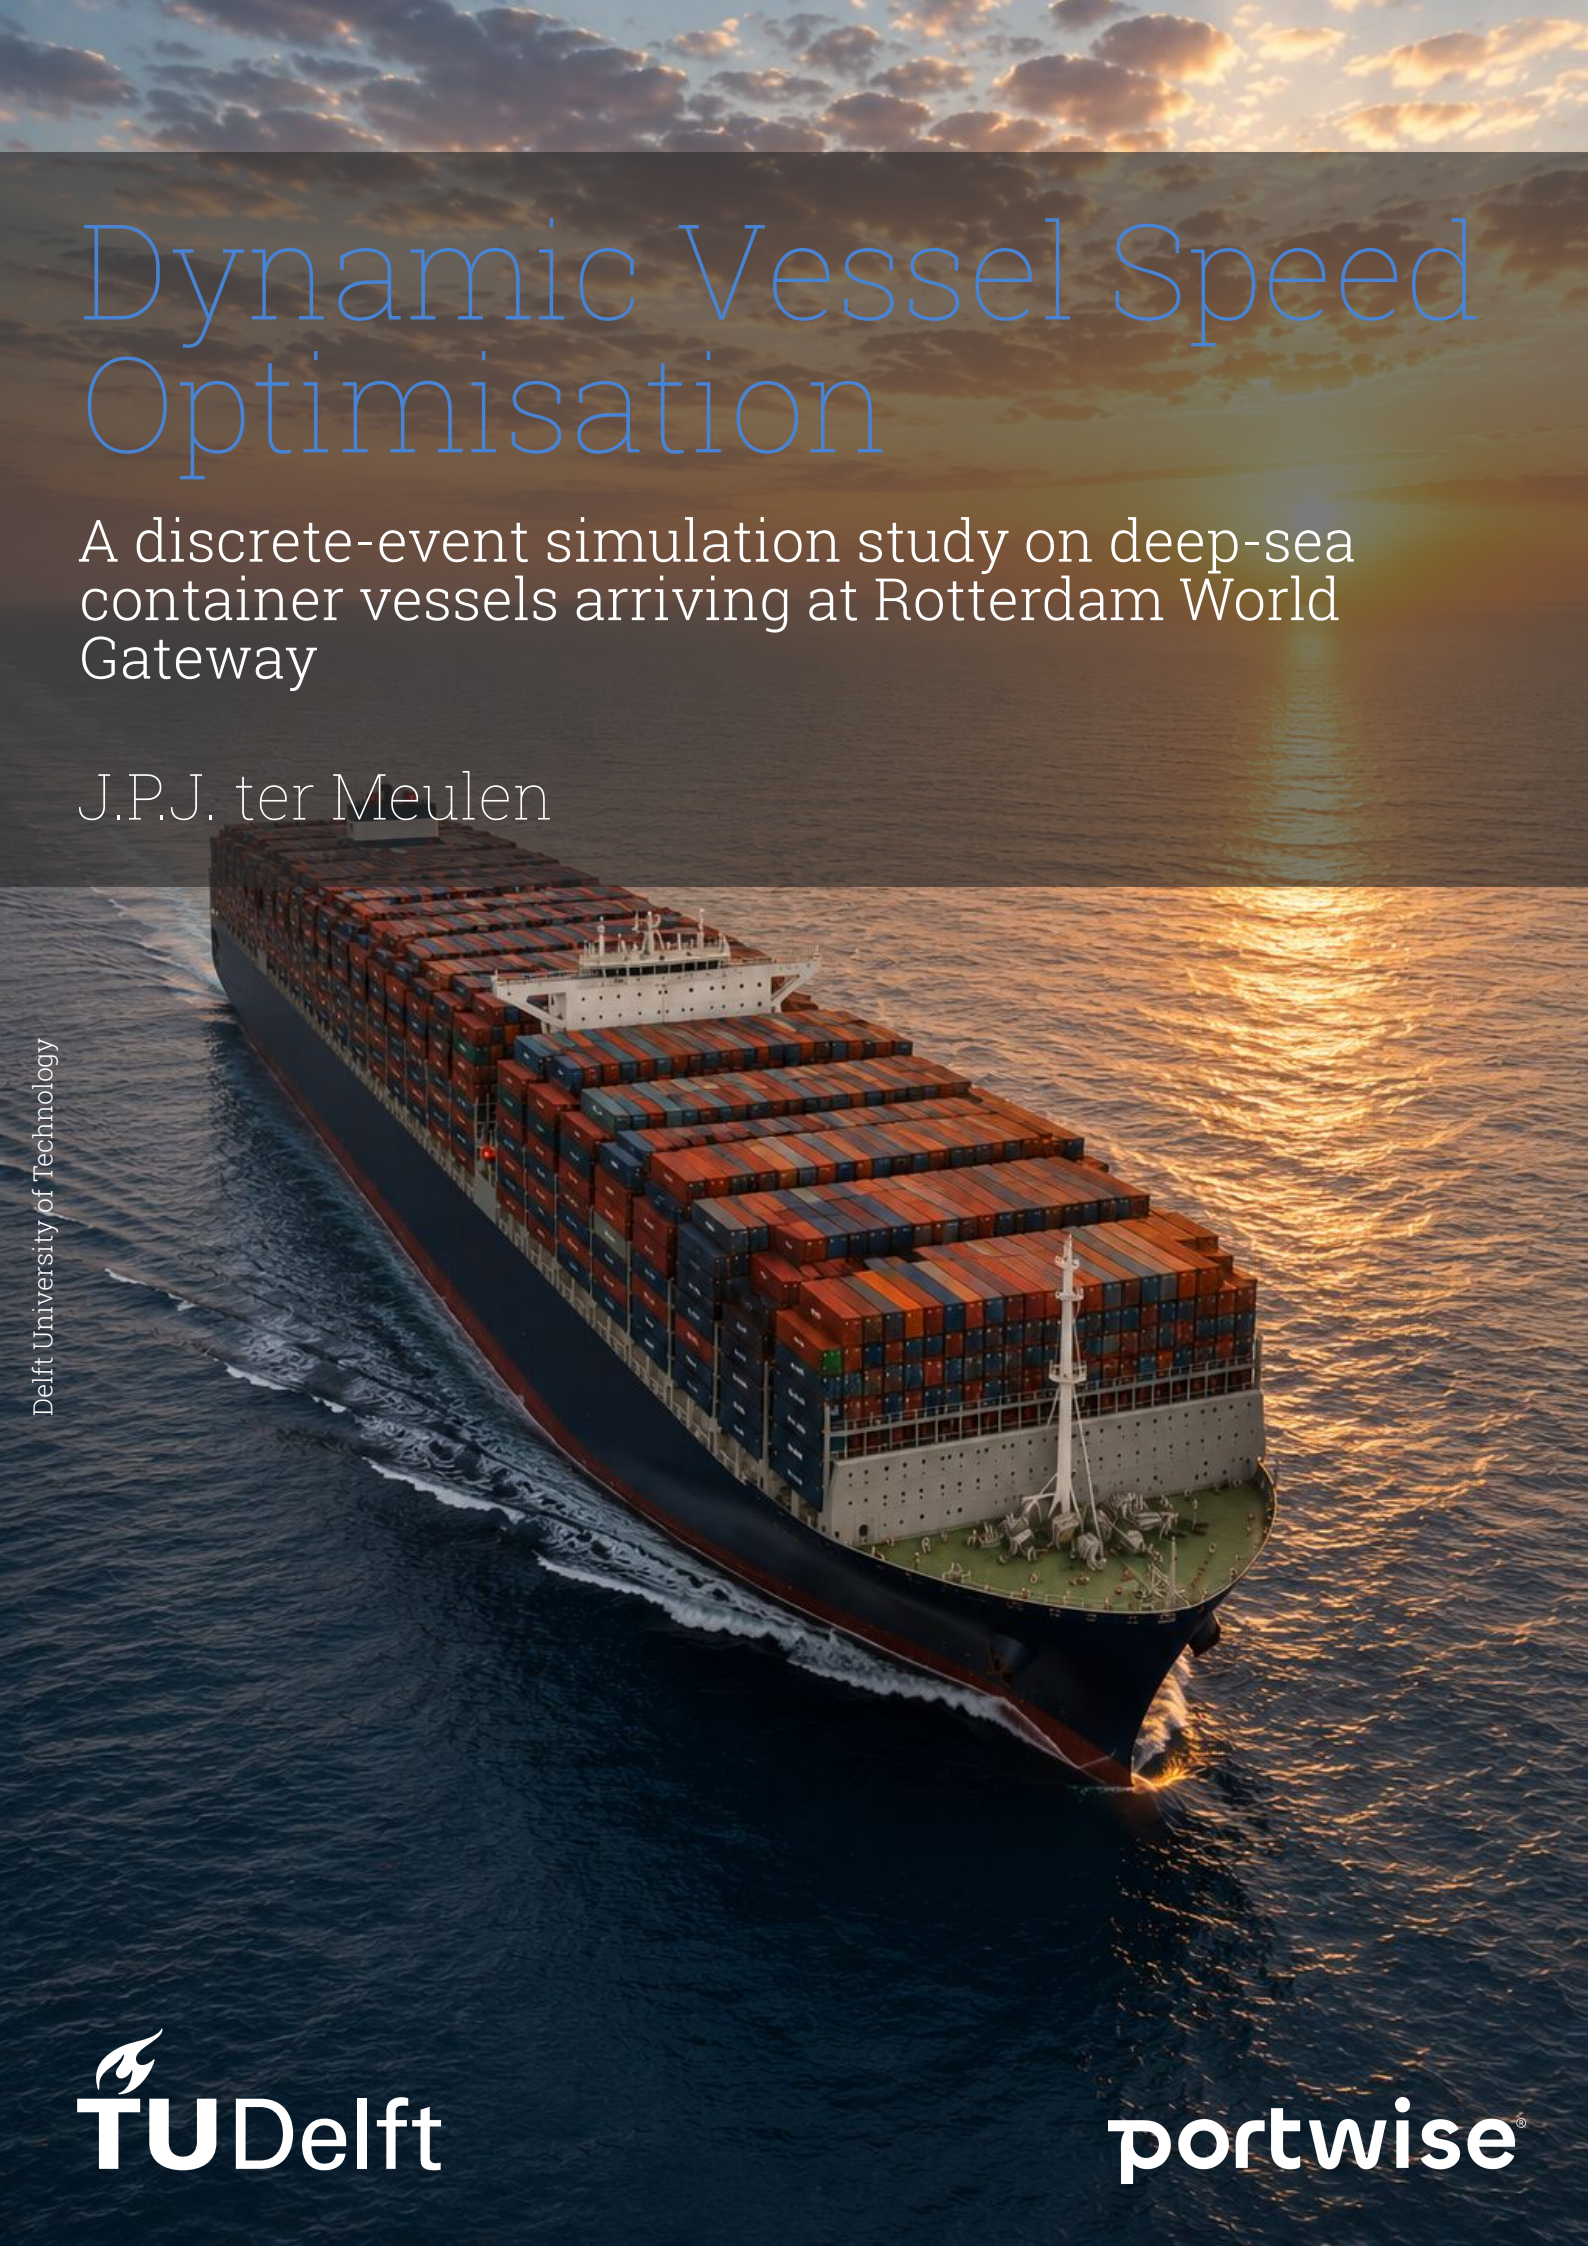
Dynamic Vessel Speed Optimisation

A discrete-event simulation study on deep-sea container vessels arriving at Rotterdam World Gateway

J.P.J. ter Meulen



Delft University of Technology

TU Delft

portwise®

Dynamic Vessel Speed Optimisation

A discrete-event simulation study on deep-sea container vessels arriving at Rotterdam World Gateway

by

J.P.J. ter Meulen

Master Thesis in partial fulfilment of the requirements for the degree of

Master of Science
in Mechanical Engineering

At the Department Maritime and Transport Technology of Faculty Mechanical Engineering of Delft University of Technology
to be defended publicly on Thursday June 11, 2026 at 1:00 PM

Name:	J.P.J. ter Meulen	
Student number:	4894170	
Report number:	2026.MME.9196	
Faculty:	Faculty of Mechanical Engineering, Delft	
Mastertrack:	Multi-Machine Engineering	
Thesis Committee:	Ir. M.B. Duinkerken	TU Delft, Chair
	Prof.dr. R.R. Negenborn	TU Delft, Committee member
	M.L. Ochoa Barnuevo, MSc	TU Delft, Committee member
	A. de Waal, MSc	Portwise, Company supervisor
	A.J. Dijkstra, MSc	Portwise, Company supervisor

