026). In Figures 3.1 and 3.2 the current state is visualised in a simplified way to explain the routes of the barges.

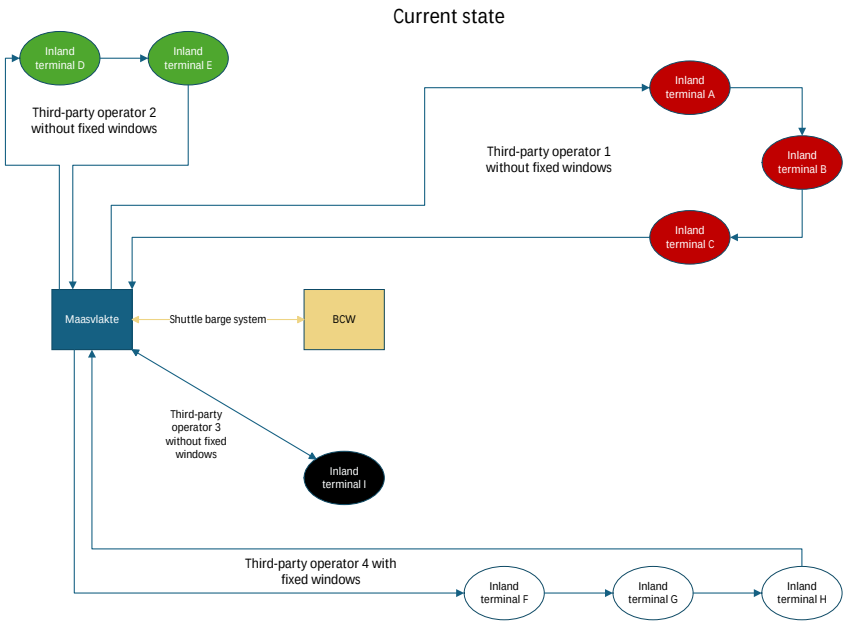

Figure 3.1: Visualisation of the current state of the barge network

Figure 3.1 shows a barge network in the current state with the routes of four third-party barge operators. These trips are made with barges between the deepsea terminals of the Maasvlakte and different inland terminals. The double sided arrow between the Maasvlakte and the BCW shows the route of the shuttle barge already implemented by MSC. Containers that arrive at the BCW can be collected by the client directly, which is mostly done by truck. The route to the client with the truck is not included in the figure, because this figure shows only the barge network and the truck routes are not included in this research.

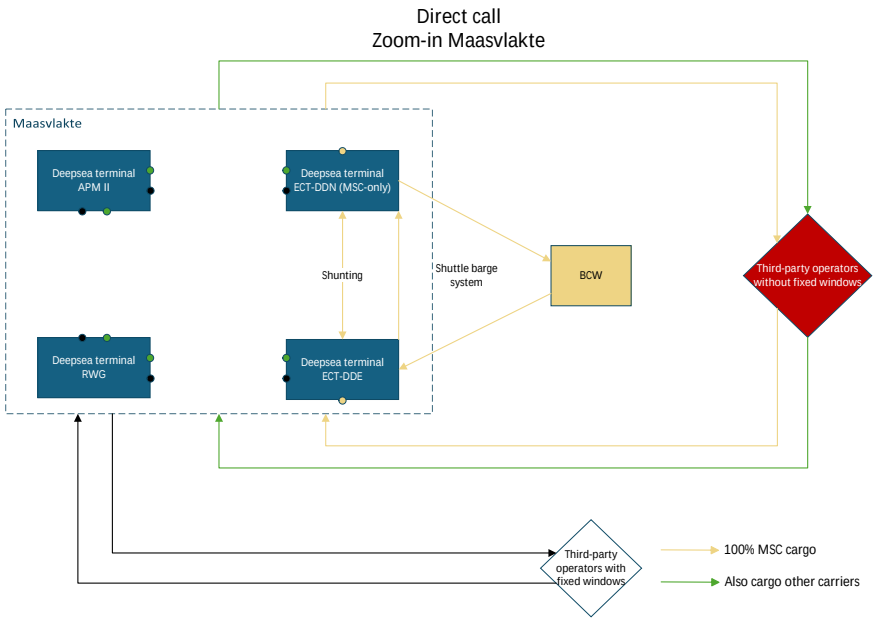

Figure 3.2: Visualisation of the current state zoom-in on the Maasvlakte

Figure 3.2 zooms in on the Maasvlakte area to show the different deepsea terminals and the kind of cargo transported between the deepsea terminals and the hinterland. The hinterland is divided into two diamonds consisting of inland terminals used by third-party barge operators with fixed windows (white) and inland terminals used by third-party barge operators without fixed windows (red). The arrows show the routes of the barges and the kind of cargo transported, and the dots connected to the different deepsea terminals present the available deepsea terminals where the different kinds of cargo can be collected by the barge. The yellow arrows present the barge routes with 100% MSC cargo. The ECT-DDN is a terminal where most of the 100% MSC cargo is handled. When this terminal is occupied, barges that collect 100% MSC cargo can deviate to the ECT-DDE to collect the MSC-cargo there. However, this is only done when absolutely necessary since the MSC cargo needs to be shunted from the ECT-DDN to ECT-DDE, which costs 70 EUR per container.

3.2. Problem statement

Due to the growth of the global container flow and the expanding capacity of the deepsea vessels, deepsea terminals have a very high throughput. In order to get these containers through the hinterland the use of Inland Waterway Transport (IWT) has become an important and efficient mode of transport. Especially because road transport is often congested and trucks have a low capacity when it comes to transporting containers. Unfortunately, barges have difficulty making calls at a deepsea terminal. This problem occurs because there is not always a fixed window within the contract between the deepsea terminal and the barge operators. One of the reasons is the high cost of fixed windows for the barge operator. The consequence is an increase in barge congestion at the deepsea terminals because of the prioritisation of deepsea vessels over the barges. Barge operators do not get a time window unless there is space for them (Oganesian et al., 2021). Another problem of the increasing congestion is the high number of barges with small call sizes at different deepsea terminals. The many direct-calls (6 to 8) at different deepsea terminals lead to a lower overall terminal productivity and therefore a longer waiting line of barges (Roux et al., 2023). Due to barge congestion, the flow of containers to the hinterland is stagnating.

Research has been done to find solutions to barge congestion (Konings et al., 2013; Nicolet et al., 2023; Roux et al., 2023; Shobayo & van Hassel, 2019). The difficult part remains the number of different stakeholders within the port that are related to the barge congestion. Certain research has been looking at the expansion of deepsea terminals through a consolidation hub that makes use of the

hub-and-spoke model (Konings et al., 2013). This consolidation hub refers to a central place where different barges go to one central place to deliver and collect containers. Between the consolidation hub and the deepsea terminal one large shuttle barge is commuting to keep the container flow running both ways. With this idea, fewer barges need to travel all the way to the deepsea terminal, which will have a decrease in the barge congestion as a result (Konings et al., 2013).

There are several conditions to this idea. The most important condition is a fixed window at the deepsea terminal for the large shuttle barge to deliver and collect containers in a continuous flow. Without this fixed window the consolidation hub is not as effective as it could be. MSC is using a certain consolidation hub at the moment, where one large barge is commuting between the deepsea terminal and the Barge Center Waalhaven (BCW). The BCW functions as the consolidation hub in this case where trucks can collect containers. For now the BCW functions as a hub for clients only, but in the future it may also be possible to welcome third-party barge operators who can load and unload their cargo to prevent the barge congestion at the deepsea terminal. Figure 3.3 shows the future scenario of the hub-and-spoke system with the BCW as hub. The barges load and unload their containers at the hub instead of sailing all the way to the deepsea terminals.

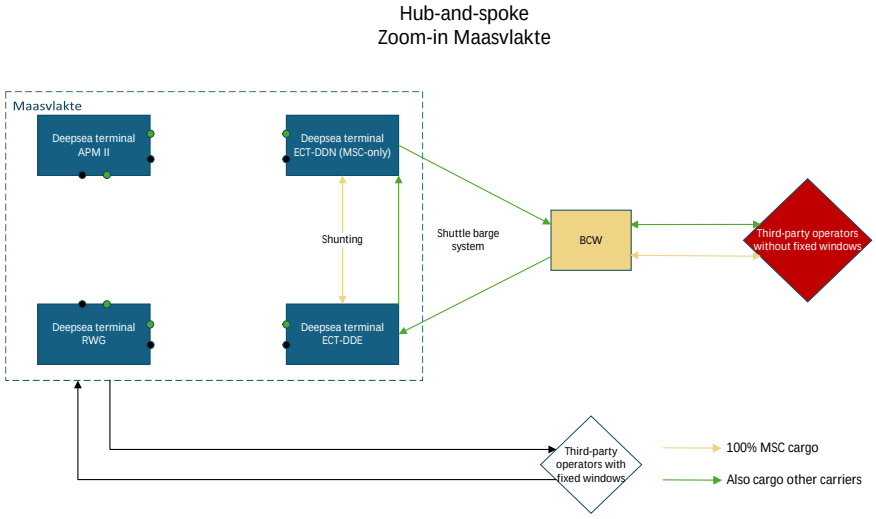

Figure 3.3: Visualisation of the future hub-and-spoke system with BCW as the hub

The shuttle system between the Maasvlakte and BCW is already in operation, so the question is no longer whether such a system can be built. The question is whether a third-party barge operator has a reason to use it. Routing through the hub removes the call at the deepsea terminal and the waiting time that comes with it. However, it also adds a handling step at the hub and a fee per container for the shuttle leg. Which of the two is larger is not known and there is no reason to expect the same answer for every operator. A small barge on a short hinterland leg and a large barge on a long leg differ in the number of containers they move per call, in their cost per hour, and in how often they return to the hub. The problem addressed in this thesis is therefore that the cost and the time effect of routing a third-party inland barge through the hub is not known and neither is the way that effect depends on the operator profile. As long as this is unknown, an operator cannot decide whether to use the hub and the hub cannot be priced or dimensioned for the operators it is meant to attract.

3.3. Research questions

The objective of this research is to quantify the operator cost and turnaround time of routing third-party inland barges through a hub-and-spoke system relative to direct deepsea calls, and to establish whether and for which inland barge profiles this reduces cost and turnaround time, given the inland-leg profile and the hub system's handling capacity.

To what extent does routing third-party inland barges through a hub-and-spoke system change their turnaround time and transport cost per TEU and for which operator is the hub routing a promising alternative?

The following sub-questions help answering the main research question:

1. What does the observed AIS data reveal about current inland-barge turnaround times and Maasvlakte in-port times in the Port of Rotterdam?
2. How do the cost and time effects of the hub routing differ across inland vessel classes?
3. How does the inland-leg profile (short vs. long) affect the cost and time benefit of routing via the hub?
4. Which constraint limits the throughput of the hub-and-spoke system and how does it shape the call size of the barges it serves?
5. What are the cost implications for the third-party barge operator and under which conditions would the hub situation become the cheaper routing compared to the direct route?

4

Methodology

This thesis adopts a quantitative case-study approach in which discrete-event simulation (DES) is the primary analytical method. The case-study approach is applied to the real hinterland network used by MSC in the Port of Rotterdam. In this network, a shuttle barge connects the deepsea terminals ECT-DDN and ECT-DDE on the Maasvlakte to the Barge Center Waalhaven (BCW), which acts as a hub from which third-party inland barges collect containers to transport into the hinterland.

Discrete-event simulation is chosen because of the stochastic interactions that drive the cause of the research problem. These stochastic interactions can be translated to randomness in container arrivals, variable sailing and handling times, barge fleets, fixed opening hours at the hub, and the priority of deepsea vessels over inland barges that do not have a fixed window at the deepsea terminal (Oganesian et al., 2021). These dynamics generate queuing and delays that a static optimisation model cannot represent. Bu and Nachtmann (2023) explicitly recommend simulation combined with case-study data for container-on-barge systems, precisely because mathematical optimisation alone is insufficient when stakeholder interactions and operational variability dominate system performance. This research follows that recommendation. A second feature of the approach defines how the comparison is constructed. By simulating both a current-state network and a hub-and-spoke network and comparing the two model outputs a baseline comparison can be made. From there hub configurations can be tested compared to the same current-state simulation.

The remainder of this chapter is organised as follows. Section 4.1 explains the research approach and the reasons for the methods used. Section 4.2 describes the system with flowcharts of the hub-and-spoke and current situation. Section 4.3 presents the conceptual model and the simplifying assumptions. Section 4.4 specifies the simulation model consisting of the components and processes. Section 4.5 describes the filtering of the AIS data resulting in a distribution used in the simulation. Section 4.6 describes the cost layer that translates time-based output into operator cost per TEU.

4.1. Research approach

This research is built around a single case study, consisting of the container import flows of MSC and the operational role of Barge Center Waalhaven (BCW) in the Port of Rotterdam. The case study design is appropriate because MSC operates a real shuttle service between the Maasvlakte and BCW and is able to provide the operational data needed for the realistic part of the model by actual movements of inland barges. Using a single operator under fixed external conditions also isolates the effect of interest. The research question concerns the relative difference between routing third-party barges directly to the deepsea terminals and routing them via the hub. Holding the operator, terminals and corridor constant removes unnecessary variation and lets the model attribute the observed differences to the hub configuration and the operator profile.

Within this case study, discrete-event simulation is the core method. It is supported by two further methodological elements: an explicit cost layer (section 4.6) and a scenario and sensitivity analysis (chapter 7). The outcome of these analyses makes it possible to answer each sub-question with quan-

titative output. The scenario comparison gives the time effect of the hub, the fleet calibration gives the appropriate fleet size per vessel class, and the cost layer applied across the scenarios presents the cost effect. The robustness of the conclusions is given by two sensitivity checks focused on the deepsea block in the current-state simulation and the shuttle fee floor in the hub-and-spoke simulation.

4.2. System description

The system is described as time-ordered flowcharts that visually represent the route a barge follows in the current situation (Figure 4.1) and in the hub-and-spoke situation. The flowcharts make explicit where waiting and handling occur in each case. Both figures show the import container flow only, since the export flow of containers is not included in this research. The route of the barges is not affected by the exclusion of the export flow.

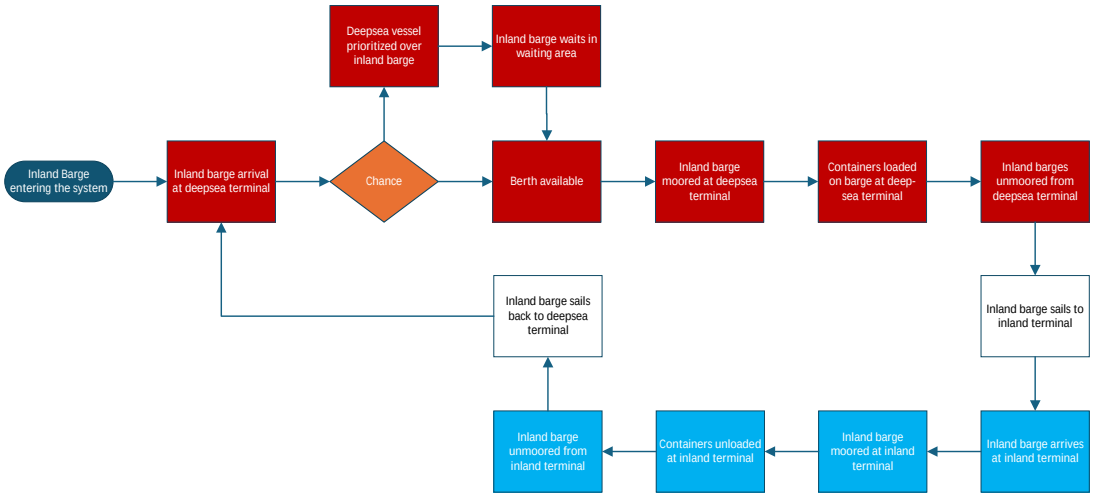

Figure 4.1: Flowchart of the current barge process

The flowchart in Figure 4.1 shows the process of the current situation. The flowchart shows the events in time of the inland barges and the terminals in order to transport containers from the deepsea terminal to the hinterland. The path shape in the flowchart presents the possibilities of situations that can occur at the deepsea terminal. The situation in which an inland barge needs to wait for a deepsea vessel that is prioritised over the inland barge, is the situation from which the inland barge congestion appears mostly (Nicolet et al., 2023).

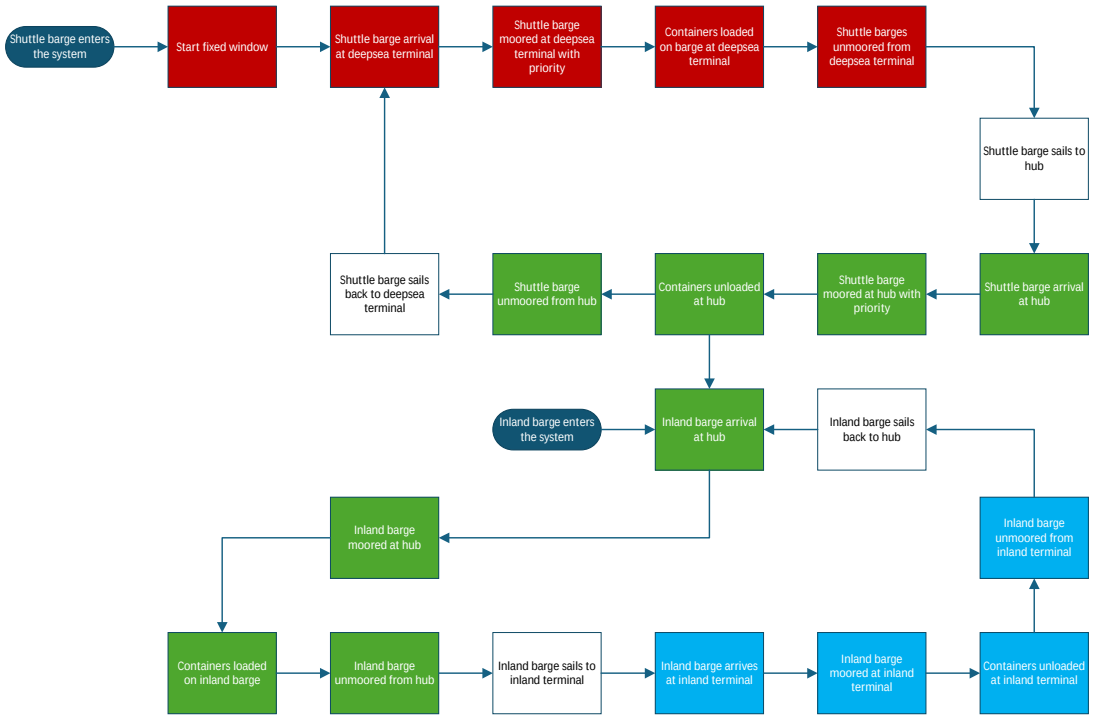

Figure 4.2: Flowchart of the hub-and-spoke process

Figure 4.2 shows the flowchart of the hub-and-spoke situation. The biggest differences between the current-state and the hub-and-spoke situation are the shuttle barges and the extra node (hub) that is implemented between the deepsea terminal and the inland terminal. The hub-and-spoke system consists of a shuttle barge loop and an inland barge loop. The two loops alternate in the transportation of containers. This alternation happens at the hub when the containers are unloaded at the hub from the shuttle barge. The hub-and-spoke flowchart shows a larger process with more steps, which would suggest a longer transportation time between the deepsea terminal and the inland terminal. The aim of the shuttle system with fixed windows is to avoid the barge congestion at the deepsea terminal. Whether this saves time and cost for the operator is tested in chapter 7.

4.3. Conceptual model and simplifying assumptions

The conceptual model abstracts the flowcharts into a tractable simulation, which is kept generic so that it can represent the hub-and-spoke configuration consistently across all scenarios and against the same current-state simulation. A visualisation is given in Figure 4.3. The red arrows show the shuttle barge route between the deepsea terminals and the hub. The black arrows show the inland barge route between the hub and the hinterland. The current-state simulation excludes the hub, the shuttle system and uses only one deepsea terminal. The inland barges make a direct call at the deepsea terminal, collect their containers and deliver them to the inland terminals.

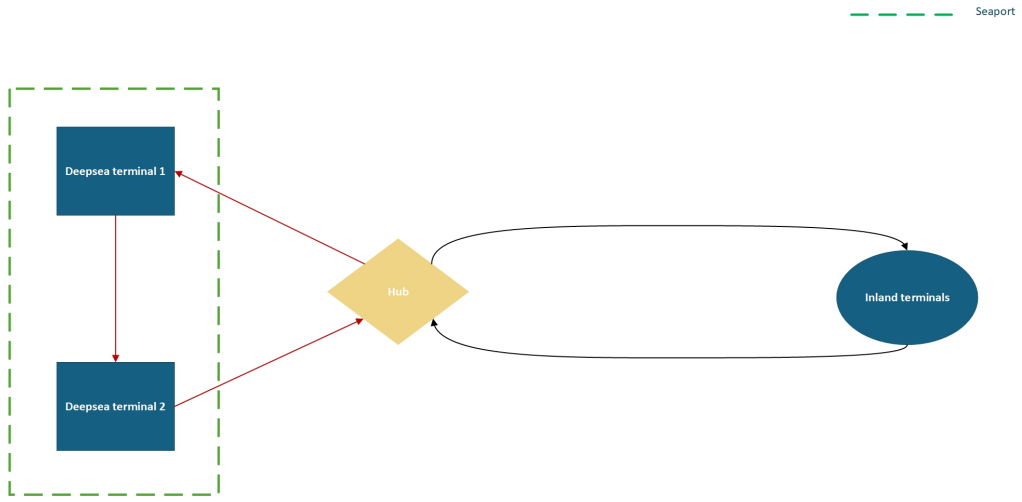

Figure 4.3: Conceptual model of the hub-and-spoke network

In the process of creating the conceptual model and flowcharts, simplifications and assumptions have been made. These assumptions and simplifications have been added to outline the system and reduce complexity in order to keep the focus on the main goal of this research. The assumptions are located in Appendix B. Table B.1 lists the assumptions made in constructing the two simulation models. Each assumption is given with its justification and the expected direction of its effect on the gap between the hub routing and the direct call. The assumptions define the scope of the model and concern the deepsea terminals, the shuttle and hub, the inland leg and the cost layer. They are explained further in Sections 4.4 to 4.6.

4.4. Simulation model

The conceptual model and the process of the flowcharts are implemented in Python. For the two simulation models, discrete-event simulation is used with SimPy as Python library. The hub simulation represents three classes of resource (deepsea terminals, the hub and the inland terminals) and three classes of moving entities (containers, the scheduled shuttle barge and the inland barges). The direct call simulation is a simpler version of the hub simulation where the shuttle system and hub terminal are excluded. The models advance in continuous simulated time and record events only after the warm-up period (Section 7.2). One simulated time unit is one hour in the system.

4.4.1. Container arrivals and deepsea terminals

Containers arrive at two deepsea terminals in batches. The batches differ in numbers of containers per deepsea terminal. The difference in volume per deepsea terminal is partly based on empirical data and suggested by Nicolet et al. (2023). Each deepsea terminal consists of a single berth and a single crane resource that handles moves per hour with a normal distribution. This is consistent with the crane productivity reported for ECT (Hutchison Ports ECT Rotterdam, 2026). Containers stay at the terminal until the shuttle barge loads them within its fixed window in the hub simulation. In the current-state simulation the containers get collected by the inland barges directly when a berth is available.

4.4.2. Shuttle service and fixed windows

A single dedicated shuttle barge serves the deepsea terminals on a fixed timetable. The shuttle operates in a number of cycles per week, aiming to arrive at the hub during its opening hours. The sailing time between the deepsea terminal and the hub is on average the same in each direction. In the model, each trip consists of a fixed minimum and a random delay drawn from a gamma distribution. The minimum represents the fastest realistic trip. The delay falls to the right, because a trip can take longer than average but cannot be much faster. At each deepsea terminal the shuttle loads within a fixed window. If the window closes, the shuttle departs immediately, even if it is not full. This is the defining feature of the hub-and-spoke service and the mechanism by which it avoids joining the barge queue without

fixed windows (Vermaat, 2026). At the hub, the shuttle has priority over inland barges for the berth, so that it can keep to its schedule. Just as the inland barges at the hub, it can only be unloaded during the hub's opening hours.

4.4.3. Hub operations and opening hours

The hub (BCW) has a single berth served by two cranes handling moves per hour with the same normal distribution as used with the deepsea cranes. The hub has a maximum capacity for containers and keeps an opening hour policy. Loading and unloading can only take place during the open hours. If the barge is moored at the berth and the hub closes during the loading time, the loading of containers will be interrupted until the hub is open again. A barge that moors outside opening hours needs to wait until the hub opens. To represent berth occupation due to traffic that is out of scope for this study, a shuttle barge occupies the hub berth several times per week ahead of the inland barges in priority. This is a shuttle barge that transports containers between the Port of Antwerp and BCW in this case. Only its occupation of the berth is modelled, its containers are not part of the hub volume. To avoid overestimation of the hub capacity, a lower capacity is taken into account instead of the official hub capacity.

4.4.4. Inland barges and the inland leg

Inland barges form the third-party barge operator fleet that collects containers at the hub and carries them to hinterland destinations. The fleet size is a scenario variable and each barge capacity is set by the vessel class chosen for the scenario (90 TEU for an M6, 192 TEU for an M8, and 256 TEU for an M9). A barge arrives at the hub, waits if no containers are yet available or the berth is occupied, requests the hub berth, moors, loads during opening hours, unmoors if the loading is complete or if the hub stock is empty, and sails the round trip to its destination. After the inland barge arrived at its destination it returns to begin a new cycle. In the current-state simulation the inland barges sail straight to the deepsea terminal instead.

The inland round trip (sailing out, dwelling at the destination, sailing back) is not modelled deterministically but sampled from the empirical AIS distribution (Section 4.5.2). Because empirical dwell already includes the real waiting and berth time at the inland terminal, inland-side waiting is treated as exogenous and is billed at the waiting rate, while hub-side waiting emerges endogenously from berth and crane contention. To keep these two mechanisms separate, the inland berth resource is sized so that it is never binding, ensuring that no artificial inland queue is mixed with the empirically measured dwell.

4.4.5. Current-state simulation

The current-state simulation represents the direct call. It uses the same inland barges, the same inland-leg distributions and the same seeds as the hub simulation, so the two routings differ only in the port-side part of the round trip. The current-state simulation has one deepsea terminal with one berth and one crane handling moves per hour with a normal distribution included. The container stock at this terminal never runs dry, so every barge loads to its full capacity.

Compared to the hub route, the inland barge sails further in each direction to reach the Maasvlakte, since the hub is not part of the current-state simulation. The barges do not have a fixed window and deepsea vessels have priority at the berth. This is represented by blocks in which the berth is reserved for deepsea vessels. A block has priority over the inland barges in the queue. It does not interrupt a barge that is already loading, but a waiting barge cannot moor until the block has ended. The same block schedule is used for all vessel classes. Experiment 4 lengthens these blocks with broken crane situations elsewhere included to test extreme congestion at the deepsea terminal. In case of a broken crane situation elsewhere, vessels with fixed windows come from other berths that experience broken crane situations to be handled according to schedule. Because the inland barges do not have a fixed window, they pay no call fee at the deepsea terminal. The fleet in each baseline scenario is the same as in the hub scenario it is compared with, so the comparison is a routing choice made by one operator with one fleet.

4.5. Input modelling

This section explains how the data is collected and cleaned. The filtering steps of the data are explained, as well as the use of the data in the simulation.

4.5.1. From AIS data to visits

The empirical foundation of the model is AIS data covering the inland barges that carry containers. Each barge transmits its position over time, together with a navigational status that distinguishes sailing (status 0) from moored (status 5). With the use of the GeoJSON format, the relevant areas of the deepsea port are drawn as polygons. The relevant areas are the deepsea terminal berths, the waiting areas, and the broader zones in which these berths and areas are located. Each change in a barge's status is then matched to the area in which it occurred. This creates a dataset in which the barge route can be followed as a sequence of labelled visits. The next step is a data analysis in Python. The dataset with the total number of barges and their visited locations in a certain status, is filtered and separated into different datasets. These datasets show the in-port time of barges that only visit the ECT terminals during their visit to the seaport, the number of cycles of each barge, and the inland turnaround time of the barges. The in-port time consists of loading, waiting, in-port sailing and approach sailing. The approach sailing describes the sailing between the deepsea terminal and the hub. Since this sailing time is also replaced by the shuttle system in the hub simulation, it is included in the total in-port time of the current-state simulation. The datasets are used to parameterise the stochastic inputs of the simulation and to validate the current-state simulation.

4.5.2. The inland leg

The inland turnaround time is sampled from the cleaned set of AIS-derived inland trips. Trips with extreme dwell times are removed from the dataset. The total inland sailing time is measured from the deepsea area to the inland terminal and back to the deepsea area. The deepsea hub sailing time is subtracted in both directions so the inland leg only covers the part both routings share. To represent the two turnaround profiles, the cleaned trips are split at the median of the combined inland leg. This results in a short subset and a long subset. Within each subset the sailing and dwell components are sampled independently.

4.5.3. Distribution selection

Due to the limited AIS data, a fitted distribution is used to smooth the limited sample. The distribution for each component is selected in three steps following standard input modelling practice (Law, 2020). Three candidate distributions with positive support are considered, namely the lognormal, the gamma and the Weibull distribution. Each candidate is fitted per component by maximum likelihood with the location parameter fixed at zero, so every fit has two free parameters. The maximum likelihood estimate maximises the log-likelihood of the observed sample.

$$\ln L(\theta) = \sum_{i=1}^{n} \ln f(x_i; \theta), \tag{4.1}$$

where $f(x;\theta)$ is the candidate density and $x_1, \ldots, x_n$ are the observed component times.

In the first step the candidates are ranked with the Akaike information criterion:

$$\mathrm{AIC} = 2k - 2 \ln L(\hat{\theta}), \tag{4.2}$$

With $k = 2$ parameters for every candidate, the ranking reduces to a likelihood comparison while remaining valid if distributions with different parameter counts were added. A difference in AIC above 10 indicates essentially no support for the weaker model, while a difference below 2 indicates comparable support. In the second step, the fitted distributions are compared with the sample moments, since the fitted distributions act as a parametric smoothing of the empirical data. After all, the purpose of the fitted distributions is to reproduce the level and spread of the observed legs rather than every local mode. In the third step, goodness of fit is assessed with the Kolmogorov-Smirnov statistic:

$$D = \max_x \left| F_n(x) - F(x; \hat{\theta}) \right|, \tag{4.3}$$

where D is the largest distance between the empirical distribution function and the fitted one.

The standard p-value of this test assumes parameters that were specified independently of the sample. Because the parameters are estimated from the same sample, the fitted distribution is pulled towards the empirical one, the observed D shrinks and the standard p-value becomes too tolerant (Lilliefors, 1967). The p-values are therefore obtained from a parametric bootstrap. A total of 1000 synthetic samples of the original size are drawn from the fitted distribution, the parameters are re-estimated on every synthetic sample and the p-value is the fraction of synthetic D values that reach the observed one. Two checks are used to support this use. First, the sum of the fitted components is compared with the observed combined leg. Second, the same fitted distributions are used in every scenario under common random numbers and the inland leg enters the hub route and the direct route identically. Therefore, any remaining misfit affects the absolute cost and time levels far more than the comparison between the two routes on which the conclusions rest.

4.6. Cost layer

The cost layer translates the time-based simulation output into operator cost per TEU. Two components make up the cost of one replication. First, the inland barge's sailing hours are charged at a vessel class specific sailing rate and its waiting hours at a lower vessel class specific waiting rate. Sailing hours are computed as turnaround time minus waiting time and loading time, so congestion shows up purely as more waiting hours at the waiting rate. The waiting hours include the waiting at the hub or the deepsea terminal and the dwell at the inland terminal. In the current-state simulation the extra hours of sailing to the Maasvlakte are part of the sailing hours. The loading hours are multiplied by the waiting rate, since the dataset where the rates are from, calculates the loading hours in this way. The sailing and waiting rates per class are extracted from a document of Ministerie van Infrastructuur en Waterstaat (2026), which updates the rates per month. Second, the shuttle transhipment is charged per container moved, since that is the way third-party barge operators need to pay for the shuttle and hub service. The costs are presented in Table 5.2. The per-move charge is the structural cost disadvantage of the hub configuration, since the hub adds an extra handling step (Otten et al., 2020).

In the current-state simulation no costs are assigned to the moves at the deepsea terminal, since the price per move is paid by the carrier MSC and inland barges without a fixed window do not have to pay for a call at the deepsea terminal. Instead the barges need to wait much longer before they can load compared to the shuttle barges in the hub simulation. The free deepsea call and the per-container shuttle compensation at the hub define the central economic trade-off of the hub: the free call with waiting time a barge avoids by not calling at the deepsea terminal versus the variable per-container charge it pays for the shuttle leg. The Value of Time and Value of Reliability used by Bedoya-Maya et al. (2025) are deliberately excluded. They are behavioural valuations rather than real expenditures, and including them alongside the waiting-hour costs would double-count delay. Cost per TEU is computed by forming the ratio of cost to delivered TEU within each replication and only then averaging across replications, so the reported confidence interval reflects run-to-run variation in the ratio itself.

4.7. Conclusion

This chapter has set out the method used to compare the two routings. The system was described first as time-ordered flowcharts, which identified the resources, the control variables and the information flows that the simulation has to represent. The conceptual model reduced that description to a tractable form and mentioned the simplifying assumptions behind it. The simulation models implemented the conceptual model in SimPy. The hub simulation consists of two deepsea terminals, a scheduled shuttle under a fixed window, a hub with a single berth and fixed opening hours, and a fleet of inland barges whose round trip is sampled from the AIS data. The current-state simulation excludes one deepsea terminal, the hub and the shuttle system. The cost layer then translated the time-based output of the model into an operator cost per TEU, using the sailing and waiting rates of the barge class together with the fee that separates the routings.

Every simulated scenario in this thesis is a hub-and-spoke configuration compared to the same direct call situation. The simulation scenarios are paired by seed, which prevents the results from having unnecessary variation between the scenarios and the models. Furthermore, the inland leg is sampled from the same distributions in both routings and enters both identically, so it cannot produce a difference between them. What separates the two is the port-side time and the fee structure, which consist of a free deepsea call against a variable fee per container for the shuttle leg. The methodology is built so that any difference the results report must originate in one of those two places.

The models described here have to be shown to work as intended before their output can be interpreted. Chapter 5 describes the data and the case parameters, chapter 6 reports the verification and validation and chapter 7 presents the results.

5

Data analysis

This chapter applies the method described in the methodology to the case study of MSC in the Port of Rotterdam. It describes the different steps that have been performed to arrive at the data that can be used for the input of the simulation models. The cleaning steps for the AIS data, the fitted distribution of the inland leg and the parameters used in the simulations are explained.

5.1. AIS data

The AIS data covers the inland barges that transport MSC containers to and from the Port of Rotterdam. The dataset contains 25,447 status events of 142 vessels over a period of 40 days. The polygons that are used to mark different areas cover the berths of the deepsea terminals, the waiting areas and broader zones such as the Maasvlakte, Eemhaven and Waalhaven. The inland terminals in the hinterland are marked as well, so that the destination of each trip can be identified.

5.1.1. Deepsea terminal visits

The data is filtered based on the inland leg and the in-port cycles of the inland barges. With Python the inland leg and in-port cycles are separated from each other. The in-port cycles are filtered on visiting only the ECT-terminals. The dataset contains barges that also visit other deepsea terminals, but those barges fall outside of the scope of this thesis. After cleaning the dataset, 140 of the 142 inland barge cycles contain a call at a single deepsea terminal. The shuttle barge system sails between both the ECT-terminals and the BCW. The current-state simulation therefore has one deepsea terminal compared with two deepsea terminals in the hub simulation.

For the inland leg, all cycles in the dataset are used including the cycles that contain visits of other terminals beside the ECT-terminals. This is possible, because only the inland leg of the cycle is important in this case, which does not include the in-port situation. Barges that sail between the Port of Antwerp and the inland terminals are removed from the dataset as well as the barges that do not visit an inland terminal at all within their cycle. The last cleaning step contains the maximum dwell time at the inland terminals. Barges that have a dwell time at the inland terminals of more than a week are filtered out of the dataset. This leaves 282 useful inland leg cycles of 54 barges. The sailing time between the Maasvlakte area and the BCW is removed from the inland leg so the same inland legs can also be used in the hub simulation. 3.5 hours of sailing time is subtracted from the inland leg in both ways. These 7 hours are the extra sailing hours of a direct call compared to the hub routing. The 282 cycles are split at the median combined inland leg of 55.7 hours. This gives a short and long leg of 141 cycles each.

The most important quantity is the time a barge currently spends inside the port at the Maasvlakte, because this is the time that routing via the hub is intended to remove. For each in-port cycle, the in-port window is decomposed into four components that sum exactly to the window: loading hours, waiting hours, in-port sailing hours and approach sailing hours. The approach sailing component is the approximately 3.5-hour leg between the hub and the deepsea terminal, which is counted as part of the in-port time in both the current and hub simulation. Two empirical observations support the modelling

choices. First, virtually all observed in-port cycles involve a call at a single deepsea terminal, which makes one berth call per cycle the empirically defensible baseline. Second, the mean in-port time at the deepsea terminal, without the approach, is 12 to 19 hours per call in the current-state simulation, depending on the class. For the comparison with the hub simulation, the approach sailing time is added to the in-port time, because the shuttle system removes the in-port time and the approach sailing time in the hub simulation.

5.1.2. Inland-leg distribution

Table 5.1 gives the result of the three selection steps of Section 4.5.3 for the four components.

Table 5.1: Candidate distributions for the inland-leg components. AIC per family (lower is better), fitted Weibull parameters and bias-corrected KS p-values.

	AIC			Weibull fit		
Component	Lognormal	Gamma	Weibull	k	λ (h)	KS p
Short profile, sailing	1050	1032	1019	2.29	21.9	<0.01
Short profile, dwell	1063	1034	1023	2.20	21.7	0.41
Long profile, sailing	1305	1287	1287	2.17	56.2	<0.01
Long profile, dwell	1297	1256	1251	2.06	47.3	0.05

In the first step the Weibull distribution has the lowest AIC for three of the four components, with margins of about 30 to 45 points over the lognormal. For the sailing component of the long leg, the Weibull and the gamma distribution are almost the same with a difference of 0.3 points. In the second step the fitted values are compared with the sample values. For the Weibull distribution the mean equals $\lambda\,\Gamma(1+1/k)$. The fitted Weibull reproduces the sample mean and standard deviation of every component within about 2%. For the dwell component of the short profile, for example, the fitted Weibull gives a mean of 19.2 hours and a standard deviation of 9.2 hours, against sample values of 19.3 and 9.2 hours. The lognormal fit of the same component gives 20.1 and 13.9 hours. Across the four components the lognormal overstates the standard deviation by about 30% to 55%. This would give the simulation more variability than the data supports. The Weibull distribution is therefore used for all four components. The bias correction in the third step makes a clear difference. The naive test accepts the lognormal dwell fit of the short profile with $p = 0.11$, while the corrected test rejects it with $p = 0.001$. The Weibull fit of the same component drops from a naive $p = 0.79$ to a corrected $p = 0.41$ and remains accepted. The corrected test accepts the dwell fits at the 5 percent level, although the long profile is marginal. It rejects the sailing fits for every candidate family. This rejection is a direct consequence of the median split. The split is applied to the combined inland leg and truncates the sailing marginals, which produces the multimodal shape in Figure 5.1. No two-parameter family reproduces these local modes.

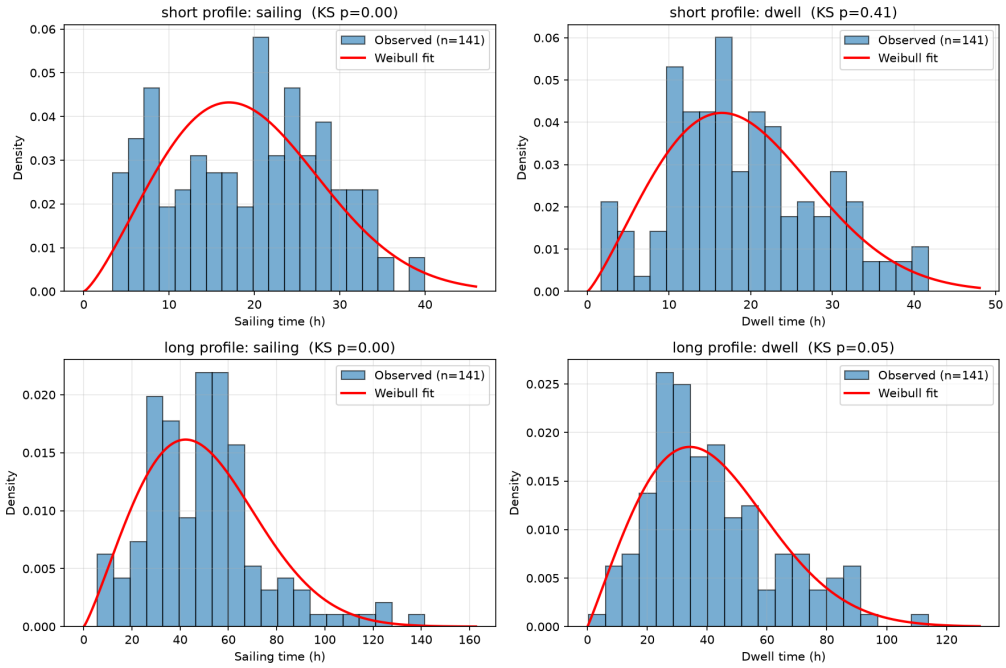

Figure 5.1: Weibull fits for the sailing and dwell components of the short and long inland-leg profiles, with bias-corrected KS p-values.

5.2. Model parameters

Table 5.2 lists the parameter values of the case study. The distribution column shows which inputs are stochastic and the source column shows where each value comes from. The run settings of the experiments, such as the run length, the number of replications and the warm-up period, are not case data and are given in the chapter of the numerical results. Most values were provided by MSC. These are the values of the shuttle service, the hub and the deepsea blocks of the current-state simulation. The crane rate, the minimum call size and the call fee at the deepsea terminals follow from the barge handling conditions of ECT (Hutchison Ports ECT Rotterdam, 2026). The inland leg and the single deepsea terminal of the current-state simulation follow from the AIS data. In 140 of the 142 observed cycles the barge calls at a single terminal, so the current-state simulation has one deepsea terminal.

Two inputs are assumptions. The first is the arrival of containers at the deepsea terminals in batches of 100 and 80 containers at a mean interval of 12 hours. This assumption does not drive the results, because the arrivals are much larger than what the shuttle can carry. The shuttle therefore sets the ceiling of the hub system. This matches the ceiling of 540 to 545 containers per week in Table 7.2. The second assumption is the gamma delay in the sailing time of the shuttle. Only the average sailing time of the shuttle barge was known, so a gamma delay was added.

The official holding capacity of BCW is 3,700 TEU. The model uses 2,700 TEU to avoid overestimating the capacity that is available for the hub flow. This lower value is used, because the hub also handles containers from the Port of Antwerp. Therefore the maximum capacity chosen in this case is not the real maximum capacity of the hub. The sailing and waiting rates per vessel class come from the cost figures for freight transport of Ministerie van Infrastructuur en Waterstaat (2026), which are updated every month. The shuttle fee of MSC is 90 EUR for one TEU and 120 EUR for one FEU. With an average of 1.6 TEU per container (Van Koningsveld et al., 2021), 40 percent of the containers are 20-foot and 60 percent are 40-foot. The average fee is then 108 EUR per container, which is 67.50 EUR per TEU.

Table 5.2: Parameter values of the case study

Parameter	Value	Distribution	Source
Deepsea terminals			
Terminals in hub simulation	2 (ECT-DDN, ECT-DDE), 1 berth and 1 crane each	fixed	MSC, Assumption
Container batch size DST1 / DST2	100 / 80 containers	fixed	Assumption
Time between batches	mean 12 h	exponential	Assumption
Shuttle load split DST1 / DST2	60 / 40 percent	fixed	Nicolet et al. (2023)
Crane rate	max 20 moves/h	factor $N(0.92, 0.08)$, bounded to [0.5, 1.0]	Hutchison Ports ECT Rotterdam (2026)
Minimum barge call size	20 moves ECT-DDN, 10 moves ECT-DDE	fixed	Hutchison Ports ECT Rotterdam (2026)
Mooring/unmooring time	0.10 h per call	fixed	Van Es (2019)
Shuttle service			
Number of shuttles	1	fixed	MSC
Shuttle capacity	288 TEU	fixed	MSC
Timetable	3 cycles per week, arrival at BCW Mon, Wed, Fri around 06:00	fixed	MSC
Latest departure from deepsea terminal	02:30	fixed	MSC
Loading window	6 h per deepsea terminal	fixed	MSC
Sailing time deepsea-hub	mean 3.5 h each way	3.0 h + gamma delay	MSC, Assumption
Sailing time DST1-DST2	0.5 h	fixed	MSC
Antwerp shuttle days	Tue, Thu, Sat from 06:00	fixed	MSC
Hub			
Berths and cranes	1 berth, 2 cranes	fixed	MSC
Crane rate	max 25 moves/h in total	factor $N(0.92, 0.08)$, bounded to [0.5, 1.0]	MSC
Holding capacity	2,700 TEU (official 3,700 TEU)	fixed	MSC
Hub initial inventory	1,000 TEU	fixed	MSC
Opening hours	Mon-Fri 06:00-22:00, Sat 07:00-13:00, Sun closed	fixed	MSC
Antwerp shuttle at hub berth	3 times per week, 7 h each	fixed	MSC
Mooring/unmooring time	0.10 h per call	fixed	Van Es (2019)
Inland barges			
Barge capacity	M6 90, M8 192, M9 256 TEU	fixed	Veldman (2019)
Fleet size	per scenario, Table 7.2	calibrated	Simulation

Parameter	Value	Distribution	Source
Inland sailing and dwell	Figure 5.1	Weibull distribution	AIS data
Current-state simulation			
Deepsea terminals	1, with 1 berth and 1 crane	fixed	AIS data
Extra sailing compared with hub	2 $\times$ 3.5 h per round trip	fixed	MSC
Deepsea vessel blocks	5 per week, 21 h each (27 h in exp 4)	fixed	MSC
Call fee	0 EUR	fixed	Hutchison Ports ECT Rotterdam (2026)
Costs			
Sailing rate M6 / M8 / M9	158.69 / 194.66 / 199.59 EUR/h	fixed	Ministerie van Infrastructuur en Waterstaat (2026)
Waiting rate M6 / M8 / M9	94.50 / 116.87 / 115.27 EUR/h	fixed	Ministerie van Infrastructuur en Waterstaat (2026)
Shuttle fee	90 EUR per 20 ft, 120 EUR per 40 ft container	fixed	MSC
TEU per container	1.6	fixed	Van Koningsveld et al. (2021)

6

Verification and validation

Confidence in the results of a simulation study rests on two separate sayings. Robinson (1997) separates them as ‘building the model right’ and ‘building the right model’. Verification is the process of ensuring that the conceptual model has been transformed into a computer model with sufficient accuracy and validation is the process of ensuring that the model is sufficiently accurate for the goal of the thesis. In this chapter the structure of the verification and validation of the thesis is explained.

6.1. Framework and scope

The research of Sargent (2011) links each form of verification and validation to a phase of the model development process. Conceptual model validation belongs to the analysis and modelling phase, computerised model verification belongs to the programming and implementation phase, and operational validation belongs to the experimentation phase. The research of Robinson (1997) uses the same division and adds a distinction that is useful in this thesis. White-box validation is a micro check of whether each part of the model represents the corresponding real world element with sufficient accuracy, while black-box validation is a macro check of whether the overall model represents the real world with sufficient accuracy. Robinson notes that black-box validation is the only form that requires a completed model, because there is little sense in comparing the overall model against the real world before the model is finished.

This thesis uses two simulations. The current-state simulation describes the direct call as it operates today. The hub simulation describes the hub routing, which does not yet exist in the form that is modelled. The shared core is the inland leg distributions, the vessel classes, the cost rates, the container arrival process and the crane rate at the deepsea terminals. Both models use these parts identically under common random numbers so any error in them hits both sides in the same way and largely cancels in the paired differences. Both simulations are verified and only the current state simulation is validated against the real world, since this is the simulation that describes the real world.

The present chapter covers the whole verification and validation part including conceptual model validation, verification, white-box validation, and black-box validation. The first three are carried out while the model is being built. Operational validation, which corresponds to the black-box validation, requires the completed model together with its output.

6.2. Conceptual model validation

According to Sargent (2011), the validity of the conceptual model is based on determining whether the theories and assumptions underlying the model are correct and that the representation of the problem entity is reasonable for the intended purpose. The conceptual model and the flowcharts are checked by the company supervisor at MSC who has knowledge about this system. This ‘face’ validation is done by asking individuals who are knowledgeable about the system whether the model and its behaviour are reasonable. According to the research of Robinson (1997) this is an effective method. The conceptual model and flowcharts were reviewed this way during model development.

6.3. Verification and white-box validation

This research treats verification and white-box validation together because both are micro checks carried out continuously during model coding, and lists the aspects that should be checked as timings, the control of elements, the control of flows, the control logic and the sampling from distributions (Robinson, 1997). The verification part consists of tracing, internal consistency checks and extreme condition tests. After that a white-box parameter check is done. The tracing and the extreme condition tests are carried out on both models. The internal consistency checks are carried out on the hub simulation, because it contains the most elements and the most container flows of the two models.

6.3.1. Traces

A verification method mentioned by Sargent (2011) recommends tracing the progress of different types of specific entities through the model as a way of checking the logic. The model writes a print statement as an event trace for every state change. Figure D.1 in Appendix D shows the loops of a shuttle barge and an inland barge in a hub simulation run. Containers are held as a counted stock at the hub rather than as individually identified entities, so the two barge types are the entities a trace can follow. Reading the two lanes together is what makes the control logic checkable, because the waiting of the inland barge can be attributed to a cause only once the state of the shuttle and the state of the opening hours are visible on the same time axis. In the same Appendix D the tracing of the current state simulation can be found in Figure D.2. Figure D.2 also shows the block mechanism of the current-state simulation. At t = 0.00 h the first block is due while Barge 1 is loading, so the block waits. It starts when Barge 1 leaves at t = 3.35 h and holds the berth for 21 hours, during which Barges 2 to 5 join the queue. The block therefore takes priority over waiting barges without interrupting a barge that is already loading.

6.3.2. Internal consistency checks

Three identities on the container flow are checked at the end of every replication, which is the internal consistency check among the techniques of dynamic testing (Sargent, 2011). Together they state that no container is created or lost anywhere in the network. Every container the shuttle discharges at the hub must be accounted for as barge-bound stock. Every container that enters the hub stock must still be there at the end of the run or have been collected by an inland barge. Every container collected by an inland barge must have been delivered at the inland terminal or still be on board. The identities are exact rather than approximate, because each lists a quantity the model computes in one place using quantities it computes elsewhere. Therefore, any leftover above the rounding of double-precision arithmetic indicates a defect. Table 6.1 shows that every residual is exactly zero.

Table 6.1: Container flow identities asserted at the end of every replication, over the 30 replications of the M6 long baseline.

#	Identity	Residual	Result
C1	The shuttle discharge at the hub splits without remainder into the barge-bound share and the share destined for other modes	0 exactly	Pass
C2	The hub stock at the end of the run equals the initial inventory plus the barge-bound receipts minus the volume collected by inland barges	0 exactly	Pass
C3	The volume collected by inland barges equals the volume delivered at the inland terminal plus the volume still on board at the end of the run, and that remainder is never negative and never exceeds the capacity of the fleet	0 exactly	Pass

6.3.3. Extreme condition tests

The structure and the output of a model need to remain plausible under any extreme and unlikely combination of factor levels. Two of such settings are tested on the hub simulation and each is chosen because it drives a different mechanism of the model to a limit at which the correct behaviour can be stated in advance rather than judged after the fact. A third extreme setting is tested on the direct call simulation. The extreme condition tests have been executed on the M6 long leg scenario. Figures 6.1, 6.2 and Table 6.2 report the outcome.