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

Cover:	AI-Generated image made with ChatGPT
Style:	TU Delft Report Style, with modifications by Daan Zwaneveld

TU Delft

portwise

Preface

This thesis is the result of my graduation assignment in Mechanical Engineering at Delft University of Technology with a specialisation in the Multi-Machine Engineering track. Completing this thesis marks the end of my time as a student, which I have enjoyed very much.

I have never before spent this much time on a singular topic. However, it has sparked my interest in the logistics surrounding container terminals. I am intrigued by how all the moving components, from deep-sea container vessels to autonomous guided vehicles, operate within the same ecosystem to move containers from one side of the world to another.

During my time at Portwise, I enjoyed working in a different environment, surrounded by colleagues working on simulation projects while I focused on my thesis. Whenever I needed help, whether with coding or logic, someone was always willing to assist me, regardless of the complexity of the question. It did not take me long to learn the coffee orders of everyone around me, although the policy of getting coffee for the room may have increased my own caffeine consumption. My time at Portwise has made the transition from studying to working feel less significant. I will miss the weekly Monday round-up and the lunches on Tuesdays and Thursdays.

I would like to thank Arjen and Age, and all the other employees at Portwise. Your expertise on the simulation software and the topic of container terminals showed me what was possible and what I should aim to achieve in this thesis.

I would also like to express my gratitude to Mark Duinkerken and Rudy Negenborn for their valuable feedback on the thesis and the guidance throughout the project. Your advice helped to shape this thesis and to improve the quality of the work.

Finally, I would like to express my thanks to my family and friends. My parents, Akosua and Rik, for their support throughout my studies. My friends, who have listened to me talk about optimising the velocity of container vessels without complaining. And to my girlfriend, Lara, for her interest in a topic far removed from her own studies, her help with writing and her confidence in me.

Jaap ter Meulen
Delft, June 2026

AI Statement

For this thesis, I have used Generative AI tools in a supportive manner. Grammarly was used to improve the grammar, style, layout, and spelling of the text. In addition, ChatGPT was used to generate the cover image and assisted in identifying coding errors. In all cases, I have reviewed and corrected the work and remain fully responsible for the analysis, methods, results and conclusions presented in this thesis.

Summary

Deep-sea container vessels play a crucial role in the world economy. Despite the enormous operating costs, these vessels often have to wait at the anchorage point due to limited berth availability at container terminals. This thesis investigates the influence of vessel speed optimisation on the performance of container vessels and on berth occupancy at the container terminal, using a discrete-event simulation model.

Most of the existing research focuses either on the berth allocation problem or on vessel speed optimisation. If the two topics are combined, research often focuses on homogeneous vessels or on vessels departing from a fixed point. This thesis focuses on a more realistic scenario: implementing vessel speed optimisation in the berth allocation problem for heterogeneous vessels departing from ports at varying distances from the focal port.

In this model, four different configurations are compared: (1) A baseline configuration, in which vessels sail at their economic speed and are served at the terminal on a *first-come, first-served* basis, (2) single speed optimisation at departure, where the speed of the vessel is optimised only once at the departure of the previous terminal, (3) re-optimisation at new departures, where the speed of the vessel is optimised, when a new vessel departs which has a smaller remaining distance and (4) global speed optimisation, where the terminal adjusts the velocities for all sailing vessels at time intervals, and new vessel departures.

These configurations are compared in a scenario that represents operations at Rotterdam World Gateway in 2024 and serves as a baseline. Besides the baseline scenario, the global speed optimisation is evaluated in two additional scenarios in which the number of vessels arriving at the port is increased.

The model is evaluated on three KPIs: the departure delay, vessel costs and berth occupancy. The vessel's costs are divided into operating, fuel and waiting costs. Together, the departure delay and costs indicate the vessel's performance. The vessel's performance is indicated both at the system level, as an average over all vessels arriving, and at the individual level, over all vessels that have departed from the same port. The terminal's berth occupancy remains the same within a scenario, but it increases across scenarios as the number of vessels arriving every 2 weeks increases.