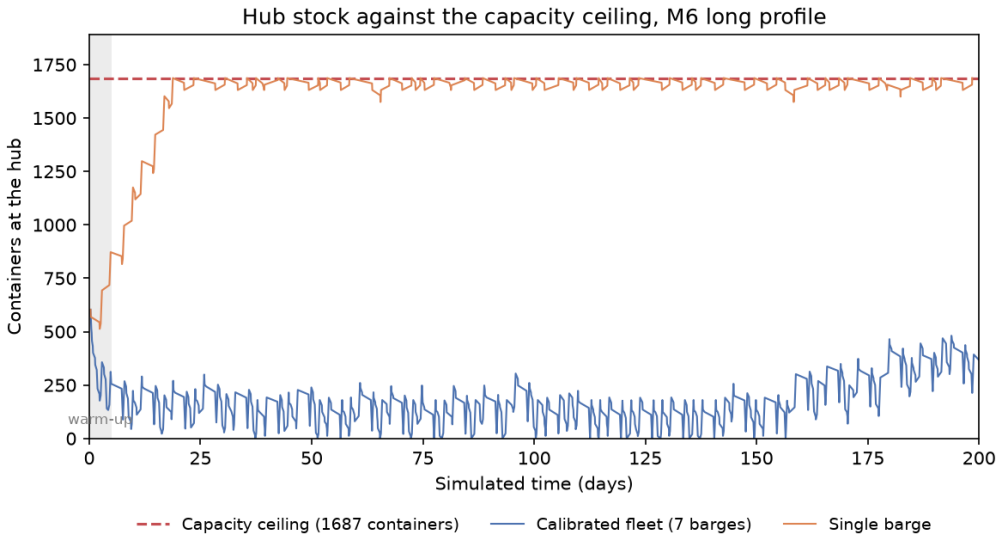

Figure 6.1: Hub stock over a full run on the M6 long profile

The first test reduces the inland fleet on the long profile to a single barge, which places the collection rate far below the rate at which the shuttle service delivers containers to the hub. The hub stock must then rise until it meets the capacity ceiling and remain against it. The containers that can no longer be transported must be held at the deepsea terminals rather than disappear. Figure 6.1 shows this behaviour. The peak stock reaches 1,687 containers, which is the ceiling of 2,700 TEU expressed in containers, and does not exceed it at any point in the run.

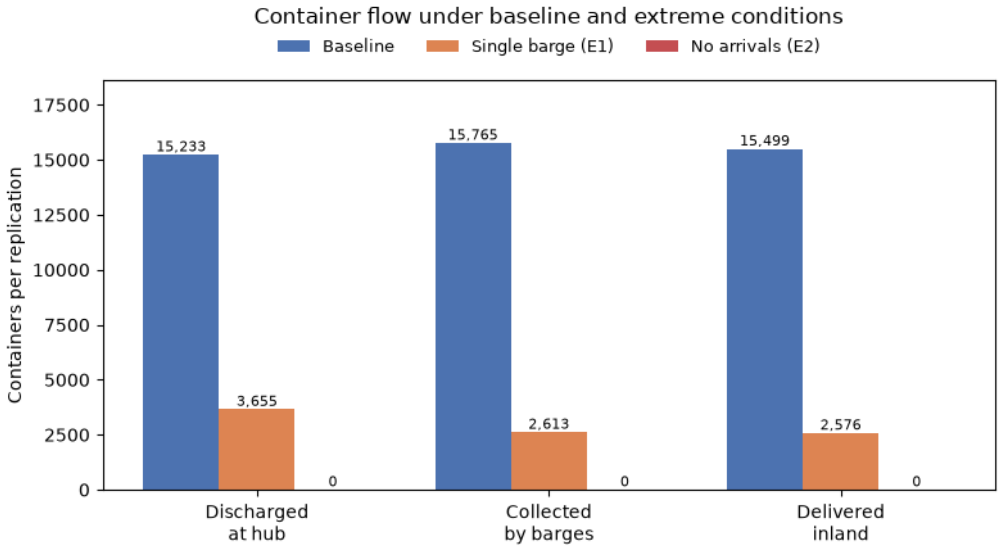

Figure 6.2: Container flow per replication under the baseline and under two extreme settings

Figure 6.2 shows two different settings to test the container flow. Reducing the inland fleet to a single barge on the long profile lowers the collected and delivered volumes while the hub operates against its ceiling. The second test removes the container arrivals at both deepsea terminals and the opening stock at the hub, which leaves the network with nothing to move. All three flow quantities fall to zero in every replication, as Figure 6.2 shows. The value of this test lies in the collapse being simultaneous. A single quantity falling to zero can be produced by a number of different faults, whereas the discharge, the collection, the inland delivery and the deepsea backlog reaching zero together is consistent only

with the network being empty. The flow counters in Figure 6.2 cover the whole run, including the warm-up. The collected volume exceeds the discharged volume by 532 containers, because the hub starts the run with an opening stock of 1,000 TEU (625 containers), which the inland barges collect in addition to the shuttle deliveries. The delivered volume is lower than the collected volume because part of the collected containers are still on board at the end of the run. In the single-barge case this is never more than 56 containers, which is exactly the capacity of one M6 barge.

The last extreme condition is tested for the current state simulation. In this test the arrival of containers at the deepsea terminal is set to zero. This would mean inland barges cannot collect any containers resulting in no cycles during the simulation run and no containers delivered at the inland terminals. Table 6.2 shows the results of the third extreme condition test.

Table 6.2: Extreme condition test on the current state simulation: no container arrivals at the deepsea terminal

Quantity (per replication)	Expected	Observed	Result
Completed cycles	0	0	Pass
Berth calls at the deepsea terminal	0	0	Pass
Containers delivered inland	0	0	Pass
Containers left at the deepsea terminal	0	0	Pass

As predicted the containers do not reach the hinterland since the containers delivered to the inland are equal to zero and the completed cycles share the same result. Because no containers appear at the deepsea terminal the containers left at the deepsea terminal are also zero. Without arrivals, all seven barges wait for supply for the entire run, so none reaches the berth. These results show that also the current-state simulation passed the extreme condition test.

6.3.4. White-box parameter checks

The checks above establish that the model is true to the conceptual model. White-box validation asks the further question of whether each constituent part is true to the real world (Robinson, 1997). This has already been presented in the previous chapter. Table 5.2 in Section 5.2 lists the parameters and states the external source against which each was checked or where the information was extracted from.

6.3.5. Coverage of the verification

The checks above concentrate on the mechanisms that carry the conclusion, which are the hub berth logic, the opening-hours logic, the waiting-time decomposition and the cost layer. Parts of the model that affect both routings identically, such as the inland terminal representation, receive less attention by design. The verification checks show that the model works as intended. Section 6.4 validates the current-state simulation.

6.4. Operational validation

The verification has checked whether the model is built right. The operational validation checks whether the right model is built for the goal of this thesis. It compares the in-port time of the current-state simulation with the in-port time observed in the AIS data. In total there are 110 cycles with in-port time observed, since these cycles contain the three classes that are taken into account in this thesis. The other cycles consist of other barge classes. In the operational validation only the current-state simulation is validated, since the hub routing does not yet exist in the form that is modelled.

Each observed barge is attributed to a vessel class with a list of barges including their lengths and beams per class of Van Koningsveld et al. (2021). The observed in-port time per class is the mean over all cycles of that class. The simulated in-port time is the mean deepsea time per call of the short and the long scenario, without the seven-hour round trip. To test whether the difference between the model and the data is larger than the spread in the data, a 95% confidence interval is computed for each observed mean with the t-distribution. The simulated value is consistent with the data when it lies inside this interval. The simulated values are means over thirty replications and their own intervals

are narrow compared to those of the observed data. The results of the 95% confidence interval are presented in Table 6.3.

Table 6.3: Operational validation of the current state simulation

Class	Cycles	Barges	Observed (h)	95% CI (h)	Simulated (h)	Difference (h)
M6	43	12	18.59	14.83 - 22.35	12.20	-6.39
M8	54	19	14.00	11.38 - 16.61	15.55	$+1.55$
M9	13	7	15.84	8.34 - 23.33	18.95	$+3.11$

For M8 and M9 the simulated in-port time lies inside the interval of the observed mean. For these classes the model cannot be distinguished from the data. For M6 the simulated value lies below the interval, so the model underestimates the M6 in-port time by 6.39 hours. The M9 interval is wide because it rests on only 13 cycles of seven barges, so the validation of this class is weak.

Two reasons explain the difference between the classes. The first is the fixed window. The dataset contains barges with and without a fixed window and the data does not show which is which. The waiting time before loading points in this direction. For M6 only 16 percent of the cycles wait less than one hour, against 59 percent for M8 and 46 percent for M9. The mean waiting time is 12.08 hours for M6, against 4.30 hours for M8 and 7.90 hours for M9. The M6 cycles therefore look like calls without a fixed window, while a large share of the M8 and M9 cycles looks like calls with a fixed window. This also explains why the observed in-port time does not rise with the size of the class, as the longer loading time of a larger barge would suggest. The second reason is the excluded export flow. The observed in-port time in the data includes both loading and unloading, while the model only contains the import flow. The model is therefore expected to give a shorter in-port time than the data.

The purpose of the current-state simulation is not to predict in-port times exactly, but to provide a fair baseline for the comparison with the hub routing. For M8 and M9 the model is consistent with the data. For M6 the model underestimates the in-port time, which makes the direct call look faster and cheaper than it is. This works against the hub, so the M6 results in chapter 7 are too unfavourable for the hub. The congestion experiment tests a longer deepsea waiting time and shows how the conclusion on time changes for the M6 class.

7

Numerical results

7.1. Experimental framework and performance indicators

The numerical chapter shows the simulation results of the tested scenarios and experiments. The output of the scenarios shows the cost and time effect of the hub routing per barge profile, compared to the direct call. Every scenario is a hub configuration compared to the simulated direct call. The direct call is the baseline of this thesis to which the hub configurations are compared. The three barge classes (M6, M8, M9) are crossed with the long and short leg, resulting in 6 baseline scenarios. The fleet size is calibrated per scenario based on the throughput of the shuttle barge in the hub-and-spoke model. The experiments show the output of operational settings and sensitivity analyses, to find conditions under which the cost gap could be closed. In total this research tests four experiments. The first three are operational experiments and the fourth is a structural experiment, used for the sensitivity analyses.

The two key performance indicators follow directly from the research question, namely the barge turnaround time in hours and the operator cost per TEU in euros. Both are collected per replication and reported with 95% confidence intervals. Every remaining indicator either decomposes these two or compares a scenario against another scenario or against the current state. The turnaround time is decomposed into the inland leg and the hub-side time per call. The hub-side time splits further into loading time and waiting at the hub. Finally, the waiting consists of three mechanisms, the berth queue, the hours during which the hub is closed, and supply starvation. This decomposition is what allows the waiting time of each vessel class to be traced to its operational cause. The cost per TEU is decomposed into sailing, waiting, and shuttle fee of the barge, following the cost layer of Section 4.6.

The comparison against the current operation is expressed in two forms. The turnaround ratio and the cost ratio give the hub value relative to the direct call and the gap gives the same difference as an absolute amount in euro. The absolute gap is the preferred measure for comparisons across the two inland-leg profiles, because the ratio is inflated by the inland sailing cost that both routes share, whereas the absolute gap isolates the part of the cost that the hub actually changes. A set of operational quantities supports the fleet calibration and the throughput analysis rather than the headline comparison. These are the number of completed cycles and the weekly deliveries they produce. The containers delivered per replication doubles as the common denominator of the cost per TEU and as a check that each configuration clears the same weekly flow. The hub inventory over time is kept as a diagnostic in Appendix C that confirms that the supply ceiling and not the inland fleet limits the system.

The chapter first explains the value of the warm-up time of the simulation model with the use of the SPC method. The final results of the scenarios and experiments with their explanations can be found in Section 7.5 and Section 7.6.

7.2. Warm-up period

Each replication starts from an empty and idle state. There are no containers transported to the hub and every inland barge is initialised at the start of its cycle. As a result, the first observations produced by the model are biased in relation to steady-state operation. Including these first observations in the averages would make the comparison between configurations inaccurate. A warm-up period is therefore used from the start of every run before any key performance indicator (KPI) is collected (Law, 2020).

The length of the warm-up period is calculated with the statistical process control (SPC) method of Robinson (2007). The method rests on the observation that, once a simulation has reached steady state, its output behaves like a process that is in statistical control. The end of the warm-up period is the moment at which the output series first enters and remains within its control limits. The control limits are placed at three standard deviations above and below the center line which presents the mean of the steady state of the turnaround time (Robinson, 2007). The warm-up period is based on the barge turnaround time because it is the slowest converging KPI in the model. The waiting time at the hub and the hub berth utilisation both settle within the first six hours. The turnaround time, which reflects a complete loop between the hub and the inland terminals, requires several cycles before the fleet reaches its steady-state distribution around the network. As the warm-up period must be long enough to remove the initialisation bias from the slowest measure reported, the turnaround time is the binding measure and is used to set the warm-up for all scenarios.

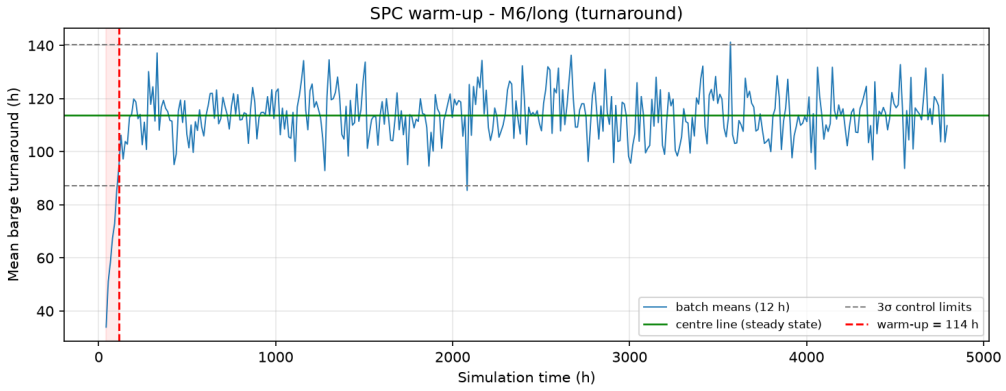

Figure 7.1: Warm-up period against mean barge turnaround time for scenario M6 with a long inland leg

Figure 7.1 shows the SPC control chart for a representative scenario (vessel class M6, long inland leg). The graph rises steeply from the empty-system state and levels out around the center line. The warm-up marker is placed where the graph enters and remains within the limits. Table 7.1 presents the warm-up period obtained for each of the six scenarios. The values range from a few hours for the fast settling short leg scenarios to 114 hours for the M6 class with a long inland leg.

Table 7.1: Warm-up period per scenario, obtained with the SPC method of Robinson (2007) on the inland barge turnaround time.

Vessel class	Inland leg	Warm-up (h)
M6	short	42
M6	long	114
M8	short	42
M8	long	90
M9	short	30
M9	long	90
Common warm-up (maximum across scenarios)		114

The largest warm-up period found across all scenarios is rounded up to 120 hours (5 days) and used as a single common warm-up period for every scenario. This is a long warm-up period for the faster settling scenarios, but it removes only about 2% of the 200-day run which means a negligible loss of data. The control charts for all six scenarios are collected in Appendix C.

7.3. Replications and statistical analysis

Each scenario is run for thirty replications of 200 days. After the warm-up of 120 hours, every replication gives about 28 weeks of output. Each replication uses its own master seed, so the replications are independent of each other. Within a replication each stochastic source draws from its own random stream, and all five streams derive from the master seed. The scenarios are run under common random numbers (Glasserman & Yao, 1992). Because all scenarios share the same thirty master seeds, every scenario sees the same random inputs for a given seed. A comparison between two configurations is therefore made as a paired difference per seed. The pairing removes the sampling noise that both runs share and gives a narrower confidence interval on the difference than two independent samples would.

Mean values and paired differences are reported with a 95% confidence interval. The interval is the mean $\pm t_{0.975,29}\, s/\sqrt{30}$, where s is the standard deviation across the thirty replications. The t-distribution assumes that the replication means are close to normal. This is reasonable here, because each replication mean is an average over about 28 weeks of output. The cost per TEU is formed as a ratio within each replication and then averaged across replications, so its interval reflects the run-to-run variation in the ratio itself. A difference is significant at the 5% level when its 95% interval does not contain zero. Whether thirty replications are enough is checked with the relative precision of the results, which is the half-width of the interval divided by the mean (Law, 2020). A relative precision below 5 percent is used as the target. The outcome of this check is reported with the baseline results in Section 7.5.

7.4. Scenarios and fleet calibration

7.4.1. Fleet-size calibration

Each scenario requires a fleet size for the third-party inland barges that is large enough to absorb the container flow generated by the shuttle service, but not so large that barges needlessly queue for the single berth at the hub. The fleet size was therefore calibrated by sweeping the number of inland barges and selecting the “knee” of the throughput curve. This is the smallest fleet at which weekly deliveries reach the system ceiling, before additional vessels only add berth queue time. Table 7.2 summarises the calibration.

Table 7.2: Fleet-size calibration per scenario at the hub.

Class	Leg	Fleet at knee (barges)	Ceiling (containers/week)	Berth queue (h/visit)	Inland berth utilisation (%)
M6	short	4	∼545	11.5	57.2
M6	long	7	∼544	14.3	57.0
M8	short	2	∼542	5.4	52.8
M8	long	3	∼543	6.5	50.6
M9	short	2	∼540	9.8	54.3
M9	long	3	∼543	12.2	54.5

Two observations follow. First, the throughput ceiling is essentially identical across all six scenarios (540–545 containers per week), confirming that the binding constraint is not the inland fleet but the supply side of the hub. The shuttle service and the opening hours regime of the hub cap the flow that can be passed on, irrespective of vessel class. Second, the fleet required to reach that ceiling increases with cycle length and decreases with vessel capacity. The small M6 class needs four (short leg) to seven (long leg) barges to sustain the flow, whereas the M8 and M9 classes reach the same ceiling with two to three vessels because each call moves two to three times as many containers. Beyond the minimum number of barges needed, deliveries remain flat while the berth queue time grows steeply. This is why in all scenarios the knee fleet sizes are used.

7.4.2. Scenario definition

All six simulated scenarios are hub-and-spoke configurations. They are based on the inland barge class and their turnaround leg.

1. **Inland vessel class:** The choice of class fixes both the simulated barge capacity and the cost rates used in the cost layer through a single setting, so the two can never drift apart. The studied classes are M6 (90 TEU), M8 (192 TEU), and M9 (256 TEU).
2. **Turnaround leg:** The short leg and long leg empirical samplers drawn from the median split of the AIS-derived inland trips (Section 4.5.2). The two legs capture the difference between short and long hinterland connections, since the hub removes deepsea in-port time rather than inland sailing time.

The current-state simulation provides the third element of the comparison. It is the reference against which each simulated hub scenario is scored, supplying the Maasvlakte in-port time per class that the hub removes. The net time effect of the hub for a given scenario is therefore the removed in-port time per class minus the additional hub-side time the scenario attracts in the hub simulation. The in-port time of the inland barges in the current-state simulation consists of the in-port time at the Maasvlakte area plus the time it takes to travel from the hub to the Maasvlakte (3.5 hours both ways). This is an important detail since the hub in the hub situation excludes the sailing time between the hub and the deepsea terminals for the inland barges.

7.5. Baseline results

7.5.1. Steady state performance of the hub-and-spoke

Table 7.3 reports the primary performance metrics of the six hub-and-spoke configurations at their calibrated fleet sizes. Each scenario delivers approximately 15,000 containers over the 28 effective weeks per replication.

Table 7.3: Hub performance of the baseline scenarios

	M6		M8		M9	
Metric	short	long	short	long	short	long
Fleet size (barges)	4	7	2	3	2	3
Turnaround time (h)	58.5	114.2	58.5	112.8	64.7	121.9
± 95% CI	0.3	0.9	0.5	1.4	0.6	1.5
Hub-side time per call (h)	19.8	22.7	19.7	20.8	25.9	30.0
Loading time (h)	3.6	3.6	6.7	7.0	9.5	11.1
Waiting at hub (h)	16.1	19.0	12.8	13.8	16.4	18.8
berth queue (h)	11.5	14.3	5.4	6.5	9.8	12.2
hub closed (h)	4.3	4.6	6.8	7.2	6.0	6.3
supply starvation (h)	0.3	0.2	0.6	0.0	0.6	0.2
Cost per TEU (EUR)	155.58	232.71	121.66	157.23	121.81	154.14
± 95% CI	0.27	0.56	0.24	0.81	0.25	0.40
Containers delivered	15,181	15,161	15,129	14,757	15,033	15,115

A few observations can be made from Table 7.3. The first observation is the increasing turnaround time with the larger vessel classes within each inland leg. The decomposition of the hub-side time shows that the loading time is the main cause. Larger barges load more containers per call, which raises the loading time from 3.6 hours for the M6 class to 6.7-7.0 hours for M8 and 9.5 to 11.1 hours for M9. The M9 class has the longest loading time which is partially due to the fact that the loading is regularly interrupted by the closing of the hub. This pushes the hub-side time per call to 25.9-30.0 hours, which is high above the other two classes. The second observation is about the waiting time of class M6, which is dominated by the berth queue (11.5-14.3 hours). This is a direct consequence of the large fleet of four to seven barges that this class needs to handle the container flow, while using the same single berth at the hub. The M8 and M9 fleets consist of only two to three barges, which shrinks the berth

queue. For M8 the hub-closed component becomes the largest waiting component. With fewer vessels in circulation, arrivals land outside opening hours relatively more often. However, the berth queue of M9 (9.8-12.2 hours) is nearly double the berth queue of M8 at the same fleet size, because each M9 call occupies the berth considerably longer. Supply starvation is close to zero in every scenario, which confirms that the fleet calibration balanced the inland fleet against the shuttle supply.

The third observation is about the cost structure. The M6 class is the most expensive per TEU despite having the lowest hourly cost rate, because its small capacity spreads the cycle cost over few containers. Therefore the price per TEU for the M6 class is higher than the price per TEU for the M9 class. This is called economies of scale. When the number of containers would be the same on all the barges, the barges of the M6 class would have the lowest price per TEU.

A final observation concerns the number of containers moved per call, reported in Table 7.4. On the short leg every class sails noticeably emptier than on the long leg and the effect is strongest for the M9 class, which departs the hub at only 66% of its full capacity. The mechanism is the result of the loading rule and the supply side of the hub. An inland barge departs when the hub inventory is empty, not when the vessel is full. The load per call equals the number of containers the shuttle service has delivered since the previous visit, until the maximum barge capacity is reached. Because the shuttle supply is fixed at approximately 540 containers per week in all scenarios, the call size is the weekly supply divided by the weekly visit frequency of the fleet. On the short leg the cycle time roughly halves, so the calibrated fleet calls at the hub almost twice as often and the same supply is split into smaller portions. In Appendix C the hub capacity per class and short/long leg is shown for one replication. These figures confirm the mechanism of the system.

Table 7.4: Call size and vessel capacity utilisation per baseline scenario. The minimum frequency is the visit rate at which the weekly supply would be moved with fully loaded vessels.

Class	Leg	Capacity (containers)	Containers per call	Fill (%)	Visits/week	Min. visits/week
M6	short	56	48.1	86	11.5	9.7
M6	long	56	53.1	95	10.3	9.7
M8	short	120	96.5	80	5.7	4.5
M8	long	120	119.0	99	4.5	4.4
M9	short	160	104.8	66	5.2	3.4
M9	long	160	131.9	82	4.1	3.4

The effect per class is explained by the fleet size. The last column of Table 7.4 gives the minimum visit frequency at which the weekly supply could be moved with fully loaded barges. For an M9 barge in a short leg this minimum is 3.4 visits per week, but a single M9 barge provides only 2.60 visits per week, which is insufficient. The calibration must therefore be increased to two barges, which supply 5.2 visits per week. This is fifty percent more visit capacity than the supply can fill. Therefore, the average call drops to two thirds of a load. The M6 class sits at the other end of the spectrum: its small capacity requires at least 9.7 visits per week, close to the 11.5 its fleet makes, so it sails nearly full.

This has two consequences for the results above. First, it explains why the M9 class realises no scale advantage over M8 on the short leg. Their cost per TEU is practically identical (121.81 versus 121.66 EUR), because the M9 barge capacity is neutralised by partial loads while its hourly cost rate is higher. On the short leg the binding constraint is not the vessel capacity but the supply of the shuttle service at the hub. Deploying a larger barge class at the current shuttle throughput has no benefits in this simulation. Second, the near-full M8 long leg scenario illustrates the opposite failure mode. At 99% fill the fleet has no slack, and it is the only scenario that undershoots the supply ceiling (approximately 530 versus 543 containers per week). The call size is thus a system outcome of supply and visit frequency rather than a barge property, and it will shift when interventions change the turnaround time.

It can be concluded that a schedule misalignment exists for M8 and M9, and a queueing problem for M6 caused by its fleet size. An extension of the opening hours or a different shuttle delivery schedule are therefore candidate interventions, which are taken up in Section 7.6.2.

7.5.2. Performance of the current state

Table 7.5: Direct-call performance of the baseline scenarios

	M6		M8		M9	
Metric	short	long	short	long	short	long
Fleet size (barges)	4	7	2	3	2	3
Turnaround time (h)	57.7	111.1	61.2	115.1	64.4	118.8
± 95% CI	0.3	0.6	0.5	1.2	0.5	1.2
Deepsea time per call (h)	19.1	19.3	22.3	22.8	25.6	26.3
Loading time (h)	3.1	3.1	6.6	6.6	8.8	8.8
Waiting at deepsea (h)	8.8	9.0	8.6	9.0	9.7	10.4
berth queue (h)	8.8	9.0	8.6	9.0	9.7	10.4
terminal closed (h)	0.0	0.0	0.0	0.0	0.0	0.0
supply starvation (h)	0.0	0.0	0.0	0.0	0.0	0.0
Cost per TEU (EUR)	79.79	158.12	47.97	93.38	37.74	72.52
± 95% CI	0.42	0.89	0.37	1.18	0.30	0.96
Containers delivered	18,196	16,546	18,384	14,678	23,269	18,978
± 95% CI	97	100	143	175	167	226

Table 7.5 shows the output of the current-state simulation. The same six scenarios are run for the current situation, in which the barge sails past Waalhaven to the Maasvlakte and calls at the deepsea terminal itself. The fleet size is taken from the hub calibration rather than recalibrated, because the comparison is a routing choice made by one operator with one fleet and not a comparison between two differently sized operators within the same class.

The structure of the port call differs from the hub in a way that matters more than the totals. Waiting at the deepsea terminal consists of the berth queue alone. The terminal has no opening hours, so the closed component is zero. Furthermore, the container stock is never empty so the supply starvation component is zero as well. The hub introduces both of these. Waiting at the deepsea berth is almost flat in the current-state simulation across the six scenarios, which has to do with the fact that the same block of deepsea vessels at the berth of the direct call is used for all six scenarios. All of the variation in the deepsea call therefore comes from loading, which rises from 3.1 hours for the M6 class to 8.8 hours for the M9 class. This is the result of the container arrival at the deepsea terminal that gives barges the opportunity to load to their full capacity. The deepsea call runs from 19.1 to 26.3 hours and includes the seven-hour round trip between Waalhaven and the Maasvlakte. That leg is included, because the shuttle system takes over those hours in the hub-and-spoke situation. The time an inland barge spends at the hub is compared with the time the inland barge does not have to deal with in the hub situation, which includes the sailing hours between the hub and the deepsea terminal, and the in-port time.

Cost per TEU falls steeply with vessel size, from 79.79 EUR for the M6 class on the short leg to 47.97 EUR for M8 and 37.74 EUR for M9. Two effects drive this. The hourly cost of a barge rises slower than its capacity and the free terminal fee per call is fixed for barges without fixed windows. The barge loads to its full capacity at every call in all six scenarios, which is the direct consequence of a container stock that never runs dry. This is the effect of economies of scale when looking at the transport costs per TEU of the different barge classes.

The confidence intervals in Table 7.3 and Table 7.5 confirm that thirty replications are enough. For the turnaround time, the cost per TEU and the containers delivered, every half-width is below 1.5 percent of its mean. This is well below the target of 5 percent.

7.5.3. Decomposition of hub-side waiting time

Table 7.6 decomposes the waiting time per hub visit into its three components: supply starvation (the barge is at the berth but no containers are available), the berth queue (another barge occupies the single berth) and the hub-closed component (the visit overlaps with hours outside the 86 h/week opening regime).

Table 7.6: Decomposition of hub-side waiting time per barge visit

Class	Leg	Starvation (h)	Berth queue (h)	Hub closed (h)	Total (h)
M6	short	0.32 (2.0%)	11.53 (71.5%)	4.27 (26.5%)	16.12
M6	long	0.17 (0.9%)	14.23 (75.0%)	4.57 (24.1%)	18.98
M8	short	0.59 (4.6%)	5.41 (42.3%)	6.80 (53.1%)	12.81
M8	long	0.02 (0.1%)	6.54 (47.6%)	7.20 (52.3%)	13.76
M9	short	0.57 (3.5%)	9.78 (59.8%)	6.00 (36.7%)	16.35
M9	long	0.25 (1.3%)	12.23 (65.2%)	6.27 (33.5%)	18.75

Three patterns stand out. First, the dominant component differs per class. For M6 the berth queue accounts for 72 to 75% of all waiting, which is the largest part of berth queue compared to the other classes. For M8 the hub-closed component is the largest share (52-53%). The berth queue therefore scales with the pressure on the single berth. Many small barges (M6) or long berth occupations per call (M9) both congest the berth, while the two to three fast loading M8 barges leave it comparatively free. Second, the hub-closed component is remarkably stable across all six scenarios (4.3-7.2 hours per visit). It is a structural property of the opening hours setting rather than of the vessel class or the inland-leg, which makes it a natural target for an opening hours experiment. The last observation that stands out is the supply starvation, which is negligible everywhere ($\leq$0.6 h) and is even smaller on the long legs. With longer cycles the hub inventory has more time to replenish between the visits of the same barge.

7.5.4. Current vs hub situation

This section compares each hub-and-spoke state against the current state, in which the inland barge sails to the Maasvlakte and calls at the deepsea terminal directly. The two routings use the same fleet, the same inland leg samplers and the same thirty seeds. The hours the hub removes are split between sailing and waiting time. The fee advantage is the free deepsea call against the average 108 EUR per container shuttle fee on the hub call for inland barges without fixed windows. The seaport-side time of the current state is the deepsea port call plus the seven-hour round trip between Waalhaven and the Maasvlakte, because that leg disappears when the shuttle takes it over. The inland leg is identical in both routes, so the difference in seaport-side time per call has to equal the difference in turnaround time.

Table 7.7: Decomposition of the time-cost advantage per deepsea call

		Direct call removed			Hub-side call added		Time-cost	Fee	Total
Class	Leg	Sailing (EUR)	Waiting (EUR)	Total (EUR)	Hours (h)	Cost (EUR)	adv. (EUR)	adv. (EUR)	adv. (EUR)
M6	short	1,111	1,139	2,250	19.8	1,874	+376	−5,195	−4,819
M6	long	1,111	1,159	2,269	22.7	2,144	+125	−5,735	−5,610
M8	short	1,363	1,794	3,157	19.7	2,298	+859	−10,422	−9,563
M8	long	1,363	1,845	3,208	20.8	2,436	+772	−12,852	−12,080
M9	short	1,397	2,146	3,543	25.9	2,991	+552	−11,318	−10,766
M9	long	1,397	2,229	3,626	30.0	3,460	+166	−14,245	−14,079

The hours the hub removes are split between sailing and waiting time, whereas the hub-side call counts waiting time only. Table 7.7 separates the two. An M6 barge on the short leg gives up seven hours of sailing at 158.69 EUR per hour and 12.05 hours at the berth at 94.50 EUR per hour, which is a total of 2,250 EUR per call. What the M6 barge gets in return are 19.8 hub-side hours at 94.50 EUR per hour, which results in 1,874 EUR per call. The hub call in the M6 scenario is 0.78 hours longer than the direct call it replaces and still produces a time-cost advantage of 376 EUR, because the hours it removes are the expensive sailing hours. However, that advantage is small against the fee. The shuttle fee is charged per container at 90 EUR per TEU and 120 EUR per FEU, which is an average of 108 EUR per move once the factor of 1.6 TEU is applied. The deepsea call is free for barges without fixed windows.

On the M6 short leg the operator therefore trades a 376 EUR saving on time for a 5,195 EUR increase in fees. The same pattern holds in all six scenarios and the fee term is between 5,195 and 14,245 EUR per call against a time-cost advantage that never exceeds 859 EUR per call. Because the direct call is free while the shuttle fee grows with every container, the fee per call grows with vessel size.

An important note to report on the numbers is the limited shuttle supply to the hub. The shuttle supply is 540 containers per week, whereas the deepsea stock is never empty. The ceiling is therefore set by the shuttle payload and frequency, while the modelled holding capacity of the hub at 2,700 TEU is never binding. Therefore, the call size at the hub-and-spoke is a system outcome of shuttle supply and visit frequency rather than a vessel property. This is why the M9 class realises no scale advantage over M8 on the short leg (121.81 versus 121.66 EUR/TEU) in the hub situation. Its extra capacity is neutralised by partial loads while its hourly cost rate is higher. The economies of scale principle is therefore fixed to a ceiling, because of the limited amount of containers delivered to the hub per week.

7.5.5. Paired differences hub and current situation

Because both routings are simulated, the comparison can be closed as a paired difference rather than a difference of means. The two routings share the same fleet, the same inland leg samplers and the same thirty master seeds, so the difference is taken per seed and the sampling noise common to both runs drops out. Table 7.8 reports the result. The sign convention is hub minus direct, so a positive value means the hub takes longer, waits longer or costs more. Almost all differences are significant at the 5 percent level. The three exceptions are marked in the table.

Table 7.8: Paired differences between the hub and the direct-call situation

Class	Leg	Δ Turnaround (h)	Δ Seaport-side (h)	Δ Waiting (h)	Δ Cost per TEU (EUR)	Δ Containers delivered
M6	short	$+0.81 \pm 0.28$	$+0.77 \pm 0.29$	$+7.25 \pm 0.28$	$+75.79 \pm 0.36$	-3015 ± 96
M6	long	$+3.11 \pm 0.62$	$+3.44 \pm 0.56$	$+9.91 \pm 0.53$	$+74.59 \pm 1.05$	$-1{,}385 \pm 110$
M8	short	-2.68 ± 0.42	-2.69 ± 0.40	$+4.14 \pm 0.38$	$+73.69 \pm 0.31$	$-3{,}255 \pm 139$
M8	long	-2.29 ± 0.53	-1.94 ± 0.51	$+4.67 \pm 0.51$	$+63.85 \pm 0.79$	$+79 \pm 126^{*}$
M9	short	$+0.36 \pm 0.48^{*}$	$+0.33 \pm 0.46^{*}$	$+6.60 \pm 0.42$	$+84.07 \pm 0.27$	$-8{,}236 \pm 165$
M9	long	$+3.14 \pm 1.03$	$+3.69 \pm 1.12$	$+8.33 \pm 0.93$	$+81.62 \pm 0.85$	$-3{,}863 \pm 219$

* not significant at the 5% level.

The turnaround and seaport-side columns agree to within 0.55 hours in every scenario. This follows from the construction. The inland leg is drawn from the same sampler in both routes, so the only place a difference can arise is the seaport side and the agreement confirms that the seven-hour round trip has been attributed correctly. On time, the hub is faster than a direct call only for the M8 class, by 2.68 hours on the short leg and 2.29 hours on the long leg. The M9 short result is the closest to neutral at 0.36 hours against an interval of 0.48, and it is the only time difference that is not comfortably clear of its interval. Waiting tells a more consistent story than the port call as a whole. The hub waits longer in every one of the six scenarios, by 4.14 to 9.91 hours per visit. Where the hub still wins on the total call comes from the seven hours of sailing the shuttle absorbs, but not from any reduction in queueing. The hub replaces a single berth queue with three waiting mechanisms and the net effect on waiting is negative on all six scenarios. On cost no scenario favours the hub. The gap runs from 63.85 EUR per TEU for M8 on the long leg to 84.07 EUR per TEU for M9 on the short leg. The primary reason for this is the shuttle fee per container that needs to be paid in the hub situation, compared to no fee at the deepsea terminal in the current situation.

The final column records what the two routings actually move with the same fleet. In five of the six scenarios the current situation delivers between 1,385 and 8,236 more containers over the same run, because the hub is limited by the shuttle supply while the deepsea terminal stock is never empty. The exception is the M8 class on the long leg, where the hub barge already departs at 99 percent of capacity and the difference of 79 containers is the smallest of the six. That scenario also carries the smallest cost gap. The hub loses least where its supply ceiling binds least.

7.5.6. Short versus long inland-leg

The effect of the inland-leg is tested with paired replications. The short and long scenarios of each class share common random numbers, so the paired differences in Table 7.9 isolate the leg effect from sampling noise. The table shows the difference between the long and short leg with a focus on cost and time.

Table 7.9: Paired differences in time and cost (long minus short per vessel class)

Class	Δ Cost per TEU (EUR)	Δ Turnaround (h)	Δ Wait at hub (h)
M6	$+77.13 \pm 0.58$	$+55.76 \pm 0.81$	$+2.87 \pm 0.57$
M8	$+35.57 \pm 0.69$	$+54.30 \pm 1.19$	$+0.96 \pm 0.56$
M9	$+32.33 \pm 0.38$	$+57.19 \pm 1.26$	$+2.43 \pm 0.95$

All differences are significant at the 5% level.

Two conclusions follow. First, the turnaround penalty of the long leg is almost entirely explained by the extra sailing itself. The mean inland leg is approximately 53 hours longer on the long leg and the turnaround differences of 54.30 to 57.19 hours match this closely. The hub-side contribution is marginal. The waiting difference is at most 2.9 hours (M6). The hub therefore treats short leg and long leg operators identically, which means a long leg operator does not queue meaningfully longer than a short leg one. Only for M9 does the turnaround difference exceed the sailing difference by a few hours, which is caused by the longer loading time (11.1 versus 9.5 hours) combined with the slightly larger hub-closed exposure. Second, the cost penalty of the long leg decreases with vessel size, from 77.13 EUR/TEU for M6 to 32.33 EUR/TEU for M9. The extra sailing hours of the long leg are a fixed time quantity per cycle, which means larger vessels spread the cost over two to three times as many containers. This explains the lower costs for larger vessels.

7.5.7. Summary of findings

The comparison between the hub and current situation presents mostly disadvantages at the expense of the hub configurations on both the cost and time performance indicators. On time, the hub is faster than a direct call only for the M8 class, by 2.68 hours on the short leg and 2.29 on the long leg. In every other scenario the direct call is quicker or almost the same time. On cost, the shuttle fee is a large factor, because the shuttle fee is charged per container compared to a free deepsea fee for inland barges without fixed windows. The M8 long leg carries the smallest gap of the six at 63.85 EUR per TEU. The highest gap can be found on the short leg of the M9 class with 84.07 EUR/TEU. A second structural result is that the hub caps the system. All six scenarios converge on the same throughput ceiling of roughly 540 containers per week. That number is not a property of the hub. The shuttle carries 288 TEU, which is 180 containers at the conversion factor of 1.6 and it runs three cycles per week.

Table 7.10: Decomposition of the cost gap into sailing, waiting and fees. Waiting includes the waiting at the hub or the deepsea terminal, the loading time and the dwell at the inland terminal, all at the waiting rate.

		Hub (EUR/TEU)			Direct (EUR/TEU)			Difference (EUR/TEU)		
Class	Leg	Sail	Wait	Fee	Sail	Wait	Fee	Sail	Wait	Gap
M6	short	40.34	47.74	67.50	47.01	32.77	0.00	-6.67	$+14.97$	$+75.79$
M6	long	94.16	71.05	67.50	101.34	56.78	0.00	-7.18	$+14.27$	$+74.59$
M8	short	24.71	29.45	67.50	26.94	21.03	0.00	-2.23	$+8.42$	$+73.69$
M8	long	51.34	38.38	67.50	58.06	35.31	0.00	-6.72	$+3.07$	$+63.85$
M9	short	23.25	31.06	67.50	20.69	17.05	0.00	$+2.56$	$+14.01$	$+84.07$
M9	long	47.45	39.19	67.50	44.70	27.82	0.00	$+2.75$	$+11.37$	$+81.62$

The results of the M8 class on cost and time would suggest that class to be the most promising class to close the gap. However, the operational experiments act on the waiting at the hub and for M8 the largest part of that waiting is caused by the closing hours of the hub. For M6 the berth queue is 72%

to 75% of the waiting, because four to seven barges share a single berth. A change at the hub that shortens this queue can free enough capacity to remove a barge and one barge is a large part of the cost of a small fleet. The M6 configuration is therefore the class where operational changes at the hub have the most effect, which is why the experiments are run with that class.

7.6. Experiments on the M6 configuration

The second stage takes the scenarios of the class whose cost disadvantage carries the largest removable operational component, which is the M6 class and subjects it to a series of experiments that vary only the settings the hub can actually control. The experiments are not fixed in advance, but follow from a diagnosis of the hub-side waiting. During each replication the number of barges waiting for the hub berth is sampled every simulated hour and folded onto the 168 hours of the week. The resulting mean queue profile shows when over the week the berth queue forms and which controllable setting drives it. This diagnosis shows two settings that the hub operates directly, which are the opening regime and the shuttle timetable. These define the operational experiments 1, 2 and 3.

Experiment 1 extends the opening hours of the hub with a Sunday window. Experiment 2 moves the shuttle arrival later on its three delivery days so that the morning queue can clear before the priority barge claims the berth. Experiment 3 opens the hub continuously, which is not a realistic proposal but a bounding case that shows the largest effect any opening-hours policy could reach. In each experiment, the fleet is recalibrated, because relieving congestion lowers the fleet needed to carry the same weekly flow. Experiment 1 is additionally run at the original baseline fleet. Experiment 1a holds the fleet fixed and separates the effect of the setting change from the effect of the fleet reduction. Experiments 1, 1a, 2, and 3 modify only the operational part of the hub and hold every cost parameter at its baseline value. The exact parameters per experiment are presented in Table 7.11.

Table 7.11: Definition of the experiments

Experiment	Change	Hub opening (h/week)	Deepsea block	Hub fleet	Direct fleet
Baseline	Mon-Fri 06:00-22:00, Sat 07:00-13:00	86	5×21 h	4 / 7	4 / 7
Exp. 1	Sunday opening 07:00-13:00	92	5×21 h	3 / 6	4 / 7
Exp. 1a	Sunday opening, baseline fleet	92	5×21 h	4 / 7	4 / 7
Exp. 2	Shuttle arrival moved to 14:00	86	5×21 h	3 / 6	4 / 7
Exp. 3	Continuous operation (bound)	168	5×21 h	3 / 6	4 / 7
Exp. 4	Congestion at the deepsea terminal	as above	5×27 h	as above	4 / 8

The sensitivity of the result to the shuttle fee is not tested as a separate experiment. The fee enters the cost model as a linear per-container term with no feedback on the simulated dynamics, so the physical outcome of a run does not depend on the fee. The break-even fee, which is the fee at which the hub cost per TEU equals the direct cost per TEU, can therefore be calculated directly from the saved output. It is compared against the tariff floor of 60 EUR per TEU, which is the lowest fee at which the shuttle still covers its own operation. This shows whether any fee exists that is low enough for the third-party barge operator and high enough for the shuttle to be sustainable.

The fourth experiment tests the sensitivity to congestion at the deepsea terminal. It lengthens the blocks of the deepsea vessels in the current-state simulation to five blocks of 27 hours per week, which also includes broken crane situations. This is the longest block at which the direct call can still deliver the same 540 containers per week as the hub. Where needed, the direct-call fleet is increased to reach this volume. The hub simulation is not changed.

Together the four experiments and the break-even fee show not only whether the hub helps, but for which operator profile it helps and under which operating and pricing conditions the cost gap could be closed.

7.6.1. Diagnosis of the berth queue

The experiments are based on a diagnosis of the berth queue over the weekly cycle. During each replication, the number of vessels waiting for the hub berth is sampled every simulated hour. The samples of all 30 replications are folded onto the 168 hours of the week. Figure 7.2 shows the resulting mean queue for the baseline settings and for experiment 1 on the long leg.

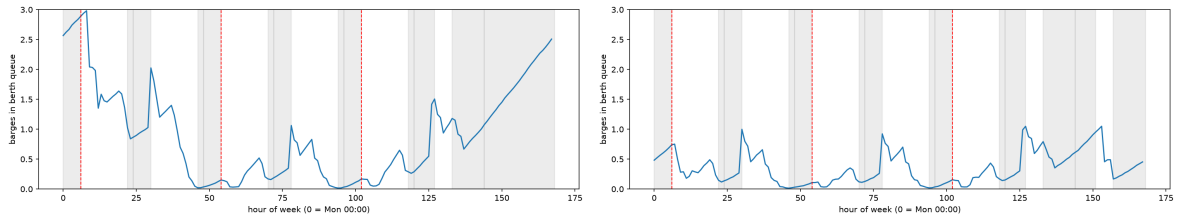

(a) Baseline regime (long leg)

(b) Experiment 1, Sunday opening (long leg)

Figure 7.2: Mean number of barges in the hub berth queue per hour of the week, M6 long leg, 30 replications.

The baseline scenario reveals a clear weekly pattern. The hub closes on Saturday at 13.00 and reopens on Monday at 6.00. The closed hours are marked as the grey bands. During this 41-hour closure the queue grows almost linearly to a peak of about 3 barges. The peak grows with the Monday shuttle arrival (red dashed lines), which claims the berth with priority at the moment of reopening. The queue then takes until Wednesday midday to return to zero barges. The overnight closures produce smaller peaks that disappear within a morning.

Two conclusions follow. First, the weekend closure is the dominant source of the berth queue. Barges arrive while the hub is closed and subsequently queue behind each other after reopening. A large part of the recorded berth-queue time is therefore manufactured by the closure regime rather than by a structural shortage of berth capacity. Second, the shuttle timetable only interacts with the queue on Monday morning. The Wednesday and Friday shuttles arrive at a nearly empty queue.

7.6.2. Results of the experiments

Table 7.12: The operational experiments against the direct-call situation

			Hub		Cost per TEU (EUR)			Excluding fees (EUR/TEU)		
Experiment	Leg	Fleet	Wait (h)	TT (h)	Hub	Direct	Gap	Hub	Direct	Gap
Baseline	short	4	16.13	58.47	155.58	79.79	$+75.79 \pm 0.36$	88.08	79.79	$+8.29$
Exp. 1	short	3	7.42	50.17	136.50	79.79	$+56.71 \pm 0.51$	69.00	79.79	-10.79
Exp. 1a	short	4	13.84	56.10	157.21	79.79	$+77.42 \pm 0.35$	89.71	79.79	$+9.92$
Exp. 2	short	3	11.30	52.75	137.22	79.79	$+57.43 \pm 0.36$	69.72	79.79	-10.07
Exp. 3	short	3	4.16	45.04	137.80	79.79	$+58.01 \pm 0.44$	70.30	79.79	-9.49
Baseline	long	7	19.00	114.23	232.71	158.12	$+74.59 \pm 1.05$	165.21	158.12	$+7.09$
Exp. 1	long	6	8.09	104.16	214.26	158.12	$+56.14 \pm 0.89$	146.76	158.12	-11.36
Exp. 1a	long	7	11.12	107.05	234.59	158.12	$+76.47 \pm 1.01$	167.09	158.12	$+8.97$
Exp. 2	long	6	11.93	106.56	216.05	158.12	$+57.93 \pm 0.69$	148.55	158.12	-9.57
Exp. 3	long	6	2.50	96.66	215.02	158.12	$+56.90 \pm 1.29$	147.52	158.12	-10.60

Table 7.12 shows the results of the three experiments against the direct call. The three experiments that recalibrate the fleet all land in the same place. Waiting at the hub over those same six runs varies far more widely, from 4.16 to 11.30 hours on the short leg and from 2.50 to 11.93 hours on the long leg. A spread of nine hours in waiting produces a spread of less than two euros per TEU in cost. Continuous operation removes almost all of the waiting and is no cheaper than six hours of Sunday opening. Experiment 1a explains why. It applies the Sunday regime at the original fleet of four and seven barges. Waiting falls by 2.29 and 7.88 hours and turnaround falls by 2.37 and 7.18 hours, yet the cost per TEU rises by 1.63 and 1.88 EUR and the gap widens rather than narrows. The reason is the supply ceiling. The hub receives roughly 540 containers per week whatever the barge does, so a

faster barge does not carry more. It simply calls more often, which results in an increase in the number of visits and a decrease in the number of containers transported per barge. The barge also spends a larger share of the run in motion, converting waiting hours charged at 94.50 EUR into sailing hours charged at 158.69 EUR. Under a fixed supply, speeding up the hub does not pay unless the fleet can shrink with it. The saving in experiments 1, 2 and 3 therefore comes from the barge that is released and not from the measure itself. All three free enough capacity to run one barge fewer and all three deliver the same saving because they all release exactly one barge. What separates them is how they get there. Sunday opening relieves the weekend closure, the retimed shuttle arrival reshapes the arrival pattern and continuous operation removes the closure entirely. However, the fleet is still what the operator pays for.

The last three columns isolate the effect that matters for the conclusion. With the fee removed from both sides, the baseline hub is 8.29 and 7.09 EUR per TEU more expensive to operate than a direct call. After any of the three experiments the sign reverses and the hub becomes the cheaper operation, by 9.49 to 11.36 EUR per TEU. The alternative operational cases for the hub are therefore sound, but what defeats it is the charge. The operator pays 67.50 EUR per TEU for the shuttle against a free single deepsea berthing and that difference of 67.50 EUR per TEU is larger than the entire operational advantage the hub can generate. Even with the lowest shuttle fee possible for MSC of 60 EUR per TEU, the gap would still be 48.64 to 50.51 EUR/TEU on the three alternative operational cases.

7.6.3. The congestion experiment

The previous experiments were focused on the operational changes of the hub to see if these different operational changes would change something in the cost gap between the hub and current situation. Despite a smaller cost gap in all three experiments, the hub still has a long way to go before being a promising alternative for the small barges compared to the direct deepsea call. The in-port time in the deepsea port is in this thesis compared to the AIS data, which presents an average in-port time of 18.59 hours on the M6 class. However, there are periods when cranes at the deepsea terminal are broken or under maintenance. When this occurs, the vessels with fixed windows get priority so they can be handled according to schedule. The result is extra waiting time for barges without fixed windows that can increase to 72 hours in some cases. This information is gathered from weekly updates done by MSC about the ECT terminals through personal communication. The literature confirms these extreme in-port times for barges without fixed windows (Nicolet et al., 2023; Vermaat, 2026). The next experiment is focused on the maximum block consisting of deepsea vessel blockage and broken cranes, in order to still be able to load the same amount of containers as in the hub situation (540 containers). The deepsea block is therefore set on 27 hours per block and the block appears 5 times per week. Due to the longer deepsea block, there are now 8 barges needed to transport the minimum of 540 containers per week. The short leg can use the same 4 barges. The hub simulation is untouched, because the shuttle loads inside its own fixed window and is therefore handled according to schedule.

Table 7.13: Results of the congestion experiment

			Turnaround (h)			Cost per TEU (EUR)		
Experiment	Leg	Direct call time (h)	Hub	Direct	Difference	Hub	Direct	Gap
Baseline	short	22.71	58.47	68.41	-9.94 ± 0.33	155.58	91.15	$+64.43 \pm 0.34$
Exp. 1	short	22.71	50.17	68.41	-18.24 ± 0.26	136.50	91.15	$+45.35 \pm 0.35$
Exp. 2	short	22.71	52.75	68.41	-15.66 ± 0.32	137.22	91.15	$+46.07 \pm 0.39$
Exp. 3	short	22.71	45.04	68.41	-23.36 ± 0.20	137.80	91.15	$+46.65 \pm 0.30$
Baseline	long	36.85	114.23	135.74	-21.51 ± 0.84	232.71	184.03	$+48.69 \pm 0.84$
Exp. 1	long	36.85	104.16	135.74	-31.58 ± 0.65	214.26	184.03	$+30.23 \pm 0.65$
Exp. 2	long	36.85	106.56	135.74	-29.18 ± 0.78	216.05	184.03	$+32.02 \pm 0.79$
Exp. 3	long	36.85	96.66	135.74	-39.08 ± 0.48	215.02	184.03	$+30.99 \pm 0.75$

As can be seen in Table 7.13, even with the longer block at the deepsea terminal, the direct call stays the cheaper routing in every row. The gap is 48.69 EUR per TEU at the unchanged hub on the long leg and falls to between 30.23 and 32.02 EUR per TEU in the three experiments. On the short leg the gap is 64.43 EUR per TEU and falls to between 45.35 and 46.65 EUR per TEU. While the hub is still

more expensive in this experiment, the time verdict reverses. In the current baseline scenario the hub is 3.11 hours slower than the direct call on the long leg. Under congestion the hub is faster in every row, by 21.51 hours at the unchanged hub and by 39.08 hours in Experiment 3 on the long leg. On the short leg the advantage runs from 9.94 to 23.36 hours. Every interval stays below 0.85 hours, so all of these differences are significant. The hub buys time under congestion because the shuttle keeps its window while the direct barge waits behind the priority vessel and experiences the congestion caused by the broken cranes.