The results show that the global speed optimisation strategy reduces the average vessel costs by approximately $30000 per trip when compared to the baseline at scenario 1. The results also indicate that with global speed optimisation, the berth occupancy can increase by up to 2.2%, corresponding to approximately 44,000 additional TEU handled annually, while maintaining lower average vessel costs and comparable departure delays compared to the baseline scenario. Further increases, however, lead to congestion and a decrease in overall vessel performance.

These findings demonstrate that improved coordination between vessels and terminals can improve vessel performance while enabling an increase in berth occupancy and throughput at container terminals.

Contents

Preface i

Summary ii

Nomenclature ix

1 Introduction 1
- 1.1 Introduction . . . 1
- 1.2 Problem statement . . . 2
  - 1.2.1 Introduction to the case study . . . 2
- 1.3 Literature review . . . 2
  - 1.3.1 Berth allocation under uncertainty . . . 3
  - 1.3.2 Berth allocation combined with vessel speed optimisation . . . 4
- 1.4 Research gap and scope . . . 4
- 1.5 Research question and sub-questions . . . 4
- 1.6 Outline of the report . . . 5

2 System Analysis on the Arrival of Container Vessels 6
- 2.1 System definition . . . 6
- 2.2 Vessel types . . . 8
  - 2.2.1 Deep-sea vessels . . . 8
  - 2.2.2 Feeder / short-sea vessels . . . 8
  - 2.2.3 Inland barges . . . 9
- 2.3 Seaside operations . . . 9
  - 2.3.1 Analysis of the seaside operations . . . 11
  - 2.3.2 Conclusion . . . 14
- 2.4 Operational pattern of deep-sea Vessels . . . 14
  - 2.4.1 Schematically defined vessel operational pattern . . . 14
  - 2.4.2 Analysis of vessel operations . . . 15
  - 2.4.3 Conclusion . . . 17
- 2.5 Combined system analysis . . . 17
- 2.6 Conclusion . . . 17

3 Literature Review 19
- 3.1 Berth allocation problem . . . 19
- 3.2 Berth allocation under uncertainty . . . 21
  - 3.2.1 Sources of uncertainty . . . 21
  - 3.2.2 Uncertainty optimisation approaches . . . 22
  - 3.2.3 Uncertainty modelling . . . 22
  - 3.2.4 Rolling horizon . . . 23
- 3.3 Vessel speed optimisation . . . 24
  - 3.3.1 Vessel speed optimisation . . . 24
  - 3.3.2 Optimisation of vessel arrival times . . . 24
- 3.4 Stakeholders . . . 25
  - 3.4.1 Vessel operators . . . 25
  - 3.4.2 Container terminal operators . . . 25
  - 3.4.3 Local government . . . 25
  - 3.4.4 Port authority . . . 25
- 3.5 Conclusion . . . 26

4 Speed Optimisation Strategies **27**
4.1 Speed optimisation strategies in literature . . . 27
4.2 Description of strategies . . . 28
4.2.1 Baseline . . . 28
4.2.2 Single-speed optimisation . . . 30
4.2.3 Re-optimise at every new departure . . . 30
4.2.4 Global speed optimisation . . . 32
4.3 Overview . . . 33
4.4 Conclusion . . . 33

5 Model **34**
5.1 Model of the seaside operations of a container terminal . . . 34
5.1.1 Model scope and goal . . . 34
5.1.2 Model inputs . . . 35
5.1.3 Outputs and performance indicators . . . 36
5.1.4 Operations . . . 36
5.2 Model of the sailing leg . . . 37
5.2.1 Model scope and goal . . . 38
5.2.2 Model inputs . . . 38
5.2.3 Outputs and performance indicators . . . 39
5.2.4 Operations . . . 39
5.3 Model integration . . . 40
5.3.1 Integration . . . 40
5.3.2 Planning frequency . . . 40
5.4 Verification and validation . . . 41
5.4.1 Verification of the speed optimisation function . . . 41
5.4.2 Validation . . . 42
5.5 Conclusion . . . 44