The three hub experiments stay close together. The long gaps are 30.23, 32.02 and 30.99 EUR per TEU and the short gaps are 45.35, 46.07 and 46.65 EUR per TEU. Experiment 3 removes the most hub waiting and still ends slightly behind experiment 1 on the short profile. Congestion at the deepsea terminal is worth far more to the hub than any opening-hours policy at the hub itself. The congestion scenario therefore separates two verdicts that the baseline reports together. The operational verdict turns in favour of the hub. The hub becomes faster on turnaround and the fleet that is needed with the hub is smaller than the fleet needed for the deepsea terminal. However, the commercial verdict does not turn. The shuttle fee still exceeds the value the hub creates and the window that congestion opens stays below the tariff floor that MSC needs. The direct call time of 36.85 hours on the long profile remains far below the level at which the cost gap would close, which is why congestion narrows the gap without reversing it. Nevertheless, the difference in time has become quite large between the hub and direct call. Therefore, one could wonder if the price of the hub route is worth it in this case. Especially when considering the long leg. One limitation belongs with this result. The 27-hour block is not calibrated on the AIS sample. It represents a period with broken cranes or cranes under maintenance and not an average week.

7.6.4. Summary of the experiments

Four experiments were run on the M6 configuration. The first three change the way the hub is operated and hold every cost parameter at its baseline value. The fourth experiment increases the block at the deepsea terminal to see the effect on the direct call simulation.

The four experiments buy a lot of time but only part of the cost. Between the baseline and the continuous case the waiting time per visit decreases as well as the turnaround ratio. The hub becomes clearly faster than the direct call. However, the cost per TEU never falls below 136.50 EUR on the short leg and the gap never becomes smaller than 30.23 EUR per TEU. Time and cost do not respond to the same lever. The cost reduction that does appear in the experiments comes from the smaller fleet needed rather than from the opening hours. Experiment 1a applies the Sunday opening at the original fleet of four and seven barges. The waiting time still falls, but the cost per TEU rises. The 19.08 and 18.45 EUR per TEU saved in experiment 1 is therefore the value of the barge that the lower congestion allows the operator to remove, not the value of the extra opening hours themselves.

Another observation recovered from experiment 3 is that the operational improvement runs into a ceiling because the binding constraint moves. Under continuous operation the cost per TEU on the short leg is higher than under the Sunday opening. Beyond the Sunday opening the berth is no longer the bottleneck and the shuttle supply is, so further opening hours convert waiting into partially loaded sailing. One intervention is worth separating from the rest on practical grounds. Experiment 2 reaches the same effect of experiment 1 and reduces the fleet by the same one barge per leg, but it only moves the arrival time of a barge that already sails, so it carries no additional cost for the hub. The Sunday opening performs slightly better but extends the opening hours, which the hub has to pay for.

The last experiment changes the handling of the deepsea terminal in the direct call simulation by increasing the deepsea block. This last experiment is used to see what the effect of a maximal deepsea block on the congestion of the deepsea terminal is and if this maximum block could give the hub the advantage in this comparison. Even with the maximum block and the most favourable operational hub configuration the hub is still more expensive. However, the time advantage has increased a lot in favour of the hub with 9.94 hours on the short leg and 21.51 hours on the long leg. When a third-party barge operator is willing to pay a little extra in order to have a shorter turnaround time, the hub scenario could be a preferred option in this case.

8

Discussion

This chapter interprets the results of chapter 7. Section 8.1 explains why the hub loses on cost. Section 8.2 shows that the size of the fleet is the main lever for both time and cost. Section 8.3 sets out the operator profiles for which the hub could still become attractive. The chapter closes with the limitations and the directions for further research.

8.1. The shuttle fee decides the cost comparison

The hub is more expensive than the direct call in all six baseline scenarios. The gap runs from 63.85 EUR per TEU for M8 on the long leg to 84.07 EUR per TEU for M9 on the short leg. The shuttle fee of 67.50 EUR per TEU occupies a large share of that gap. A barge without a fixed window pays nothing for a call at the deepsea terminal. The hub therefore has no call fee to save. What the operator receives for the fee is time, since the hub removes the seven-hour round trip to the Maasvlakte and the congested call at the deepsea terminal.

The operation of the hub itself can be made cheaper than the direct call. With the fee left out, the baseline M6 hub is 8.29 and 7.09 EUR per TEU more expensive than the direct call on the short and the long leg. After any of the three operational experiments this turns around and the hub is the cheaper operation by 9.49 to 11.36 EUR per TEU. The fee then decides the outcome. Even at 60 EUR per TEU, the lowest fee at which the shuttle still covers its own operation, the gap stays between 48.64 and 50.51 EUR per TEU. The level of the fee is also high compared to the literature. Konings et al. (2013) calculate how much the extra handling step in a hub may cost before the hub network stops being cheaper than direct calls. For the Eem-Waalhaven location, which is the relation the shuttle serves today, this margin lies between roughly 1 and 19 EUR per TEU. The fee charged today is more than three times the top of that range. A side note is the inflation that is not taken into account. The margins of Konings et al. (2013) are prices from 2013, which differ from the prices with the same value handled in 2026. Bedoya-Maya et al. (2025) also find negative cost savings for consolidation at seaports, because the handling costs dominate. This thesis reaches the same result with a different method and with the tariff that is charged in practice.

On time the picture is more mixed. In the baseline the hub is faster only for the M8 class, by 2.68 hours on the short leg and 2.29 hours on the long leg. In every scenario the hub waits 4.14 to 9.91 hours longer per visit than the direct call, because it replaces one berth queue with three waiting mechanisms. These are the berth queue, the closing hours of the hub and supply starvation. The hub can win on time, as Section 8.2 shows, but the fee keeps it from winning on cost.

8.2. Fleet size is the main lever for time and cost

The results show that the size of the fleet decides a large part of both the turnaround time and the cost. The hub and the deepsea terminal each have a single berth for the inland barges. Every barge that is added to the fleet joins the same queue. The operator pays for the extra barge and for the hours it waits. The baseline shows this at the hub. The M6 class needs four to seven barges to move the

weekly flow of about 540 containers. As a result the berth queue is 72 to 75 percent of the M6 waiting time. The M8 class moves the same flow with two or three barges and has a much shorter berth queue.

The operational experiments show the same mechanism. The Sunday opening, the later shuttle arrival and the continuous opening each free enough capacity to run one barge fewer. The gap then falls from 75.79 to between 56.71 and 58.01 EUR per TEU on the short leg, and from 74.59 to between 56.14 and 57.93 EUR per TEU on the long leg. Experiment 1a shows where this saving comes from. It keeps the original fleet with the Sunday opening. The waiting time falls, but the cost per TEU rises by 1.63 and 1.88 EUR. The hub receives about 540 containers per week whatever the barges do, so a faster barge does not carry more. It only calls more often with a smaller load. The saving therefore comes from the barge that is removed and not from the extra opening hours.

The congestion experiment shows the effect of the fleet most clearly. In this experiment the berth at the deepsea terminal is blocked by deepsea vessels and broken cranes five times per week for 27 hours. The direct call still has to deliver the 540 containers per week that the hub moves. On the long leg this is only possible when the direct-call fleet grows from seven to eight barges. The extra barge joins the same single berth, so the queue at the deepsea terminal grows. The direct call now takes 36.85 hours at the terminal on the long leg and 22.71 hours on the short leg, against about 12 hours in the baseline. The hub side does not change, because the shuttle loads within its own fixed window and is therefore handled as scheduled. On the long leg the hub now needs fewer barges than the direct call, seven at the baseline hub and six after the operational experiments against eight. The hub is also faster in every case. It saves 21.51 to 39.08 hours per round trip on the long leg and 9.94 to 23.36 hours on the short leg. The cost gap narrows, but does not close. On the long leg it falls to 48.69 EUR per TEU at the baseline hub and to 30.23 to 32.02 EUR per TEU after the operational experiments. On the short leg it falls to 64.43 and to 45.35 to 46.65 EUR per TEU. The hub is still more expensive. However, on the long leg an operator now pays about 30 EUR per TEU extra for a round trip that is more than a day shorter. A third-party barge operator that values a shorter and more predictable turnaround time may accept that price. This is the case in which the hub comes closest to being an attractive option. The 27-hour block is not calibrated on the AIS data. This result should be read as a situation that represents a period with broken cranes or cranes under maintenance, which is not the average week that can be found in the AIS data.

Together these results give one practical lesson. A change that lets the hub operator remove a barge, or that forces the direct-call operator to add one, moves the comparison toward the hub. Congestion at the deepsea terminal is worth more to the hub than any opening-hours policy at the hub itself.

8.3. Operator profiles for which the hub could be attractive

For the scenarios tested in this thesis the hub-and-spoke system is not profitable for the third-party barge operator. The results do point at profiles for which the answer could be different. All of them work through the same terms. The fee per TEU stays the same. The hub becomes more attractive when the time it saves per call grows, when that time is spread over fewer containers, or when the fee itself falls. Only the first profile below has been simulated. The others follow from the structure of the cost comparison and are identified rather than quantified.

8.3.1. Barges with a long waiting time in port

Barges without a fixed window that wait long at the deepsea terminal gain the most from the hub. The congestion experiment shows how large that gain in time can be. The literature reports in-port times far above the values in the AIS data used here, with an average of 52 hours for barges from the upper Rhine (Vermaat, 2026). Congestion alone cannot close the gap in this model, however. Above a block of about 27 hours a single berth can no longer move 540 containers per week at any fleet size.

8.3.2. Barges with more than one call

In the AIS data, 140 of the 142 cycles contain a call at a single terminal. Roux et al. (2023) report that small barges make six to eight calls at different deepsea terminals to fill up. Each extra call adds sailing between terminals and another berth queue. The hub replaces all of these with one call at BCW. This profile is the one most likely to change the answer. It would also bring more containers to the hub, so the shuttle service would have to grow with it.

8.3.3. Smaller barges and smaller calls

An hour saved is worth more per TEU for a small barge, because the barge carries fewer containers. A class below M6, if such barges are still in use, would therefore need fewer saved hours to cover the fee. This agrees with Nicolet et al. (2023), who find that a consolidation terminal suits barges that carry less than 60 TEU. A similar argument holds for operators with a small call size. ECT applies a minimum call size of 10 moves at DDE and 20 moves at DDN. For an operator that loads fewer containers per call the direct call is not a real option, while the hub has no minimum call size in this model. Two conditions apply. First, a smaller call multiplies the time difference per call in whichever direction it points, so it only helps where the hub is the faster route. Second, a smaller barge needs more calls to move the same 540 containers. This adds to the berth queue at the hub, which is exactly the waiting that has to fall.

8.3.4. A cheaper shuttle service

The fee is the most direct lever. It is also the one that fails by the widest margin. For the M6 baseline, the break-even fee is below zero, so even a free shuttle would leave the hub behind. After the operational experiments it rises to 9.49 to 11.36 EUR per TEU. In the most favourable case tested, the congested terminal with the Sunday opening on the long leg, it reaches 37.27 EUR per TEU. That is still 22.73 EUR per TEU below the 60 EUR per TEU at which the shuttle covers its own operation. A cheaper shuttle therefore needs a cheaper operation and not only a lower price. Examples are a larger shuttle, a higher use of the existing one, or direct transhipment with push barges as proposed by Konings et al. (2013).

8.3.5. The value of reliability

The cost layer only counts the hours an operator pays for. The value of time and the value of reliability used by Bedoya-Maya et al. (2025) are left out to avoid counting the same delay twice. The shuttle keeps its fixed window, so the hub route is less exposed to congestion at the deepsea terminal. Under congestion it is up to 39 hours faster per round trip. If an operator puts a money value on a shorter and more reliable turnaround, part of the remaining gap of about 30 EUR per TEU could be covered. This thesis does not test that.

8.4. Limitations

The AIS data does not show which barges hold a fixed window at the deepsea terminal. The observed in-port times are therefore a mix of barges with and without a fixed window, while the simulated direct call is always a barge without one. The data on the import flow is also limited. The observed in-port time includes both loading and unloading, while the model only contains the import flow. This is a likely reason why the simulated in-port time of the M6 class is 6.39 hours shorter than observed. Finally, the data covers about 40 days. Congestion at the deepsea terminals changes over the year and the literature reports in-port times several times larger than the mean found here.

When focusing on the model, the barges in the model sail continuously. Every barge starts a new cycle as soon as it returns. A real operator may keep a barge out of service at times, which would raise the cost per TEU of both routings. Furthermore, the throughput of the hub is set by the number of shuttle trips per week. All containers that the shuttle delivers are collected by the inland barges and an inland barge only departs when it is full or when the hub is empty. In practice part of the volume leaves the hub by truck, which this thesis does not include. Both choices tie the call size of the inland barges directly to the shuttle supply. In addition, the internal process of the hub is not modelled. The containers are moved to the inland barge as one stock at the crane rate, without stacking or yard moves. The time at the hub is therefore a lower bound. In the simulation of the direct call each inland barge visits only one deepsea terminal. This follows the AIS data, but in practice a barge can visit many more terminals. That means more calls and more waiting. Relaxing this assumption would favour the hub. However, the shuttle barge would also need to visit more deepsea terminals, which would increase the shuttle fee.

Finally, the barge capacities of 90, 192 and 256 TEU are class values from the literature and not the precise capacities of the observed barges. The factor of 1.6 TEU per container from Van Koningsveld et al. (2021) converts the shuttle charge of 108 EUR per container into 67.50 EUR per TEU. For a factor

between 1.4 and 1.8 the fee moves between about 73 and 63 EUR per TEU. This is small compared to the gaps of 63.85 to 84.07 EUR per TEU.

8.5. Further research

The most promising direction is to extend the model to more deepsea terminals, served by a shuttle that also sails between those terminals. Together with barges that make several deepsea calls per cycle and carry MSC cargo, this would show whether the advantage per barge turns positive when only cost is counted. It requires the operating cost of the MSC shuttle, so that a new fee can be calculated that includes the added terminals. It also requires data that contains such multi-call cycles. Furthermore, the system could be tested for barges smaller than the M6 class and for operators with call sizes below the minimum of the deepsea terminal. The lower the number of containers per barge, the greater the effect of the waiting and sailing time on the cost per TEU. However, it is important to prevent the barge congestion at the deepsea terminals from moving to the hub. This could be the case with smaller barges or call sizes and the same throughput of containers.

The simulation model can be improved in different ways. The first is low water. Low water periods restrict the draught of a barge and therefore the number of containers it can load. If these periods force operators to lower their load to keep sailing, the call size falls, which favours the hub where it is the faster route. The second is the export flow, including empty containers from the inland terminals. This makes the model complete and gives a more realistic loading and unloading time. Another direction is adding the internal hub process. In the model the hub is a station where nothing happens, while its internal operations could change the outcome of the system. Furthermore, adding the value of reliability to the cost comparison would show whether the time saving under congestion is worth the remaining gap. Finally, better data would strengthen the comparison. Data that separates barges with and without a fixed window would give a cleaner reference for the direct call. A period longer than 40 days would show how representative the observed in-port times are.

9

Conclusion

The objective of this research was to quantify the operator cost and turnaround time of routing third-party inland barges through a hub-and-spoke system relative to direct deepsea calls, and to establish for which barge profiles this reduces cost and turnaround time. A discrete-event simulation was built of both routings. In the hub routing the shuttle service carries the containers between the Maasvlakte and the Barge Center Waalhaven. In the direct routing the inland barge calls at the deepsea terminal itself. Both routings use the same fleet, the same inland-leg samplers and the same thirty seeds, so every difference is taken as a paired difference per seed. The AIS data is used to set the parameters and to validate the direct-call simulation. Six scenarios were run, formed by three vessel classes and two inland-leg profiles. After that, four experiments were carried out on the M6 class, which has the largest waiting cost that the hub can remove.

9.1. Answers to the sub-questions

Sub-question 1

The cleaned record contains 142 in-port cycles at the two ECT terminals and 282 usable inland trips. In 140 of the 142 cycles the barge calls at a single terminal, which makes one call per cycle the empirically defensible baseline. The observed in-port time is 18.59 hours for the M6 class, 14.00 hours for M8 and 15.84 hours for M9. The data does not show which barges hold a fixed window. This is a likely reason why the smallest class spends the longest time in port. The inland leg splits at a median of 55.7 hours into a short and a long profile of 141 trips each. The Weibull distribution reproduces the level and the spread of both.

Sub-question 2

The direct call is the cheaper routing for all three classes. The gap runs from 63.85 EUR per TEU for M8 on the long leg to 84.07 EUR per TEU for M9 on the short leg. It is not ordered by vessel size, since the M6 class sits in between at 75.79 and 74.59 EUR per TEU. On time the hub is faster only for the M8 class, by 2.68 hours on the short leg and 2.29 hours on the long leg. The M6 class is 0.81 and 3.11 hours slower. The M9 class is 3.14 hours slower on the long leg, while on the short leg the difference is not significant.

Furthermore, the waiting at the hub arises through different mechanisms per class. For M6 the berth queue is the largest part, because four to seven barges compete for a single berth. For M8 the closing hours of the hub are the largest part, because two or three barges arrive outside the opening hours relatively often. For M9 the berth queue dominates again, because each call loads for 9.5 to 11.1 hours. In every scenario the hub waits 4.14 to 9.91 hours longer per visit than the direct call.

Sub-question 3

The hub treats the two legs almost identically. The paired difference in waiting time at the hub between the long and the short leg is at most 2.9 hours, so a long-leg operator does not queue meaningfully

longer than a short-leg one, and the difference in turnaround is explained by the extra sailing itself. What the leg does change is the cost per TEU, which is 77.13 EUR higher on the long leg for M6 and 32.33 EUR higher for M9, because the additional sailing hours are a fixed quantity per cycle that larger vessels spread over more containers. The leg does matter under congestion at the deepsea terminal. The long leg needs a larger fleet, so a change in berth capacity more easily adds or removes a whole barge. In experiment 4 the direct call on the long leg needs eight barges instead of seven to move 540 containers per week, while the short leg keeps four. As a result the hub gains more on the long leg, with a gap of 30.23 to 32.02 EUR per TEU against 45.35 to 46.65 on the short leg, and a time saving of 21.51 to 39.08 hours against 9.94 to 23.36 hours.

Sub-question 4

The handling capacity of the hub is not the binding constraint. The berth and the cranes are never exhausted and the holding capacity of 2,700 TEU is reached only in the extreme condition test. What binds is the supply of the shuttle. The shuttle carries 288 TEU, which is 180 containers at 1.6 TEU per container, and it runs three cycles per week. The ceiling of about 540 containers per week is therefore a property of the shuttle and not of the hub. As a result the call size is a system outcome and not a vessel property. It is the weekly supply divided by the number of visits of the fleet. This is why the M9 barge leaves the hub at 66 percent of its capacity on the short leg, and why M9 has no cost advantage over M8 there, at 121.81 against 121.66 EUR per TEU.

The direct call has no supply ceiling. With the same fleet it delivers 1,385 to 8,236 more containers over the run in five of the six scenarios. It does have a berth ceiling. When the deepsea berth is blocked five times per week for 27 hours, the direct call needs eight barges instead of seven on the long leg to move 540 containers per week. With a longer block a single berth cannot move this volume at any fleet size. The hub runs out of supply and the direct call runs out of berth.

Sub-question 5

For the third-party barge operator the hub costs 63.85 to 84.07 EUR per TEU more than the direct call. The cause is the fee. The operator pays 108 EUR per container for the shuttle, which is 67.50 EUR per TEU, while a barge without a fixed window pays nothing for a call at the deepsea terminal.

The operation of the hub can be made cheaper than the direct call. With the fee left out, the Sunday opening, the later shuttle arrival and the continuous opening each make the M6 hub 9.49 to 11.36 EUR per TEU cheaper than the direct call. That is also the highest fee the operator could pay at those settings. In the most favourable case tested, a congested deepsea terminal together with the Sunday opening on the long leg, the break-even fee reaches 37.27 EUR per TEU. All of these lie below the 60 EUR per TEU at which the shuttle covers its own operation. No fee therefore exists that is low enough for the operator and high enough for MSC, even with an extra congestion element at the deepsea terminal.

The congestion experiment is the condition that comes closest. The direct call then takes 36.85 hours at the terminal on the long leg and 22.71 hours on the short leg. The gap narrows to 48.69 EUR per TEU on the long leg at the baseline hub and to 30.23 to 32.02 EUR per TEU after the operational experiments. On the short leg it narrows to 64.43 and to 45.35 to 46.65 EUR per TEU. The hub stays more expensive, but it becomes faster in every case, by 21.51 to 39.08 hours per round trip on the long leg and by 9.94 to 23.36 hours on the short leg.

9.2. Answer to the main research question

Routing a third-party inland barge through the hub changes the turnaround time by between minus 2.68 and plus 3.14 hours and raises the transport cost by 63.85 to 84.07 EUR per TEU. For every operator profile tested the direct call is the cheaper routing. No operational setting at the hub and no tested state of the deepsea terminal changes that.

The main reason is the shuttle fee. The operator pays 67.50 EUR per TEU for the shuttle against a free call at the deepsea terminal, and the hub does not save enough time to cover that. The operation of the hub itself is not the problem. After the operational experiments the hub is cheaper to operate than the direct call when the fee is left out. The second finding is that the size of the fleet is the main lever

for both time and cost. The hub and the deepsea terminal each have a single berth, so every extra barge joins the same queue. The operational experiments lower the gap from about 75 EUR per TEU to between 56.14 and 58.01 EUR per TEU, because they free enough capacity to remove one barge. Experiment 1a shows that the same opening hours at the original fleet raise the cost per TEU instead of lowering it. Under congestion the same mechanism works against the direct call. It needs an extra barge on the long leg to keep 540 containers per week moving, the queue at the deepsea terminal grows, and the hub becomes faster in every case.

The hub is therefore a promising alternative on time, but not on cost. It suits operators without a fixed window that face long waiting times at the deepsea terminal. Under congestion an operator on the long leg pays about 30 EUR per TEU extra for a round trip that is more than a day shorter. An operator that values a short and predictable turnaround may accept that price. The congestion experiment represents a period with broken cranes or cranes under maintenance and not an average week that is calibrated with the AIS-data. On cost the hub is not a promising alternative for any profile tested. The profiles for which this could change are barges that make several deepsea calls per cycle, barges smaller than the M6 class, operators with call sizes below the minimum of the deepsea terminal, and a cheaper shuttle operation. These are identified in the discussion but not tested.

In summary, the hub-and-spoke system as it runs today is not profitable for third-party inland barge operators in the scenarios tested. The fee is too high for the time the hub saves. When the deepsea terminal is congested the hub becomes clearly faster, and the question for the operator then becomes whether that time is worth the extra cost.

9.3. Recommendations

For MSC the structure of the level of the shuttle fee and not the operation of the hub is the triggering factor. In every tested case the break-even fee lies below the level at which the shuttle covers its own operation. A lower fee therefore requires a cheaper shuttle operation, for example a larger shuttle, a higher use of the existing one, or direct transhipment with push barges. MSC can also focus on the operators for whom the hub offers the most, which are barges without a fixed window that wait long at the deepsea terminal. Of the operational settings, the shuttle retiming of experiment 2 can be adopted directly. Like the Sunday opening it allows one barge fewer per leg, but it only moves the arrival time of a shuttle that already sails. It therefore carries no extra cost for the hub. The Sunday opening performs slightly better but extends the opening hours. The hub pays for those hours while the saving goes to the operator, so such a change needs an arrangement between the two parties. For third-party barge operators the hub is worth considering when the waiting at the deepsea terminal is long and a short and predictable turnaround has value. The decision should be made at fleet level, because the saving comes from the barges the operator can remove or does not have to add. The hub cannot relieve congestion at the deepsea terminals if the operators it is meant to attract have no cost reason to use it. The concept therefore needs either a cheaper shuttle operation or an operator profile with more to gain.

9.4. Contribution

The feasibility of a consolidation hub for inland barges has so far been assessed qualitatively or with modelled handling costs. This thesis provides a quantitative treatment for a hub that operates partly and charges a real tariff. It uses vessel-level AIS data and compares both routings as paired simulations. It quantifies the price of the extra handling step at 67.50 EUR per TEU against a free direct call for barges without a fixed window.

Four results generalise beyond the case study. First, a fixed window for the shuttle is a necessary but not a sufficient condition for a consolidation hub. This thesis shows that it gives the hub its time advantage under congestion, but that the fee also has to stay below the value of that time. Second, the throughput of such a system is set by the shuttle and not by the hub, so the call size of the operators is a system outcome and not a vessel property. A larger vessel class then gains little, as the M9 class shows on the short leg. Third, the fleet size is the main lever for both time and cost. The hub gains when it lets the operator remove a barge and when congestion forces the direct-call operator to add one. Fourth, the benefit of the hub per TEU increases as the call size decreases, provided that the

hub is the faster route. This points toward the smaller vessel classes rather than the larger ones. That direction runs against the common position that larger inland vessels are the answer for hinterland container transport, without contradicting the economies of scale on which that position rests.

References

Bedoya-Maya, F., Shobayo, P., Nicolet, A., Christopoulou, E., Majoor, I., van Hassel, E., & Vanelslander, T. (2025). Cargo consolidation in port-hinterland container transport: A spatial economic assessment for inland waterways. *Research in Transportation Business and Management*, *59*.

Bu, F., & Nachtmann, H. (2023). Literature review and comparative analysis of inland waterways transport: "Container on Barge". *Maritime Economics and Logistics*, *25*(1), 140–173.

Caris, A., Macharis, C., & Janssens, G. K. (2011). Network analysis of container barge transport in the port of Antwerp by means of simulation. *Journal of Transport Geography*, *19*(1), 125–133.

European Commission. (2021, June). *Communication from the commission to the European Parliament, the council, the European economic, and social committee and the committee of the regions* (tech. rep.). European Commission. Brussels.

Friedhoff, B., Ramne, B., Alias, C., van Hassel, E., Martens, S. E., Felde, J. z., & Samuel, L. (2023). Elevating the logistics resilience of the Rhine-Alpine Corridor with the help of innovative vessel and cargo handling concepts. *Transportation Research Procedia*, *72*, 1998–2005.

Glasserman, P., & Yao, D. D. (1992). Some Guidelines and Guarantees for Common Random Numbers. *Management Science*, *38*(6), 884–908.

Hu, Q., Wiegmans, B., Corman, F., & Lodewijks, G. (2019). Critical literature review into planning of inter-terminal transport: In port areas and the hinterland. *Journal of Advanced Transportation*, *2019*.

Hutchison Ports ECT Rotterdam. (2026, January). *Handling of barges ECT Rotterdam* (tech. rep.). https://www.nextlogic.nl/zowerkthet/barge-operator/.

Janic, M., Reggiani, A., & Nijkamp, P. (1999). Sustainability of the European freight transport system: evaluation of innovative bundling networks. *Transportation Planning and Technology*, *23*(2), 129–156.

Konings, R. (2003). Network Design for Intermodal Barge Transport. *Transportation Research Record: Journal of the Transportation Research Board*, *1820*(1), 17–25.

Konings, R. (2006). *Hub-and-Spoke Networks in Container-on-barge Transport* (tech. rep.).

Konings, R., Kreutzberger, E., & Maraš, V. (2013). Major considerations in developing a hub-and-spoke network to improve the cost performance of container barge transport in the hinterland: The case of the port of rotterdam. *Journal of Transport Geography*, *29*, 63–73.

Law, A. M. (2020). Statistical Analysis of Simulation Output Data: The Practical State of the Art. *2020 Winter Simulation Conference (WSC)*, 1117–1127.

Lilliefors, H. W. (1967). On the Kolmogorov-Smirnov Test for Normality with Mean and Variance Unknown. *Journal of the American Statistical Association*, *62*(318), 399–402.

Ministerie van Infrastructuur en Waterstaat. (2026). Bedrijfseconomische kostenkengetallen voor het goederenvervoer, update 2026. https://www.kimnet.nl/documenten/2026/05/28/bedrijfseconomische-kostenkengetallen-voor-het-goederenvervoer-update-2026

Nicolet, A., Shobayo, P., van Hassel, E., & Atasoy, B. (2023). An assessment methodology for a modular terminal concept for container barging in seaports. *Case Studies on Transport Policy*, *14*.

Oganesian, V., Sys, C., Vanelslander, T., & van Hassel, E. (2021). *Container barge (un)reliability in seaports: a company case study at the port of Antwerp* (tech. rep. No. 6).

Otten, M., Scholten, P., Tol, E., van de Lande, P., Verbeek, M., & Wagter, H. (2020, August). *Outlook hinterland and continental freight 2020* (tech. rep.). Topsector Logistiek. https://ce.nl/publicaties/outlook-hinterland-and-continental-freight-2020/

Port of Rotterdam. (2025, February). Cargo throughput in the port of Rotterdam slightly decreased in 2024.

Port of Rotterdam. (n.d.). Container Exchange Route (CER). https://www.portofrotterdam.com/en/building-port/ongoing-projects/container-exchange-route-cer

Robinson, S. (1997). Simulation model verification and validation: increasing the users' confidence. *Proceedings of the 1997 Winter Simulation Conference*, 53–59.

Robinson, S. (2007). A statistical process control approach to selecting a warm-up period for a discrete-event simulation. *European Journal of Operational Research*, *176*(1), 332–346.

Roux, L., Krämer, I., & Konings, R. (2023). Handling of Inland Vessels in Seaports – Necessary Actions and Additional Options to Support Container Transport on Inland Waterways. In Y. Li, P. Rigo, G. Zhao, Y. Hu, & F. E. Lefler (Eds.), *Proceedings of pianc smart rivers 2022* (PIANC 2022, pp. 1388–1397, Vol. 264). Springer.

Sargent, R. G. (2011). Verification and validation of simulation models. *Proceedings of the 2011 Winter Simulation Conference*, 183–198.

Shenkar, N., & Rosen, D. (2018). How has the invention of the shipping container influenced marine bioinvasion? *Management of Biological Invasions*, *9*(3), 187–194.

Shobayo, P., & van Hassel, E. (2019). Container barge congestion and handling in large seaports: a theoretical agent-based modeling approach. *Journal of Shipping and Trade*, *4*(1).

Shobayo, P., Yang, C., Bedoya-Maya, F., Nicolet, A., van Hassel, E., Atasoy, B., & Vanelslander, T. (2026). Improving a cargo consolidation strategy for port-hinterland transport via inland waterways. *Journal of Transport Geography*, *130*.

van der Horst, M. (2016). *Coordination in hinterland chains: An institutional analysis of port-related transport* [Doctoral dissertation, TU Delft].

Van Es, T. (2019, November). *Reducing container barge in-port times: A Simulation Study* (tech. rep.). TU Delft. Delft. http://repository.tudelft.nl/.

Van Koningsveld, M., Verheij, H. J., Taneja, P., & De Vriend, H. J. (2021). *Ports and Waterways: Navigating the changing world* (tech. rep.).

Veldman, J. (2019). *Optimal vessel class for container transport from hinterland to seaport, costs comparison between Class IV and Class Va.* (tech. rep.).

Vermaat, T. (2026). *Development of a feasibility assessment framework for consolidation hubs in hinterland container barge transport* (tech. rep.). Delft University of Technology. https://resolver.tudelft.nl/uuid:7b0ab80e-6afb-43f3-b544-94ff6d940a96

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

Zhang, Y., Atasoy, B., Souravlias, D., & Negenborn, R. R. (2020). Pickup and Delivery Problem with Transshipment for Inland Waterway Transport. In E. Lalla-Ruiz, S. Voß, & M. Mes (Eds.), *Computational logistics* (pp. 18–35, Vol. 12433). Springer.

Zheng, J., & Yang, D. (2016). Hub-and-spoke network design for container shipping along the Yangtze River. *Journal of Transport Geography*, *55*, 51–57.

Zhou, S., Ji, B., Song, Y., Yu, S. S., Zhang, D., & Van Woensel, T. (2023). Hub-and-spoke network design for container shipping in inland waterways. *Expert Systems with Applications*, *223*.

A

Scientific paper

This appendix contains the scientific paper based on this thesis. The paper has its own page numbers and its own reference list.

Hub routing as an alternative to direct calls for third-party barge profiles

N.L. Linssen

TU Delft, Master student, Transport, Infrastructure and Logistics, Faculty of Civil Engineering and Geosciences (CEG)

Inland barges without a fixed window lose many hours in the queue at deepsea terminals. A consolidation hub with a shuttle service is often proposed as a remedy. This paper asks whether a third-party barge operator has a reason based on cost and time to use a hub that already runs and charges a real tariff. The direct call and the hub-and-spoke routings are modelled with discrete-event simulation (DES) for the shuttle service between the Maasvlakte and the Barge Center Waalhaven (BCW) in Rotterdam. Automatic Identification System (AIS) data of inland barges sets the inland-leg distributions and validates the direct call. Six operator profiles are formed by three vessel classes and two inland-leg profiles. The direct call is cheaper in all six profiles, by 63.85 to 84.07 EUR per twenty-foot equivalent unit (TEU). The shuttle fee of 67.50 EUR per TEU explains the largest part of this gap, while a barge without a fixed window pays nothing for its deepsea call. On time the hub wins only for the M8 class, by 2.29 to 2.68 hours per round trip. When the deepsea terminal is congested the hub becomes faster for the M6 class by up to 39.08 hours, but it stays more expensive by at least 30.23 EUR per TEU. The size of the operator fleet is the main lever for both time and cost.

Keywords inland waterway transport, hub-and-spoke, container barge, hub, discrete-event simulation, Port of Rotterdam

I. Introduction

The container is at the core of the global transport system and the container throughput of ports is very large [1]. The Port of Rotterdam handled 13.8 million twenty-foot equivalent unit (TEU) in 2024 [2]. From the deepsea terminals, the containers can be moved into the hinterland by truck, train and barge. Inland waterway transport with barges carried about 3 million TEU to and from Rotterdam in 2020, which is a market share of 32 percent [3]. The large market share is a result of the low infrastructure and external costs per ton-kilometer compared to rail and road [4]. The European Commission aims to increase transport by inland waterways and short sea shipping by 25 percent by 2030 and by 50 percent by 2050 [5].

However, the sector faces a persistent problem at the deepsea terminals. Barges compete for berth and crane capacity with deepsea vessels and feeders, and the larger vessels get priority [6] due to the large container throughput and contractual relationships with the deepsea terminals. Many barge operators have no contract with the terminal, so they get a berth only when there is space [7]. Barges often have small call sizes at six to eight terminals per port visit, so a delay at one terminal spreads easily to the next [3]. This results in long in-port times. Vermaat [8] reports an average of 52 hours in-port time for barges from the upper Rhine and Nicolet et al. [9] mention waiting times of up to 120 hours in Rotterdam.

A hub-and-spoke network is a common proposal to relieve this congestion [10]. Inland barges drop and collect their containers at a consolidation hub instead of sailing all the way to the deepsea terminal. Shuttle barges with a fixed window at the deepsea terminals carry the bundled flow between the hub and the terminals.

The inland barge then avoids the queue at the deepsea terminal. The price is an extra node with an extra handling step [11]. In Rotterdam Mediterranean Shipping Company (MSC) runs a part of this hub-and-spoke model with a shuttle between the Maasvlakte terminals and the Barge Center Waalhaven (BCW) and charges a fee per container for it.

Earlier studies evaluated hub-and-spoke networks as designs at network level, with a modelled cost for the extra handling step [9, 10]. In this research that cost is a tariff the third-party barge operator pays. The operator pays a fee per container for the shuttle leg and in return avoids the call at the deepsea terminal. Whether that exchange pays off depends on how many containers the operator moves per call, on the hourly cost of the barge and on how much congestion the direct call really causes. These are properties of the operator and not of the hub. No study was found that quantifies this choice from the side of the third-party barge operator.

This paper addresses that gap. It asks to what extent routing a third-party inland barge through a hub-and-spoke system changes its turnaround time and its transport cost per TEU, and for which operator profiles the hub is a promising alternative. Both routings are modelled with discrete-event simulation. AIS data of individual barges sets the inland-leg distributions and validates the direct call. The two routings share the same fleet and the same random numbers, so every result is a paired difference. A cost layer converts time into cost per TEU with published cost rates per vessel class and the tariff charged in practice.

The paper makes three contributions. First, it gives a quantitative operator-level comparison for a hub that exists and charges a real tariff. Second, it decomposes the cost gap into a sailing, a waiting and a fee term, which shows where the gap comes from. Third, it shows that the throughput of such a system is set by the shuttle and that the size of the operator fleet is the main lever for both time and cost.

Section II reviews the literature and states the research gap. Section III describes the method. Section IV presents the case, the data and the experiments. Section V reports the results and Section VI discusses them. Section VII concludes.

II. Literature review

A. Barge congestion and consolidation

Barge handling at deepsea terminals has several structural disadvantages [12]. The cranes are sized for ocean-going vessels and lose productivity on small barges. Deepsea vessels get priority and regularly delay barges that hold no fixed window [6]. A barge also has to visit several terminals to collect a full load [3]. Shobayo and van Hassel [13] used an agent-based model of barge handling in large seaports and concluded that more dedicated barge infrastructure is needed. Who pays for that infrastructure remains an open question. Consolidation is a second line of research. Bedoya-Maya et al. [14] assessed cargo consolidation in port-hinterland container transport and found lower average costs and emissions per container. Consolidation at seaports or at single inland terminals, however, gave negative savings in all their scenarios, because handling and logistics costs dominated. Roux et al. [3] describe the bundling of flows between seaports and inland ports as a way to raise call sizes and to reduce the number of terminal visits. They state the core trade-off directly. The benefit of bundling at a hub must exceed the cost of the extra transhipment. They, however, do not quantify it. Hu et al. [15] reviewed inter-terminal transport and concluded that quantitative methods are needed that reflect the financial interests of the different actors.

Konings et al. [10] developed a cost model for a hub-and-spoke barge network in the hinterland of Rotterdam. The model reports the margin for handling at the hub. This is the largest amount the exchange at the hub may cost before the network stops being cheaper than direct calls. For the Eem-Waalhaven location, about 40 km from the Maasvlakte, the margin lies between roughly 1 and 19 EUR per TEU. It narrows as the hinterland vessel grows larger. The waiting times in that model are inputs. The present situation has nine terminal calls with about one hour of waiting each and the hub situation has a single call with one hour of waiting. The margin has therefore never been tested against waiting that emerges from vessels competing for

a real berth. Caris et al. [16] used discrete-event simulation to compare four hub-and-spoke scenarios in the port of Antwerp on waiting time and utilisation. They did not address the cost for the barge services. Nicolet et al. [9] proposed a floating consolidation hub, the Modular Mobile Terminal, and found it most suitable for barges that carry less than 60 TEU. Their assessment rests on regional container flows. They recommend a vessel-level analysis so that waiting at terminals and hubs can be represented in more detail. Network design studies treat the hub location and the fleet as decision variables [17–19]. Times and demand are deterministic there, and cost is the technical cost of running the system rather than a price someone is charged. Vermaat [8] built a qualitative feasibility framework for consolidation hubs. That study names fixed windows for the shuttle as a key requirement and recommends adding quantitative data.

Table 1 positions this paper against the most relevant studies. Hub-and-spoke barge networks have been evaluated as designs, with modelled handling costs and with waiting as an input or as an aggregate. It is not known for which kind of barges and fleets the hub-and-spoke model is a promising alternative. This paper gives answers to that gap with a paired simulation in which waiting emerges from competition for a berth, with vessel-level AIS data and with the fee that is charged in practice.

Table 1. Position of this paper in the literature.

Study	Method	Waiting at terminals	Cost basis	Perspective
Konings et al. [10]	Cost model	Input per call	Modelled handling margin	Network
Caris et al. [16]	Discrete-event simulation	Simulated	Not assessed	Hubs and terminals
Nicolet et al. [9]	Assessment on regional flows	Aggregate	Modelled, incl. capital cost	Business case of the hub
Bedoya-Maya et al. [14]	Spatial economic assessment	Not simulated	Modelled	Transport chain
Zhou et al. [19]	Network design optimisation	Deterministic	Technical system cost	Network designer
Vermaat [8]	Qualitative framework	Not quantified	Not quantified	Multiple actors
This paper	Paired simulation of both routings with AIS data	Simulated per vessel	Tariff charged in practice	Third-party barge operator

III. Methodology

A. Research design

The study is a quantitative single case study with discrete-event simulation (DES) as the core method. DES suits the problem because congestion at terminals comes from the interaction of random arrivals, variable handling, opening hours and priority rules. Bu and Nachtmann [4] recommend simulation combined with case data for container-on-barge systems for this reason. Two models are built. The hub model represents the routing through the hub. The direct-call model represents the current routing in which the inland barge calls at the deepsea terminal itself. Both models share the inland-leg samplers, the vessel classes, the cost rates and the random number streams. A difference between the routings can therefore only come from the port-side part of the round trip and from the fees. The models are written in Python with the SimPy library. One time unit is one hour.

B. Hub-and-spoke model

Figure 1 shows the conceptual model. Containers arrive in batches at two deepsea terminals. Each terminal has one berth and one crane. A single shuttle serves both terminals on a fixed timetable. At each terminal the shuttle loads within a fixed window and leaves when the window closes, also when it is not full. This window is what allows the shuttle to avoid the barge queue. The shuttle then sails to the hub, where it has priority at the berth.

The hub has one berth and two cranes. It is open on weekdays and on Saturday morning and closed on Sunday. Handling stops when the hub closes and resumes when it opens. A second shuttle from Antwerp

occupies the hub berth three times per week. Only its berth time is modelled, since its cargo is outside the scope.

The inland barges form the fleet of one third-party operator. A barge arrives at the hub, waits for containers and for the berth, loads during opening hours and leaves when it is full or when the hub stock is empty. It then sails the inland leg, which consists of a sailing time and a dwell time at the destination. Both are sampled from distributions fitted to the AIS data. The empirical dwell already contains the real waiting at the inland terminal. The inland berth is therefore sized so that it never binds, and inland waiting is treated as exogenous. Waiting at the hub is endogenous. The model splits it into three mechanisms, which are the berth queue, the hours the hub is closed and supply starvation.

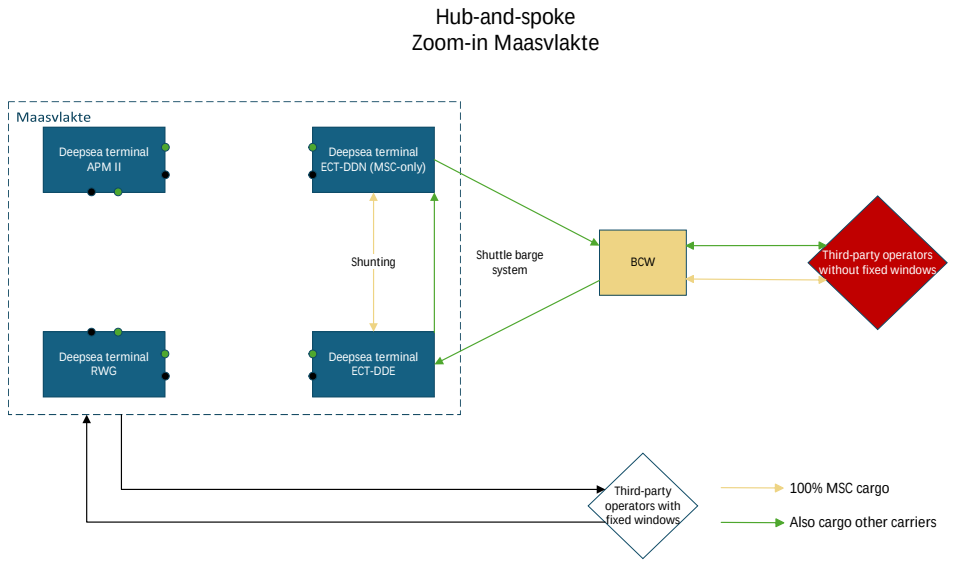

Figure 1. Conceptual model of the hub-and-spoke network. The shuttle sails between the deepsea terminals and the hub. The inland barges sail between the hub and the hinterland.

C. Direct-call model

The direct-call model removes the hub and the shuttle. In 140 of the 142 observed cycles the barge calls at a single terminal. The model therefore has one deepsea terminal with one berth and one crane. The container stock at this terminal never runs dry, so every barge loads to its full capacity. Compared to the hub route the barge sails 3.5 hours further in each direction. Deepsea vessels have priority at the berth. This is modelled as blocks in which the berth is reserved for deepsea vessels. A block does not interrupt a barge that is already loading, but a waiting barge cannot moor until the block ends. A barge without a fixed window pays no call fee at the deepsea terminal. Each direct-call scenario uses the same fleet as the hub scenario it is compared with. The comparison is thus a routing choice made by one operator with one fleet. Figure 2 visualises the direct call.

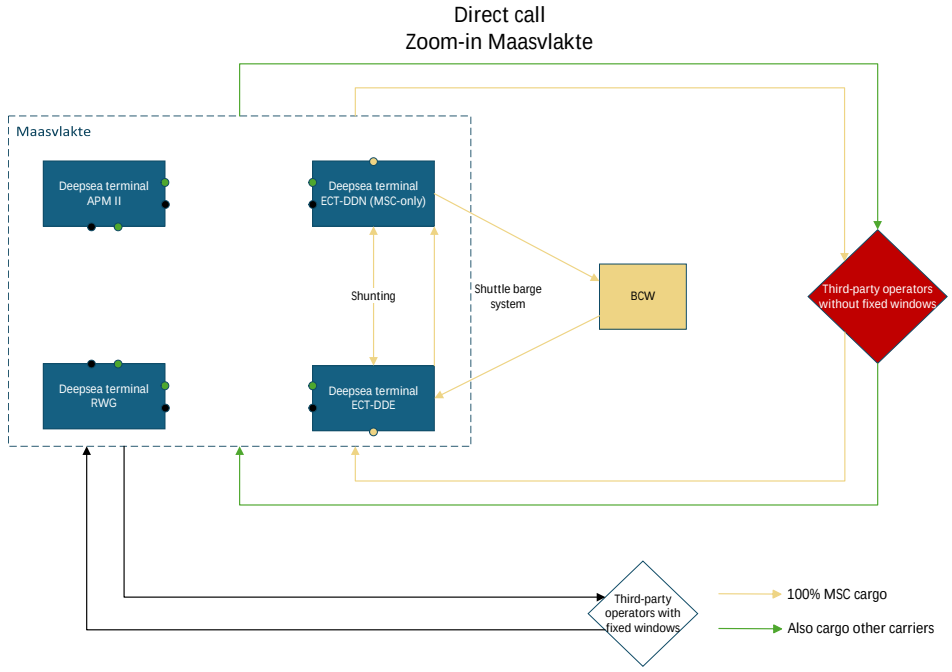

Figure 2. Conceptual model of the direct call. The inland barges sail between the deepsea terminals and the hinterland.

D. Stochastic inputs

Table 2 lists the random inputs of both models. Each source draws from its own random stream and all streams derive from one master seed per replication. A given seed therefore gives the same random inputs in every scenario and in both routings. This is the basis for the paired comparison with common random numbers [20].

Table 2. Stochastic inputs of the two models.

Input	Distribution	Model	Source
Container batch inter-arrival time	Exponential, mean 12 h per terminal	Both	Assumption
Crane moves per hour	Nominal rate × normal factor (mean 0.92, sd 0.08), truncated to [0.5, 1.0]	Both	Assumption
Shuttle sailing time	3.0 h plus a gamma-distributed delay, with a mean of 3.5 h	Hub	MSC
Inland sailing time (round trip)	Weibull per profile, see Section IV.B	Both	AIS
Inland dwell time	Weibull per profile, see Section IV.B	Both	AIS

E. Cost layer

The cost layer translates the time-based simulation output into operator cost per TEU. Two components make up the cost. First, the sailing hours are charged at a sailing rate and the waiting hours at a lower waiting rate per vessel class. Sailing hours are computed as turnaround time minus waiting time and loading time, so congestion shows up purely as more waiting hours at the waiting rate. The waiting hours include the waiting at the hub or the deepsea terminal and the dwell at the inland terminal. In the direct-call simulation the seven extra hours of sailing to the Maasvlakte are part of the sailing hours. The loading hours are multiplied by the

waiting rate. The sailing and waiting rates per class are taken from Ministerie van Infrastructuur en Waterstaat [21], which updates the rates per month. Second, the shuttle transhipment is charged per container moved, since that is the way third-party operators need to pay for the shuttle service. The costs are 90 EUR/TEU and 120 EUR per forty-foot equivalent unit (FEU). With the use of the TEU per move factor of 1.6, this comes down to an average of 108 EUR per container. The per-move charge is the structural cost disadvantage of the hub configuration, since the hub adds an extra handling step [11].

In the direct-call simulation no costs are assigned to the moves at the deepsea terminal, since the price per move is paid by the carrier MSC and inland barges without a fixed window do not have to pay for a call at the deepsea terminal. Instead the barges need to wait much longer before they can load compared to the shuttle barges in the hub simulation. The free deepsea call and the shuttle fee therefore define the central trade-off. The barge avoids the waiting at the deepsea terminal and pays a charge per container for the shuttle leg instead.

The Value of Time and Value of Reliability used by Bedoya-Maya et al. [14] are deliberately excluded. They are behavioural valuations rather than real expenditures, and including them alongside the waiting-hour costs would double-count delay.

Cost per TEU is computed by forming the ratio of cost to delivered TEU within each replication and only then averaging across replications, so the reported confidence interval reflects run-to-run variation in the ratio itself.

F. Output analysis

Each scenario is run for 30 replications of 200 days. The warm-up period is set with the statistical process control method of Robinson [22] on the barge turnaround time, which is the slowest converging indicator. The longest warm-up over all scenarios is 114 hours. A common warm-up of 120 hours is used, which removes about 2 percent of each run. Means and paired differences are reported with a 95 percent confidence interval,

$$\bar{\Delta} \pm t_{0.975,\,29}\,\frac{s_{\Delta}}{\sqrt{30}}, \tag{1}$$

where s_{Δ} is the standard deviation of the 30 paired differences [23]. A difference is significant at the 5 percent level when its interval does not contain zero. The half-width of every interval for turnaround time, cost per TEU and containers delivered is below 1.5 percent of its mean, which is well below the target of 5 percent.

IV. Case study

A. Setting and parameters

The case is the shuttle service of MSC between the Europe Container Terminals (ECT); Delta Dedicated North (ECT-DDN) and Delta Dedicated East (ECT-DDE) terminals on the Maasvlakte and the BCW in the Waalhaven. The service runs today for MSC clients with trucks. The case asks whether third-party inland barges that now call at the Maasvlakte themselves would gain from collecting their containers at BCW instead. Only the import flow is modelled. Table 3 gives the main parameters.

Table 3. Main parameters of the case.

Parameter	Value	Source
Container batches ECT-DDN / ECT-DDE	100 / 80 containers, mean interval 12h	Assumption
Deepsea crane rate	20 moves per hour	Hutchison Ports ECT Rotterdam [24]
Shuttle capacity and frequency	288 TEU, 3 trips per week, at BCW Mon, Wed, Fri 06:00	MSC
Fixed window per deepsea terminal	6h	MSC
Sailing time Maasvlakte to BCW	3.5h each way	MSC, Assumption
Hub berth and cranes	1 berth, 2 cranes, 25 moves per hour	MSC
Hub holding capacity	2,700 TEU	MSC
Hub opening hours	Mon-Fri 06:00-22:00, Sat 07:00-13:00	MSC
Other shuttle at the hub berth	3 times per week, 7 h each	MSC
Deepsea blocks (direct call)	5 per week, 21 h each	MSC
Inland barge capacity M6 / M8 / M9	90 / 192 / 256 TEU	Veldman [25]
Sailing rate M6 / M8 / M9	158.69 / 194.66 / 199.59 EUR per hour	Ministerie van Infrastructuur en Waterstaat [21]
Waiting rate M6 / M8 / M9	94.50 / 116.87 / 115.27 EUR per hour	Ministerie van Infrastructuur en Waterstaat [21]
TEU per container	1.6	Van Koningsveld et al. [12]
Shuttle fee	90 EUR per TEU, 120 EUR per FEU (67.50 EUR per TEU)	MSC
Deepsea call fee without fixed window	0 EUR	Hutchison Ports ECT Rotterdam [24]

B. Data and inland-leg distributions

The empirical basis is AIS data of the inland barges that carry MSC containers over about 40 days. Port areas, berths and waiting areas were drawn as polygons and every status change of a barge was matched to an area. This gives 142 in-port cycles at the two ECT terminals and 285 inland trips. Three trips with a dwell of more than one week are removed, which leaves 282 trips. The trips are split at the median combined inland leg of 55.7 hours into a short and a long profile of 141 trips each.

For each profile the sailing and the dwell component are fitted separately. The lognormal, gamma and Weibull distributions are fitted by maximum likelihood and ranked with the Akaike information criterion [23]. The Weibull distribution has the lowest value for three of the four components and reproduces the sample mean and standard deviation within 2 percent. Goodness of fit is tested with a Kolmogorov-Smirnov test with a parametric bootstrap, because the parameters are estimated from the same sample [26]. The dwell fits are accepted at the 5 percent level. The sailing fits are rejected for every family, because the median split truncates the sailing component. The fitted distributions therefore act as a smooth version of the data that keeps its level and its spread.

C. Scenarios and experiments

Six baseline scenarios are formed by three vessel classes (M6, M8 and M9) and two inland-leg profiles (short and long). The fleet in each scenario is calibrated in the hub model. It is the smallest fleet at which the weekly deliveries reach the throughput ceiling of the hub-and-spoke system. This gives four and seven barges for M6 on the short and the long leg, and two and three barges for M8 and M9.