6 Experimental Plan **45**
6.1 Input . . . 45
6.1.1 Scenarios . . . 45
6.1.2 Configurations . . . 47
6.2 Key performance indicators . . . 47
6.3 Influence of speed optimisation strategies . . . 47
6.4 The performance of global speed optimisation under scenarios 1, 2 and 3 . . . 48
6.5 Sensitivity analysis . . . 48

7 Results **50**
7.1 Influence of speed optimisation strategies . . . 50
7.2 Influence of scenarios . . . 53
7.3 Sensitivity analysis . . . 57
7.4 Conclusion . . . 58

8 Conclusion and Discussion **60**
8.1 Conclusion . . . 60
8.1.1 Answer to sub-question 1 . . . 60
8.1.2 Answer to sub-question 2 . . . 60
8.1.3 Answer to sub-question 3 . . . 60
8.1.4 Answer to sub-question 4 . . . 61
8.1.5 Answer to sub-question 5 . . . 61
8.1.6 Answer to sub-question 6 . . . 61
8.1.7 Answer to the main research question . . . 61
8.2 Discussion . . . 62
8.3 Recommendations . . . 62
8.3.1 Recommendations for future research . . . 62
8.3.2 Recommendations for data collection . . . 63

References **64**

A Research Paper **66**

B Results **75**
B.1 T-test on sensitivity analysis . . . 75
B.2 Velocity changes . . . 76
B.2.1 All configurations, scenario 1 . . . 76
B.2.2 Baseline scenario 1 and global speed optimisation scenario 1, 2 and 3 . . . 78
B.2.3 Maximum number of velocity changes . . . 79
B.3 Average velocity . . . 80
B.3.1 All configurations, scenario 1 . . . 80
B.3.2 Baseline scenario 1 and global speed optimisation scenario 1, 2 and 3 . . . 81
B.4 Vessel waiting time . . . 82
B.4.1 All configurations, scenario 1 . . . 82
B.4.2 Baseline scenario 1 and global speed optimisation scenario 1, 2 and 3 . . . 83

C Verification checks **85**
C.1 Realised berth schedule . . . 85
C.1.1 Weeks 1 and 2 . . . 85
C.1.2 Weeks 3 and 4 . . . 88
C.1.3 Conclusion . . . 91
C.2 Single-speed optimisation . . . 91
C.3 Global speed optimisation . . . 93
C.3.1 Current sailing vessels . . . 93
C.3.2 Select order . . . 94
C.3.3 Select velocity . . . 94
C.3.4 Conclusion . . . 95
C.4 Berth occupancy . . . 95
C.4.1 Number of vessels handled per simulation . . . 95
C.4.2 Increase in berth occupancy . . . 96

List of Figures

2.1 Schematic sideview of a container terminal [27] . . . 6
2.2 Areas of operation in a container terminal (ECT Delta, Port of Rotterdam) . . . 7
2.3 Three different vessel types at Rotterdam World Gateway (Google Earth 23-07-2024), Left to right, deep-sea vessel, inland vessel, feeder vessel . . . 8
2.4 Double berthing of two barges at the barge quay of Rotterdam World Gateway (Google Earth 23-07-2024) . . . 10
2.5 Seaside operations . . . 10
2.6 Blackbox of the seaside operations of a container terminal . . . 11
2.7 Proforma berth schedule . . . 12
2.8 Realised berth schedule . . . 12
2.9 Deep-sea and feeder/short-sea quay at RWG . . . 13
2.10 Schematically drawn vessel operations . . . 14
2.11 Blackbox of vessel operations . . . 16
2.12 Combined blackbox of vessel operations and seaside operations . . . 18

4.1 A schematic drawing of communications between stakeholders in the baseline strategy 29
4.2 A schematic drawing of communications between stakeholders in the single speed optimisation strategy . . . 30
4.3 A schematic drawing of communications between stakeholders in the re-optimise at new departures strategy . . . 31
4.4 A schematic drawing of communications between stakeholders in the global speed optimisation strategy . . . 32