Four experiments follow on the M6 class, because its waiting at the hub is dominated by the berth queue that the hub operator can influence. Experiment 1 adds a Sunday opening from 07:00 to 13:00. Experiment 1a applies the same Sunday opening at the baseline fleet, to separate the effect of the measure from the effect of a smaller fleet. Experiment 2 moves the shuttle arrival at BCW to 14:00. Experiment 3 opens the hub

continuously, which is not a realistic proposal but a bound on what any opening-hours policy can reach. In Experiments 1, 2 and 3 the hub fleet is recalibrated, which gives three and six barges. The direct call keeps its baseline fleet and results. Experiment 4 congests the deepsea terminal. Such periods occur when cranes at the terminal are broken or under maintenance, and barges without a fixed window can then wait up to 72 hours. The experiment lengthens the blocks to 27 hours, which is the longest block at which the direct call can still move the 540 containers per week that the hub moves. The direct-call fleet on the long leg then has to grow from seven to eight barges. The hub model is unchanged, because the shuttle loads within its own window.

D. Verification and validation

The models are verified and validated following Sargent [27] and Robinson [28]. Event traces of the shuttle and of an inland barge were checked against the conceptual model. Three container flow identities were asserted at the end of every replication and their residuals are exactly zero in all 30 replications. Extreme condition tests behaved as predicted. With a single inland barge the hub stock rises to its ceiling and stays there. Without container arrivals all flows fall to zero in both models. The conceptual model was reviewed by the supervisor at MSC.

Only the direct-call model is validated against the real world, because the hub routing does not yet exist in the modelled form. The simulated in-port time per class is compared with the observed mean and its 95 percent confidence interval. For M8 the model gives 15.55 hours against an observed 14.00 hours (interval 11.38 to 16.61). For M9 it gives 18.95 hours against 15.84 hours (interval 8.34 to 23.33). Both lie inside the interval. For M6 the model gives 12.20 hours against 18.59 hours (interval 14.83 to 22.35), which is 6.39 hours too short. The most likely reason is that the AIS data mixes barges with and without a fixed window. Only 16 percent of the M6 cycles wait less than one hour, against 59 percent for M8 and 46 percent for M9, so the M6 cycles look like calls without a window. Another reason is that only the import flow of containers is used in this research, while the data contains the export flow as well. The underestimate makes the direct call look faster and cheaper than it is. The M6 results are therefore conservative for the hub.

V. Results

A. Baseline comparison of the two routings

Table 4 compares the two routings for the six operator profiles. The direct call is cheaper in every profile. The gap runs from 63.85 EUR per TEU for M8 on the long leg to 84.07 EUR per TEU for M9 on the short leg. All cost differences are significant.

Table 4. Baseline results per operator profile. Differences are hub minus direct, taken per seed, with the 95 percent half-width. A positive value means the hub takes longer or costs more. Waiting is per port visit.

			Turnaround (h)				Cost per TEU (EUR)		
Class	Leg	Fleet	Hub	Direct	Δ	ΔWait (h)	Hub	Direct	Δ
M6	short	4	58.5	57.7	+0.81 ± 0.28	+7.25 ± 0.28	155.58	79.79	+75.79 ± 0.36
M6	long	7	114.2	111.1	+3.11 ± 0.62	+9.91 ± 0.53	232.71	158.12	+74.59 ± 1.05
M8	short	2	58.5	61.2	−2.68 ± 0.42	+4.14 ± 0.38	121.66	47.97	+73.69 ± 0.31
M8	long	3	112.8	115.1	−2.29 ± 0.53	+4.67 ± 0.51	157.23	93.38	+63.85 ± 0.79
M9	short	2	64.7	64.4	+0.36 ± 0.48*	+6.60 ± 0.42	121.81	37.74	+84.07 ± 0.27
M9	long	3	121.9	118.8	+3.14 ± 1.03	+8.33 ± 0.93	154.14	72.52	+81.62 ± 0.85

* Not significant at the 5 percent level.

On time the picture is mixed. The hub is faster only for the M8 class, by 2.68 hours on the short leg and 2.29 hours on the long leg. The M6 class is 0.81 and 3.11 hours slower and the M9 class is 3.14 hours slower on the long leg. The M9 difference on the short leg is not significant. The inland leg is drawn from the same sampler in both routings, so the turnaround difference has to equal the difference in port-side time. The two agree within 0.55 hours in every profile, which confirms that the seven-hour round trip to the Maasvlakte is attributed correctly.

Waiting tells a more consistent story. In the hub situation the barge waits longer in all six profiles, by 4.14 to 9.91 hours per visit. Where the hub situation still wins on the total port call, it wins through the seven sailing hours that the shuttle takes over and not through less queueing. The hub replaces one berth queue with three waiting mechanisms. For M6 the berth queue is 72 to 75 percent of the waiting at the hub, because four to seven barges share one berth. For M8 the closing hours of the hub are the largest part, at 52 to 53 percent, because two or three barges arrive outside the opening hours relatively often. For M9 the berth queue dominates again, because each call loads for 9.5 to 11.1 hours.

B. Decomposition of the cost gap

Table 5 splits the gap. The fee term of 67.50 EUR per TEU is between 80 and 106 percent of the total gap. The sailing term is small and mostly favours the hub, because the shuttle takes over the expensive sailing hours to the Maasvlakte. The waiting term favours the direct call in all six profiles. Even if the hub waited as little as the direct call, it would still cost 60.32 to 70.25 EUR per TEU more.

Table 5. Decomposition of the cost gap into sailing, waiting and fee (EUR per TEU). Waiting includes the waiting at the hub or the deepsea terminal, the loading time and the dwell at the inland terminal, all at the waiting rate.

		Hub			Direct			Difference		
Class	Leg	Sail	Wait	Fee	Sail	Wait	Fee	Sail	Wait	Gap
M6	short	40.34	47.74	67.50	47.01	32.77	0.00	−6.67	+14.97	+75.79
M6	long	94.16	71.05	67.50	101.34	56.78	0.00	−7.18	+14.27	+74.59
M8	short	24.71	29.45	67.50	26.94	21.03	0.00	−2.23	+8.42	+73.69
M8	long	51.34	38.38	67.50	58.06	35.31	0.00	−6.72	+3.07	+63.85
M9	short	23.25	31.06	67.50	20.69	17.05	0.00	+2.56	+14.01	+84.07
M9	long	47.45	39.19	67.50	44.70	27.82	0.00	+2.75	+11.37	+81.62

The gap is not ordered by vessel size. The M6 class lies between M8 and M9. Two effects explain this. The direct call has economies of scale, so its cost per TEU falls from 79.79 EUR for M6 to 37.74 EUR for M9 on the short leg. The hub has a ceiling on its supply, which takes away part of the scale advantage of the larger classes, as the next section shows.

C. Throughput and call size

The handling capacity of the hub is not the binding constraint. The holding capacity of 2,700 TEU is reached only in the extreme condition test. What binds is the shuttle. It carries 288 TEU, which is 180 containers at 1.6 TEU per container, and it makes three trips per week. All six scenarios therefore converge on a throughput of about 540 containers per week.

As a result the call size of the inland barge is a system outcome and not a vessel property. A barge leaves the hub when the stock is empty, so its load equals what the shuttle delivered since the previous visit. The call size is the weekly supply divided by the number of visits of the fleet. On the short leg an M9 barge leaves the hub at only 66 percent of its capacity. One M9 barge makes too few visits to carry the flow, and two barges make 50 percent more visits than the supply can fill. The M9 class therefore has no cost advantage over M8 on the short leg, at 121.81 against 121.66 EUR per TEU. The direct call has no supply ceiling. With the same fleet it delivers 1,385 to 8,236 more containers over a run in five of the six profiles. The exception is M8 on the long leg, where the hub barge already leaves at 99 percent of capacity. That profile also has the smallest cost gap.

D. Experiments on the M6 class

Table 6 reports the experiments. The three operational measures leave very different waiting times at the hub, between 4.16 and 11.30 hours per visit on the short leg and between 2.50 and 11.93 hours on the long leg. Their cost results are nearly the same. Each lets the operator run one barge fewer, and the gap falls from about 75 to between 56.14 and 58.01 EUR per TEU. Continuous operation removes almost all waiting and is no cheaper than six hours of Sunday opening.

Table 6. Experiments on the M6 class. Gaps are hub minus direct in EUR per TEU, taken per seed.

				Baseline terminal		Congested terminal	
Experiment	Leg	Fleet	Wait (h)	Gap	Gap excl. fee	Gap	ΔTT (h)
Baseline	short	4	16.13	+75.79	+8.29	+64.43	−9.94
1 Sunday opening	short	3	7.42	+56.71	−10.79	+45.35	−18.24
1a Sunday, baseline fleet	short	4	13.84	+77.42	+9.92	–	–
2 Shuttle at 14:00	short	3	11.30	+57.43	−10.07	+46.07	−15.66
3 Continuous operation	short	3	4.16	+58.01	−9.49	+46.65	−23.36
Baseline	long	7	19.00	+74.59	+7.09	+48.69	−21.51
1 Sunday opening	long	6	8.09	+56.14	−11.36	+30.23	−31.58
1a Sunday, baseline fleet	long	7	11.12	+76.47	+8.97	–	–
2 Shuttle at 14:00	long	6	11.93	+57.93	−9.57	+32.02	−29.18
3 Continuous operation	long	6	2.50	+56.90	−10.60	+30.99	−39.08

Experiment 1a shows where the saving comes from. It applies the Sunday opening at the original fleet. Waiting falls by 2.29 and 7.88 hours and the turnaround falls by 2.37 and 7.18 hours. Yet the cost per TEU rises by 1.63 and 1.88 EUR. The hub receives about 540 containers per week whatever the barges do, so a faster barge does not carry more. It calls more often with a smaller load and converts waiting hours into more expensive sailing hours. The saving in Experiments 1, 2 and 3 therefore comes from the barge that is removed and not from the measure itself.

Without the fee, the baseline hub is 8.29 and 7.09 EUR per TEU more expensive to operate than the direct call. After any of the three measures the sign reverses and the hub becomes the cheaper operation by 9.49 to 11.36 EUR per TEU. The operation of the hub is thus not the problem. The fee is.

Under congestion the time verdict reverses. The direct call now spends 22.71 hours at the terminal on the short leg and 36.85 hours on the long leg, against about 12 hours in the baseline. The direct fleet on the long leg needs an eighth barge to keep 540 containers per week moving, and that barge joins the same single berth. The hub is faster in every row, by 9.94 to 23.36 hours on the short leg and by 21.51 to 39.08 hours on the long leg. The cost verdict does not reverse. The gap narrows to 30.23 to 32.02 EUR per TEU on the long leg after the operational measures, but it does not close. The 27-hour block represents a period with broken cranes or cranes under maintenance and not an average week.

E. Break-even fee

For the M6 baseline the break-even fee is below zero, so even a free shuttle would leave the hub behind. After the operational measures it rises to between 9.49 and 11.36 EUR per TEU. In the most favourable case tested, the congested terminal with the Sunday opening on the long leg, it reaches 37.27 EUR per TEU. All of these values lie below 60 EUR per TEU, which is the lowest fee at which the shuttle still covers its own operation. No fee exists in the tested cases that is low enough for the operator and high enough for the shuttle.

VI. Discussion

The result confirms a theoretical prediction rather than contradicting it. Konings et al. [10] found that the handling step at a hub on the Eem-Waalhaven relation may cost at most 1 to 19 EUR per TEU before the network loses its advantage. The fee charged today is 67.50 EUR per TEU, which is more than three times the top of that range. A side note is that inflation has not been taken into account. Bedoya-Maya et al. [14] also

found negative savings for consolidation at seaports because handling costs dominate. This paper reaches the same conclusion with a different method, with waiting that emerges from berth competition and with the price that is charged in practice. A barge without a fixed window pays nothing for its deepsea call. The hub therefore has no call fee to save. What the operator buys with the fee is time and in the baseline the hub delivers little of it.

The hub and the deepsea terminal each have one berth for inland barges. Every barge that is added joins the same queue, and the operator pays for the barge and for the hours it waits. This mechanism explains both groups of experiments. The operational measures help the hub only when they let the operator remove a barge. Congestion helps the hub because it forces the direct-call operator to add one. A change that removes a barge from the hub routing, or adds one to the direct routing, moves the comparison toward the hub. Congestion at the deepsea terminal is worth more to the hub than any opening-hours policy at the hub itself. The decision to use a hub should therefore be made at fleet level and not per call.

A. Future research on barge profiles

The time saved per call is worth more per TEU when the call is small. The benefit of the hub per TEU thus rises as the call size falls, provided the hub is the faster route. This points toward smaller barges rather than larger ones and agrees with Nicolet et al. [9], who find that a consolidation terminal suits barges that carry less than 60 TEU. It runs against the common view that larger inland vessels are the answer for hinterland container transport, without contradicting the economies of scale on which that view rests. Four profiles were identified but not simulated. The first is the barge that makes several deepsea calls per cycle, which Roux et al. [3] report as common for small barges. Each extra call adds sailing and another queue that the hub would remove. The second is the barge below the M6 class or with calls below the minimum call size of the terminal, which is 10 moves at DDE and 20 moves at DDN [24]. The third is a cheaper shuttle, for example a larger shuttle, a higher use of the existing one or direct transhipment with push barges [10]. The fourth is an operator who values a short and predictable turnaround. Under congestion such an operator pays about 30 EUR per TEU extra for a round trip that is more than a day shorter and may accept that price.

B. Limitations

The AIS data covers about 40 days and does not show which barges hold a fixed window. This likely explains the underestimate of the M6 in-port time, which works against the hub. The model contains only the import flow, while the observed in-port times include loading and unloading. In the direct call each barge visits one terminal, as in 140 of the 142 observed cycles. More terminal calls per cycle would favour the hub but would also require a shuttle that visits more terminals. All containers at the hub leave by barge, while in practice part of them leave by truck, which would lower the call size of the inland barges. The internal process of the hub is not modelled, so the time at the hub is a lower bound. The barges sail continuously and water levels are not modelled.

VII. Conclusion

This paper compared the routing of third-party inland barges through a hub-and-spoke system with the direct call at the deepsea terminal. Both routings were simulated with the same fleet and the same random numbers for six operator profiles in the Port of Rotterdam. Routing through the hub changes the turnaround time by between -2.68 and $+3.14$ hours and raises the transport cost by 63.85 to 84.07 EUR per TEU. For every profile tested the direct call is the cheaper routing.

The main reason is the shuttle fee. The operator pays 67.50 EUR per TEU for the shuttle against a free call at the deepsea terminal, and the hub does not save enough time to cover that. The operation of the hub itself is not the problem, since after the operational measures it is cheaper than the direct call when the fee is

left out. The throughput of the system is set by the shuttle, which makes the call size of the inland barge a system outcome. The size of the fleet is the main lever for both time and cost. When the deepsea terminal is congested, the hub becomes clearly faster but remains more expensive.

The recommendations of this research consist of four directions. The model should be extended with barges that make several deepsea calls per cycle and with a shuttle that serves more terminals. The value of reliability should be added to the cost comparison, to test whether the time saved under congestion is worth the remaining gap. Data that separates barges with and without a fixed window, over a longer period, would give a cleaner reference for the direct call. And the final direction is the fluctuation in water levels.

For practice, the structure of the tariff matters more than the operation of the hub. In every tested case the break-even fee lies below the level at which the shuttle covers its own cost. A lower fee therefore needs a cheaper shuttle operation, for example a larger shuttle, a higher use of the existing one or direct transhipment with push barges. Of the operational measures, the later shuttle arrival can be adopted directly, because it releases a barge at no extra cost for the hub. MSC can focus the service on operators without a fixed window that face long waiting at the deepsea terminal. Third-party barge operators should judge the hub at fleet level, because the saving comes from the barges they can remove or do not have to add.

References

[1] Shenkar, N., and Rosen, D., "How has the invention of the shipping container influenced marine bioinvasion?" *Management of Biological Invasions*, Vol. 9, No. 3, 2018, pp. 187–194. doi: 10.3391/mbi.2018.9.3.02.

[2] Port of Rotterdam, "Cargo throughput in the port of Rotterdam slightly decreased in 2024," , 2 2025.

[3] Roux, L., Krämer, I., and Konings, R., "Handling of Inland Vessels in Seaports – Necessary Actions and Additional Options to Support Container Transport on Inland Waterways," *Proceedings of PIANC Smart Rivers 2022*, Vol. 264, edited by Y. Li, P. Rigo, G. Zhao, Y. Hu, and F. E. Lefler, Springer, Singapore, 2023, pianc 2022 ed., pp. 1388–1397. doi: 10.1007/978-981-19-6138-0_123.

[4] Bu, F., and Nachtmann, H., "Literature review and comparative analysis of inland waterways transport: "Container on Barge"," *Maritime Economics and Logistics*, Vol. 25, No. 1, 2023, pp. 140–173. doi: 10.1057/s41278-021-00195-6.

[5] European Commission, "Communication from the commission to the European Parliament, the council, the European economic, and social committee and the committee of the regions," , 6 2021.

[6] Friedhoff, B., Ramne, B., Alias, C., van Hassel, E., Martens, S. E., Felde, J. z., and Samuel, L., "Elevating the logistics resilience of the Rhine-Alpine Corridor with the help of innovative vessel and cargo handling concepts," *Transportation Research Procedia*, Vol. 72, Elsevier B.V., 2023, pp. 1998–2005. doi: 10.1016/j.trpro.2023.11.681.

[7] Oganesian, V., Sys, C., Vanelslander, T., and van Hassel, E., "Container barge (un)reliability in seaports: a company case study at the port of Antwerp," , 2021.

[8] Vermaat, T., "Development of a feasibility assessment framework for consolidation hubs in hinterland container barge transport," , 2026. URL `https://resolver.tudelft.nl/uuid:7b0ab80e-6afb-43f3-b544-94ff6d940a96`.

[9] Nicolet, A., Shobayo, P., van Hassel, E., and Atasoy, B., "An assessment methodology for a modular terminal concept for container barging in seaports," *Case Studies on Transport Policy*, Vol. 14, 2023. doi: 10.1016/j.cstp.2023.101103.

[10] Konings, R., Kreutzberger, E., and Maraš, V., "Major considerations in developing a hub-and-spoke network to improve the cost performance of container barge transport in the hinterland: The case of the port of rotterdam," *Journal of Transport Geography*, Vol. 29, 2013, pp. 63–73. doi: 10.1016/j.jtrangeo.2012.12.015.

[11] Otten, M., Scholten, P., Tol, E., van de Lande, P., Verbeek, M., and Wagter, H., "Outlook hinterland and continental freight 2020," , 8 2020. URL `https://ce.nl/publicaties/outlook-hinterland-and-continental-freight-2020/`.

[12] Van Koningsveld, M., Verheij, H. J., Taneja, P., and De Vriend, H. J., "Ports and Waterways: Navigating the changing world," , 2021.

[13] Shobayo, P., and van Hassel, E., "Container barge congestion and handling in large seaports: a theoretical agent-based modeling approach," *Journal of Shipping and Trade*, Vol. 4, No. 1, 2019. doi: 10.1186/s41072-019-0044-7.

[14] Bedoya-Maya, F., Shobayo, P., Nicolet, A., Christopoulou, E., Majoor, I., van Hassel, E., and Vanelslander, T.,

"Cargo consolidation in port-hinterland container transport: A spatial economic assessment for inland waterways," *Research in Transportation Business and Management*, Vol. 59, 2025. doi: 10.1016/j.rtbm.2024.101254.

[15] Hu, Q., Wiegmans, B., Corman, F., and Lodewijks, G., "Critical literature review into planning of inter-terminal transport: In port areas and the hinterland," *Journal of Advanced Transportation*, Vol. 2019, 2019. doi: 10.1155/2019/9893615.

[16] Caris, A., Macharis, C., and Janssens, G. K., "Network analysis of container barge transport in the port of Antwerp by means of simulation," *Journal of Transport Geography*, Vol. 19, No. 1, 2011, pp. 125–133. doi: 10.1016/j.jtrangeo.2009.12.002.

[17] Zheng, J., and Yang, D., "Hub-and-spoke network design for container shipping along the Yangtze River," *Journal of Transport Geography*, Vol. 55, 2016, pp. 51–57. doi: 10.1016/j.jtrangeo.2016.07.001.

[18] Ypsilantis, P., and Zuidwijk, R., "Collaborative fleet deployment and routing for sustainable transport," *Sustainability (Switzerland)*, Vol. 11, No. 20, 2019. doi: 10.3390/su11205666.

[19] Zhou, S., Ji, B., Song, Y., Yu, S. S., Zhang, D., and Van Woensel, T., "Hub-and-spoke network design for container shipping in inland waterways," *Expert Systems with Applications*, Vol. 223, 2023. doi: 10.1016/j.eswa.2023.119850.

[20] Glasserman, P., and Yao, D. D., "Some Guidelines and Guarantees for Common Random Numbers," *Management Science*, Vol. 38, No. 6, 1992, pp. 884–908. doi: 10.1287/mnsc.38.6.884.

[21] Ministerie van Infrastructuur en Waterstaat, "Bedrijfseconomische kostenkengetallen voor het goederenvervoer, update 2026," , 2026. URL https://www.kimnet.nl/documenten/2026/05/28/bedrijfseconomische-kostenkengetallen-voor-het-goederenvervoer-update-2026.

[22] Robinson, S., "A statistical process control approach to selecting a warm-up period for a discrete-event simulation," *European Journal of Operational Research*, Vol. 176, No. 1, 2007, pp. 332–346. doi: 10.1016/j.ejor.2005.07.014.

[23] Law, A. M., "Statistical Analysis of Simulation Output Data: The Practical State of the Art," *2020 Winter Simulation Conference (WSC)*, IEEE, 2020, pp. 1117–1127. doi: 10.1109/WSC48552.2020.9383993.

[24] Hutchison Ports ECT Rotterdam, "Handling of barges ECT Rotterdam," , 1 2026. URL https://www.nextlogic.nl/zowerkthet/barge-operator/.

[25] Veldman, J., "Optimal vessel class for container transport from hinterland to seaport, costs comparison between Class IV and Class Va." , 2019.

[26] Lilliefors, H. W., "On the Kolmogorov-Smirnov Test for Normality with Mean and Variance Unknown," *Journal of the American Statistical Association*, Vol. 62, No. 318, 1967, pp. 399–402. doi: 10.1080/01621459.1967.10482916.

[27] Sargent, R. G., "Verification and validation of simulation models," *Proceedings of the 2011 Winter Simulation Conference*, IEEE, 2011, pp. 183–198.

[28] Robinson, S., "Simulation model verification and validation: increasing the users' confidence," *Proceedings of the 1997 Winter Simulation Conference*, 1997, pp. 53–59.

B

Assumptions and simplifications

Table B.1: Assumption register for the conceptual model

#	Assumption	Justification	Direction of effect
A1	Two deepsea terminals in the hub simulation, one in the current-state simulation	The shuttle collects at both ECT terminals, while 140 of the 142 observed cycles call at a single terminal	Favours the direct call if barges call at more terminals
A2	One inland terminal represents the destination side	The empirical dwell already contains real berth waiting at the destination	Neutral, the same leg is charged in both routings
A3	Only the import flow is modelled instead of barges carrying containers both ways	The research question concerns the collection leg from the deepsea terminal	Neutral
A4	The hub location is taken as given	The question is the performance of routing via this hub and not its siting	Neutral by construction
A5	One berth call per deepsea cycle	Nearly all observed cycles touch a single terminal, remoorings are merged	Understates the direct call if multi-terminal calls occur
A6	Inland-side waiting is exogenous and the inland berth is never binding	The empirical dwell already contains the real waiting at the destination	Would inflate the hub routing if violated
A7	Sailing and dwell are sampled independently within each subset	Within-subset dependence is weak	Would understate the variance of the leg if violated
A8	Batch arrivals of 100 and 80 containers at a mean interval of 12 hours with a 60/40 split	Volume split based on the empirical record and on Nicolet et al. (2023)	Sets the supply of the shuttle. The container stock of the direct call never runs dry
A9	The shuttle holds priority at the hub berth	The shuttle operates to a fixed window agreed with the terminal	Inflates hub-side waiting for third-party barges
A10	The shuttle fee enters the cost model as a linear per-container term	The fee has no feedback on the simulated physical dynamics	Enables the break-even fee calculation on the saved output

#	Assumption	Justification	Direction of effect
A11	Water levels are not modelled	Water levels were outside of the scope	Favours the direct call at normal water levels
A12	Deepsea vessels block the berth five times per week for 21 hours in the current-state simulation	Represents the priority of deepsea vessels over barges without a fixed window, with the same schedule for all classes	Sets the waiting of the direct call.
A13	Block the berth five times per week for 27 hours in the current-state simulation	Represents a block caused by deepsea vessel prioritisation and broken cranes	Sets extra waiting time on the direct call
A14	Direct-call barges have no fixed window and pay no call fee	Barges without a fixed window pay no call fee at ECT	Favours the direct call on cost
A15	All containers at the hub are collected by inland barges	Truck collection from the hub is outside the scope of this research	Favours the hub, since trucks would lower the call size of the inland barges
A16	Barges sail at full capacity in current-state simulation	Economies of scale play a large role in transportation via barge	Favours direct-call since container stock at deepsea terminal is larger than at hub
A17	The internal process of the hub is not modelled	Containers are moved as one counted stock at the crane rate	Favours the hub, since the hub-side time is a lower bound
A18	Inland barges sail continuously	Every barge starts a new cycle as soon as it returns	Affects both routings in the same way
A19	Shuttle sailing time consists of a gamma distribution delay plus a minimum of 3.0 h	No observed distribution of the shuttle system.	Neutral
A20	Crane moves follow normal distribution	Crane maintenance and crane operator change	Neutral, every crane follows the same distribution
A21	The shuttle barge between Antwerp and BCW only occupies the berth without unloading containers	The route between Antwerp and BCW is outside the scope of this thesis	Favours hub since hub capacity is not affected

C

Model graphs

C.1. SPC warm-up control charts

Figure C.1 shows the SPC control charts of the mean barge turnaround time for all six scenarios, from which the warm-up periods in Table 7.1 were determined.