5.1 Visualisation of the seaside operation model . . . 34
5.2 Swimlane diagram of the berth assignment . . . 37
5.3 Swimlane diagram of the loading and unloading of containers at the quay . . . 37
5.4 Swimlane diagram of the sailing operations of the vessel under the baseline speed optimisation strategy . . . 39
5.5 Swimlane diagram of the sailing operations of the vessel under the single-speed optimisation strategy . . . 39
5.6 Swimlane diagram of the sailing operations of the vessel under the re-optimise at new departures, and global speed optimisation strategy . . . 40
5.7 Average vessel service time per TEU category . . . 43
5.8 Average vessel waiting time per TEU category . . . 43

7.1 Average vessel cost per configuration . . . 50
7.2 Average vessel cost per port and TEU category per configuration . . . 51
7.3 Average departure delay per configuration . . . 52
7.4 Average departure delay per port and TEU category per configuration . . . 53
7.5 Average vessel cost per configuration and scenario without additional vessels . . . 54
7.6 Average vessel cost per port and TEU category per configuration and scenario . . . 55
7.7 Average departure delay per configuration and scenario with additional vessels . . . 55
7.8 Average departure delay per port and TEU category per configuration and scenario . . . 56
7.9 Berth occupancy per configuration and scenario . . . 57
7.10 Sensitivity per class . . . 58

B.1 The number of velocity changes in scenario 1, under the four different configurations per port and TEU category . . . 76

B.2 The average number of velocity changes in scenario 1, under the four different configurations . . . 77
B.3 The number of velocity changes in baseline scenario 1 and global speed optimisation scenarios 1, 2 and 3 per port and TEU category . . . 78
B.4 The number of velocity changes in baseline scenario 1 and global speed optimisation scenarios 1, 2 and 3 . . . 78
B.5 Velocity changes in a TEU 7 vessel . . . 79
B.6 The average velocity in scenario 1, under the four different configurations per port and TEU category . . . 80
B.7 The average velocity in scenario 1, under the four different configurations . . . 80
B.8 The average velocity in baseline scenario 1 and global speed optimisation scenario 1, 2 and 3 per port and TEU category . . . 81
B.9 The average velocity in baseline scenario 1 and global speed optimisation scenarios 1, 2 and 3 . . . 81
B.10 The average waiting time in scenario 1, under the four different configurations per port and TEU category . . . 82
B.11 The average waiting time in scenario 1, under the four different configurations . . . 82
B.12 The average waiting time in baseline scenario 1 and global speed optimisation scenario 1, 2 and 3 per port and TEU category . . . 83
B.13 The average waiting time in baseline scenario 1 and global speed optimisation scenarios 1, 2 and 3 . . . 83

C.1 Realised berth schedule of week 1 and 2 in scenario 1, configuration baseline . . . 86
C.2 Realised berth schedule of week 1 and 2 in scenario 1, configuration single speed optimisation . . . 87
C.3 Realised berth schedule of week 1 and 2 in scenario 1, configuration re-optimise at new departures . . . 87
C.4 Realised berth schedule of week 1 and 2 in scenario 1, configuration global speed optimisation . . . 88
C.5 Realised berth schedule of week 3 and 4 in scenario 1, configuration baseline . . . 89
C.6 Realised berth schedule of week 3 and 4 in scenario 1, configuration single speed optimisation . . . 89
C.7 Realised berth schedule of week 3 and 4 in scenario 1, configuration re-optimise at new departures . . . 90
C.8 Realised berth schedule of week 3 and 4 in scenario 1, configuration global speed optimisation . . . 91
C.9 Cost function without expected delay . . . 92
C.10 Cost function with expected delay . . . 92
C.11 Cost function in which the lowest cost would be achieved by speeding up . . . 93
C.12 All sailing vessels of which the velocities have to be re-optimised . . . 93
C.13 Arrival list with all possible times of arrival, fuel cost and operating cost for the TEU 4 vessel departing from Antwerp . . . 94
C.14 Evaluation of all different orders of arrival of sailing vessels, from left to right: the table in which all different orders are evaluated, the order of the table with the lowest cost, the cost of individual vessels of the order with the lowest cost . . . 94
C.15 Evaluation of all different speed-up possibilities within the order with the lowest costs, from left to right: the table in which all different speed-ups are evaluated, the indices of the vessels' velocity with the lowest cost, and the cost of the velocity of individual vessels of the combination with the lowest costs . . . 95

List of Tables

2.1 Table of class of sea-going vessels, with their average length based on AIS data of vessels travelling to Rotterdam World Gateway in 2024 . . . 9