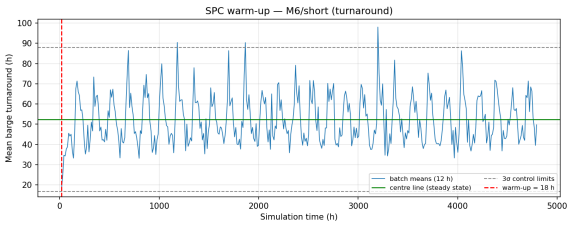

(a) M6, short inland leg

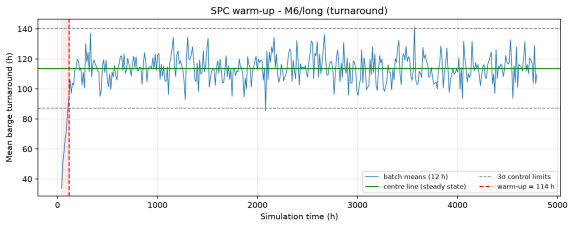

(b) M6, long inland leg

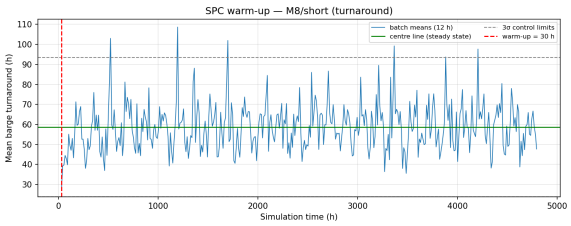

(c) M8, short inland leg

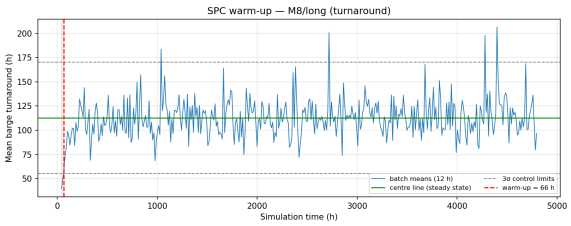

(d) M8, long inland leg

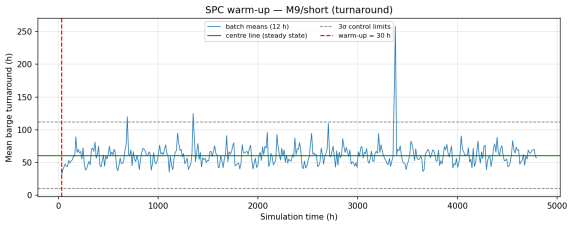

(e) M9, short inland leg

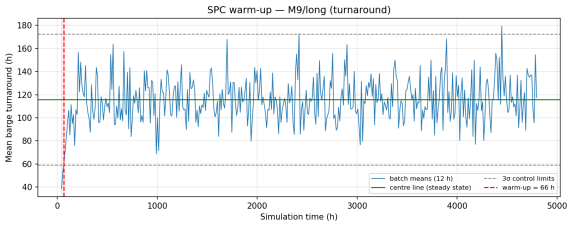

(f) M9, long inland leg

Figure C.1: Warm-up period of the mean barge turnaround time for all six scenarios

C.2. Hub stock

In B.2 the hub stock per scenario is presented. The six different graphs show the number of containers within the simulated 200 days per scenario. One thing to notice is that the graphs show one run of 200 days and not an average of the 30 replications.

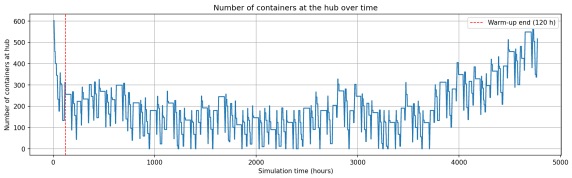

(a) Hub stock, 1 replication, class M6, long leg

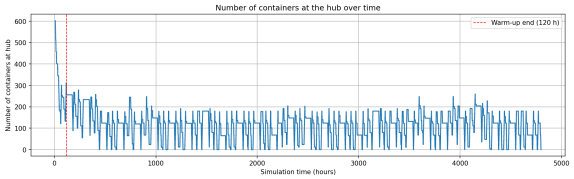

(b) Hub stock, 1 replication, class M6, short leg

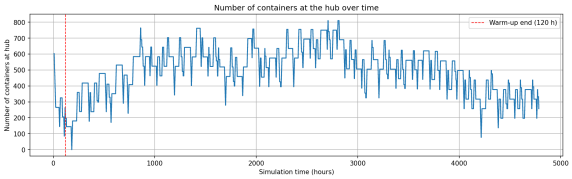

(c) Hub stock, 1 replication, class M8, long leg

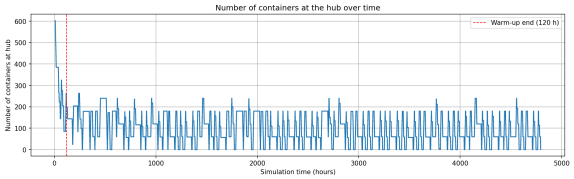

(d) Hub stock, 1 replication, class M8, short leg

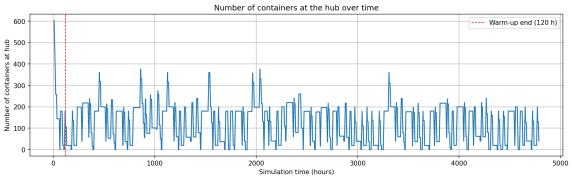

(e) Hub stock, 1 replication, class M9, long leg

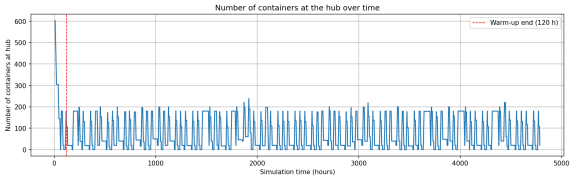

(f) Hub stock, 1 replication, class M9, short leg

Figure C.2: Stock of the hub per class on the short and long leg

C.3. Waiting-time

The following figures show the waiting times of the inland barges resulting from the six tested scenarios at the hub. The waiting time is decomposed in four different waiting times; supply starvation, berth-queue, closed-hub, and loading time.

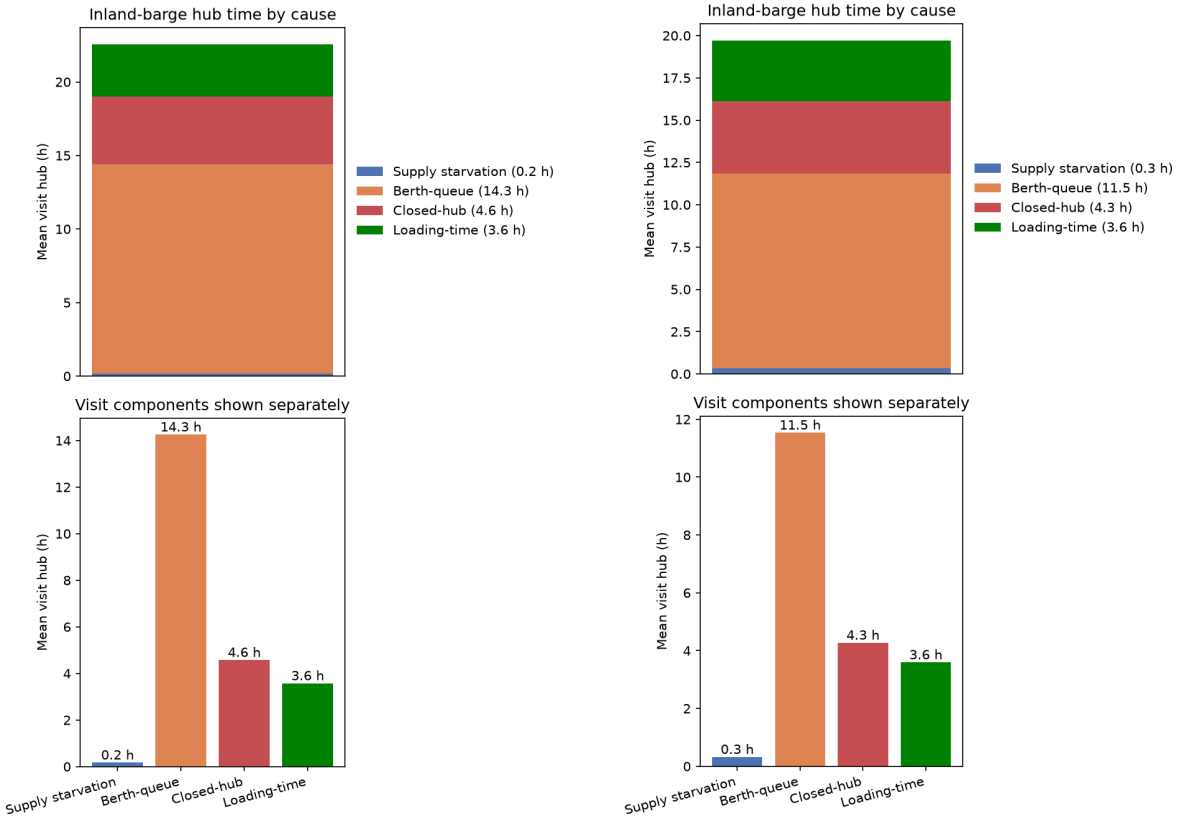

(a) Time at hub, class M6, long leg

(b) Time at hub, class M6, short leg

Figure C.3: Waiting time at the hub per class on the short and long leg

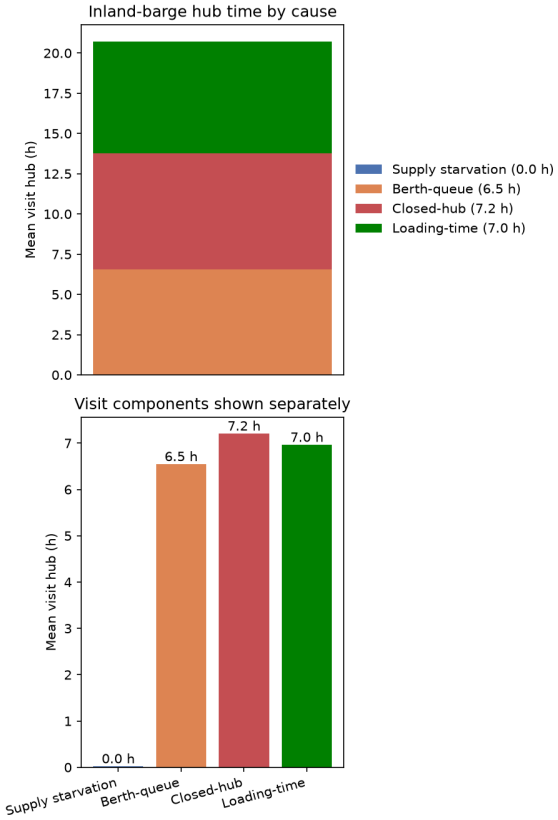

(c) Time at hub, class M8, long leg

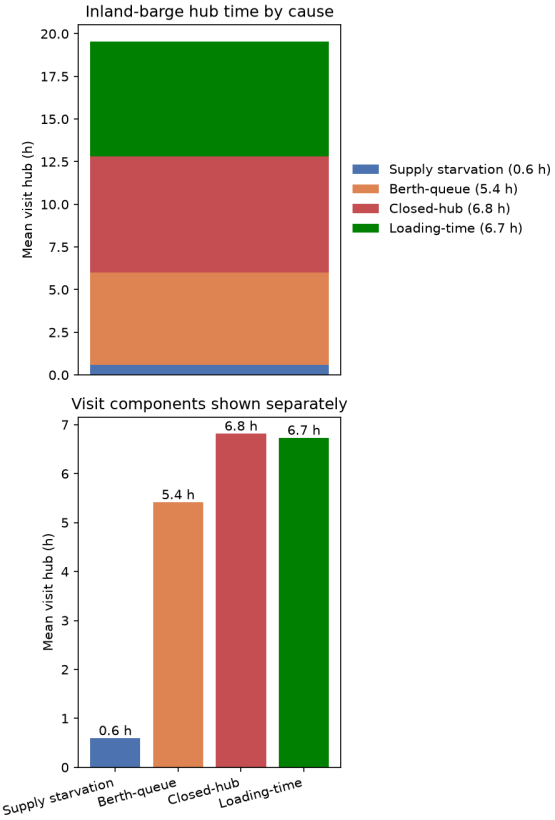

(d) Time at hub, class M8, short leg

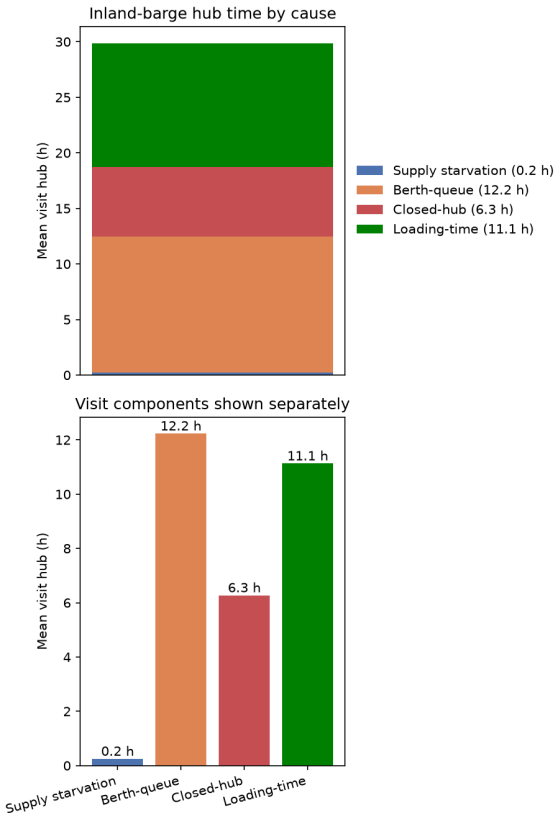

(e) Time at hub, class M9, long leg

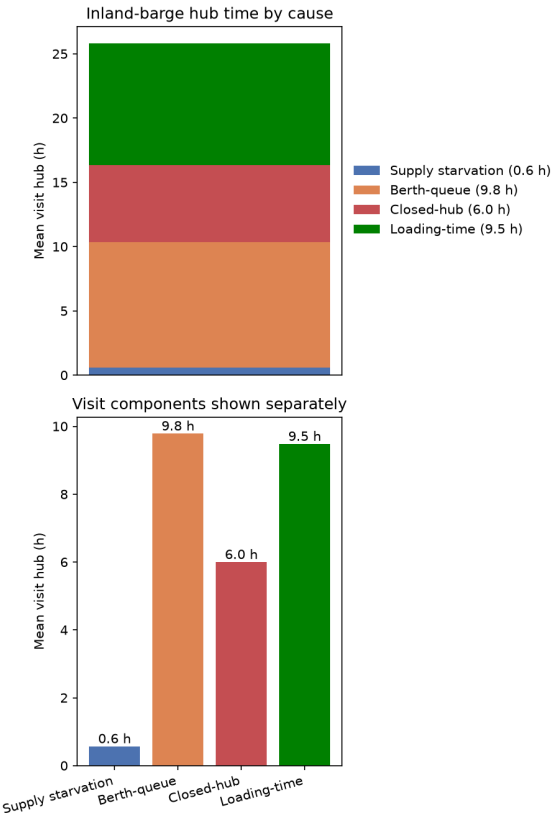

(f) Time at hub, class M9, short leg

Figure C.3: Waiting time at the hub per class on the short and long leg (continued)

D

Verification and validation

D.1. Trace verification of the model

A part of the verification of the model is the tracing part of certain units within the simulation. By using print statements after commands that are given in the code, the units which are necessary to follow can be traced. With the print statements listed as in figure D.2 and figure D.1, both the models can be checked on their correctness.

```
[t=0.0] Shuttle schedule built: 85 departures (3/week target, fleet 1 barges)
[t=0.0] Inland barge 1 arrives at hub (cycle 1)
[t=46.0] Shuttle trip 1 departs
[t=46.0] Shuttle 1 arrives at DST1
[t=46.0] Shuttle 1 berths at DST1(window open for 6 h)
[t=46.0] Shuttle 1 mooring at DST1 (0.10 h)
[t=52.0] Shuttle 1 un-mooring at DST1 (0.10 h)
[t=52.1] Shuttle 1 finished at DST1 (108 containers)
[t=52.1] Shuttle 1 sails DST1 -> DST2
[t=52.6] Shuttle 1 arrives at DST2
[t=52.6] Shuttle 1 berths at DST2(window open for 6 h)
[t=52.6] Shuttle 1 mooring at DST2 (0.10 h)
[t=56.5] Shuttle 1 un-mooring at DST2 (0.10 h)
[t=56.6] Shuttle 1 finished at DST2 (72 containers)
[t=56.6] Shuttle 1 leaves seaport with 180 containers
[t=60.7] Shuttle 1 arrives at hub
[t=60.7] Shuttle 1 mooring at hub (0.10 h)
[t=61.9] Inland barge 1 requests hub berth
[t=68.7] Shuttle 1 finished unloading at hub | at hub: 180
[t=68.7] Shuttle 1 un-mooring at hub (0.10 h)
[t=68.8] Inland barge 1 berths at hub for loading
[t=68.8] Inland barge 1 mooring at hub (0.10 h)
[t=68.9] Inland barge 1 starts loading (waiting time at hub = 68.9 h)
[t=70.9] Hub CLOSED -> Inland barge 1 waits 7.1 h for opening
[t=73.0] Shuttle trip 1 completed (cycle = 27.0 h)

[t=78.0] Hub OPEN -> Inland barge 1 resumes
[t=78.5] Inland barge 1 un-mooring at hub (0.10 h)
[t=78.6] Inland barge 1 finished loading (loading time = 9.7 h)
[t=78.6] Inland barge 1 departs hub
[t=81.9] Inland barge 1 arrives at inland terminal
[t=81.9] Inland barge 1 berths at inland terminal for unloading
[t=93.9] Inland barge 1 delivered 56 containers in 12.0 h dwell
[t=93.9] Inland barge 1 sails back to hub
[t=94.0] Shuttle trip 2 departs
[t=94.0] Shuttle 2 arrives at DST1
[t=94.0] Shuttle 2 berths at DST1(window open for 6 h)
[t=94.0] Shuttle 2 mooring at DST1 (0.10 h)
[t=97.2] Inland barge 1 completed cycle (turnaround time = 97.2 h)
```

Figure D.1: Output print statements hub-and-spoke simulation

```
scenario: M6 / long / 7 barges | cap=56 moves | BARGE_SHARE=1.0
[t=0.0] Container batch 1 arrives at DST (180 containers, total waiting: 180)
[t=   0.00 h | day  1 00:00] Barge 1   enters the simulation at the deep-sea terminal
[t=   0.00 h | day  1 00:00] Barge 1   cycle 1: arrives at the deep-sea terminal, 180 containers on the quay
[t=   0.00 h | day  1 00:00] Barge 1   berth is free, takes it straight away
[t=   0.00 h | day  1 00:00] Window    window 1 is due, but Barge 1 is on the berth, so it waits
[t=   0.00 h | day  1 00:00] Barge 1   moors (0.10 h)
[t=   1.10 h | day  1 01:06] Barge 1   crane at 17 moves/h loads 17 containers in 1.00 h: 17/56 on board, 163 left on the quay
[t=   2.10 h | day  1 02:06] Barge 1   crane at 17 moves/h loads 17 containers in 1.00 h: 34/56 on board, 146 left on the quay
[t=2.8] Container batch 2 arrives at DST (180 containers, total waiting: 326)
[t=   3.10 h | day  1 03:06] Barge 1   crane at 19 moves/h loads 19 containers in 1.00 h: 53/56 on board, 307 left on the quay
[t=   3.25 h | day  1 03:15] Barge 1   crane at 20 moves/h loads 3 containers in 0.15 h: 56/56 on board, 304 left on the quay
[t=   3.25 h | day  1 03:15] Barge 1   loading done (barge full): 56 containers in 3.15 h, unmoors (0.10 h)
[t=   3.35 h | day  1 03:21] Barge 1   leaves the berth. Deep-sea call 3.35 h = supply wait 0.00 + queue 0.00 + (un)mooring 0.20 + loading 3.15
[t=   3.35 h | day  1 03:21] Barge 1   sails to the inland terminal: 11.03 h (3.5 h Maasvlakte leg + 7.53 h inland leg)
[t=   3.35 h | day  1 03:21] Window    window 1 starts: berth blocked for 21.0 h (priority over inland barges)
[t=   6.00 h | day  1 06:00] Barge 2   enters the simulation at the deep-sea terminal
[t=   6.00 h | day  1 06:00] Barge 2   cycle 1: arrives at the deep-sea terminal, 304 containers on the quay
[t=   6.00 h | day  1 06:00] Barge 2   berth is taken by window 1, joins the queue (order: Barge 2)
[t=  12.00 h | day  1 12:00] Barge 3   enters the simulation at the deep-sea terminal
[t=  12.00 h | day  1 12:00] Barge 3   cycle 1: arrives at the deep-sea terminal, 304 containers on the quay
[t=  12.00 h | day  1 12:00] Barge 3   berth is taken by window 1, joins the queue (order: Barge 2, Barge 3)
[t=14.2] Container batch 3 arrives at DST (180 containers, total waiting: 484)
[t=  14.38 h | day  1 14:22] Barge 1   arrives at the inland terminal, dwell 28.88 h
[t=  18.00 h | day  1 18:00] Barge 4   enters the simulation at the deep-sea terminal
[t=  18.00 h | day  1 18:00] Barge 4   cycle 1: arrives at the deep-sea terminal, 484 containers on the quay
[t=  18.00 h | day  1 18:00] Barge 4   berth is taken by window 1, joins the queue (order: Barge 2, Barge 3, Barge 4)
[t=22.7] Container batch 4 arrives at DST (180 containers, total waiting: 664)
[t=  24.00 h | day  2 00:00] Barge 5   enters the simulation at the deep-sea terminal
[t=  24.00 h | day  2 00:00] Barge 5   cycle 1: arrives at the deep-sea terminal, 664 containers on the quay
[t=  24.00 h | day  2 00:00] Barge 5   berth is taken by window 1, joins the queue (order: Barge 2, Barge 3, Barge 4, Barge 5)
[t=  24.35 h | day  2 00:21] Window    window 1 ends: berth free again
[t=  24.35 h | day  2 00:21] Barge 2   gets the berth after 18.35 h in the queue
[t=  24.35 h | day  2 00:21] Barge 2   moors (0.10 h)
[t=  25.45 h | day  2 01:27] Barge 2   crane at 17 moves/h loads 17 containers in 1.00 h: 17/56 on board, 647 left on the quay
[t=  26.45 h | day  2 02:27] Barge 2   crane at 17 moves/h loads 17 containers in 1.00 h: 34/56 on board, 630 left on the quay
[t=  27.45 h | day  2 03:27] Barge 2   crane at 15 moves/h loads 15 containers in 1.00 h: 49/56 on board, 615 left on the quay
[t=  27.84 h | day  2 03:50] Barge 2   crane at 18 moves/h loads 7 containers in 0.39 h: 56/56 on board, 608 left on the quay
[t=  27.84 h | day  2 03:50] Barge 2   loading done (barge full): 56 containers in 3.39 h, unmoors (0.10 h)
[t=  27.94 h | day  2 03:56] Barge 2   leaves the berth. Deep-sea call 21.94 h = supply wait 0.00 + queue 18.35 + (un)mooring 0.20 + loading 3.39
[t=  27.94 h | day  2 03:56] Barge 2   sails to the inland terminal: 17.63 h (3.5 h Maasvlakte leg + 14.13 h inland leg)
[t=  27.94 h | day  2 03:56] Barge 3   gets the berth after 15.94 h in the queue
[t=  27.94 h | day  2 03:56] Barge 3   moors (0.10 h)
[t=  29.04 h | day  2 05:02] Barge 3   crane at 19 moves/h loads 19 containers in 1.00 h: 19/56 on board, 589 left on the quay
[t=  30.00 h | day  2 06:00] Barge 6   enters the simulation at the deep-sea terminal
[t=  30.00 h | day  2 06:00] Barge 6   cycle 1: arrives at the deep-sea terminal, 589 containers on the quay
[t=  30.00 h | day  2 06:00] Barge 6   berth is taken by Barge 3, joins the queue (order: Barge 4, Barge 5, Barge 6)
[t=  30.04 h | day  2 06:02] Barge 3   crane at 20 moves/h loads 20 containers in 1.00 h: 39/56 on board, 569 left on the quay
[t=  30.93 h | day  2 06:56] Barge 3   crane at 19 moves/h loads 17 containers in 0.89 h: 56/56 on board, 552 left on the quay
[t=  30.93 h | day  2 06:56] Barge 3   loading done (barge full): 56 containers in 2.89 h, unmoors (0.10 h)
[t=  31.03 h | day  2 07:02] Barge 3   leaves the berth. Deep-sea call 19.03 h = supply wait 0.00 + queue 15.94 + (un)mooring 0.20 + loading 2.89
[t=  31.03 h | day  2 07:02] Barge 3   sails to the inland terminal: 30.72 h (3.5 h Maasvlakte leg + 27.22 h inland leg)
[t=  31.03 h | day  2 07:02] Barge 4   gets the berth after 13.03 h in the queue
[t=  31.03 h | day  2 07:02] Barge 4   moors (0.10 h)
[t=  32.13 h | day  2 08:08] Barge 4   crane at 16 moves/h loads 16 containers in 1.00 h: 16/56 on board, 536 left on the quay
[t=  33.13 h | day  2 09:08] Barge 4   crane at 19 moves/h loads 19 containers in 1.00 h: 35/56 on board, 517 left on the quay
[t=  33.60 h | day  2 09:36] Window    window 2 is due, but Barge 4 is on the berth, so it waits
[t=  34.13 h | day  2 10:08] Barge 4   crane at 20 moves/h loads 20 containers in 1.00 h: 55/56 on board, 497 left on the quay
[t=  34.19 h | day  2 10:11] Barge 4   crane at 18 moves/h loads 1 containers in 0.06 h: 56/56 on board, 496 left on the quay
[t=  34.19 h | day  2 10:11] Barge 4   loading done (barge full): 56 containers in 3.06 h, unmoors (0.10 h)
[t=  34.29 h | day  2 10:17] Barge 4   leaves the berth. Deep-sea call 16.29 h = supply wait 0.00 + queue 13.03 + (un)mooring 0.20 + loading 3.06
[t=  34.29 h | day  2 10:17] Barge 4   sails to the inland terminal: 15.43 h (3.5 h Maasvlakte leg + 11.93 h inland leg)
[t=  34.29 h | day  2 10:17] Window    window 2 starts: berth blocked for 21.0 h (priority over inland barges)
[t=  36.00 h | day  2 12:00] Barge 7   enters the simulation at the deep-sea terminal
[t=  36.00 h | day  2 12:00] Barge 7   cycle 1: arrives at the deep-sea terminal, 496 containers on the quay
[t=  36.00 h | day  2 12:00] Barge 7   berth is taken by window 2, joins the queue (order: Barge 5, Barge 6, Barge 7)
[t=  43.26 h | day  2 19:15] Barge 1   has delivered 56 containers, sails back to the deep-sea terminal: 11.03 h
```

Figure D.2: Output print statements current state simulation