3.1 BAP classification according to Bierwirth and Meisel [4] . . . 19

4.1 Overview of the baseline strategy and the different speed optimisation strategies . . . 33

5.1 Number of vessel arrivals per class from AIS data and simulation . . . 44
5.2 Total waiting time of all deep-sea vessels . . . 44

6.1 Vessel input parameters; the days of arrival marked with a star (*) are part of scenario 2 and scenario 3 . . . 46
6.2 Fuel consumption parameters used by Brouer et al. [5]. FFE: Forty-foot equivalent unit, s*: vessel design speed in NM/h, Fs*: bunker consumption at design speed in metric tonnes per day . . . 46
6.3 Fuel consumption parameters used in this research. . . . 46
6.4 Changed variables for sensitivity test . . . 49

7.1 Pairwise t-test p-values for the different vessel speed configurations in scenario 1 for the average vessel cost . . . 51
7.2 Pairwise t-test p-values for the different vessel speed configurations in scenario 1 for the average departure delay . . . 52
7.3 Pairwise t-test p-values for the different vessel speed configurations and scenarios for the average vessel costs . . . 54
7.4 Pairwise t-test p-values for the different vessel speed configurations and scenarios for the average departure delay . . . 56

B.1 Pairwise t-test p-values for the sensitivity tests on scenario 1, global speed optimisation for the average departure delay on the TEU 4 and 5 vessels . . . 75
B.2 Pairwise t-test p-values for the sensitivity tests on scenario 1, global speed optimisation for the average departure delay on the TEU 6 and 7 vessels . . . 75
B.3 Pairwise t-test p-values for the sensitivity tests on scenario 1, global speed optimisation for the average departure delay on all vessels . . . 76

C.1 Average number of arrivals per scenario and configuration . . . 96
C.2 Additional vessels in scenarios 2 and 3 . . . 96

Nomenclature

Abbreviations

Abbreviation	Definition
AGV	Automatic Guided Vehicle
AIS	Automatic Identification System
ATA	Actual Time of Arrival
ATB	Actual Time of Berth
ATD	Actual Time of Departure
BACAP	Berth Allocation and Quay Crane Assignment Problem
BACASP	Berth Allocation and Quay Crane Assignment and Scheduling Problem
BAP	Berth Allocation Problem
ECT	European Container Terminals
ETA	Estimated Time of Arrival
FCFS	First-Come, First-Served
FFE	Forty-foot Equivalent Unit
IMO	International Maritime Organization
JIT	Just-In-Time
KPI	Key Performance Indicator
NM	Nautical Mile
RMG	Rail Mounted Gantry
RTG	Rubber Tired Gantry
RWG	Rotterdam World Gateway
RO	Robust Optimisation
SFTW	Sail Fast, Then Wait
SP	Stochastic Programming
TEU	Twenty-foot Equivalent Unit
TRAFALQUAR	Portwise discrete-event berth simulation software
UNCTAD	United Nations Conference on Trade and Development
VSO	Vessel Speed Optimisation

Symbols

Symbol	Definition	Unit
a	vessel fuel consumption parameter	[-]
b	vessel fuel consumption parameter	[-]
$C(v)$	Cost for sailing at velocity v	[$]
C_f	Fuel cost individual vessel	[$]
C_o	Operating cost individual vessel	[$]
C_t	Total cost individual vessel	[$]
C_w	Waiting cost individual vessel	[$]
c_f	Coefficient of fuel cost	[$ per metric tonne]
c_o	Coefficient of operating cost	[$ per hour]

Symbol	Definition	Unit
c_w	Coefficient of waiting cost	[$ per hour]
D	Distance	[NM]
d	departure delay	[hour]
$\tilde{f}$	Fuel consumption at $\tilde{v}$	[metric tonne per day]
l_j	Vessel length	[m]
l_q	Quay length	[m]
Q	Berth occupancy	[-]
R	Remaining distance	[NM]
$RC(v)$	Cost over remaining distance at velocity v	[$]
T	Total simulation time	[hour]
t_o	Operating time for a vessel	[hour]
t_s	Service time for a vessel	[hour]
t_w	Waiting time for a vessel	[hour]
v	Velocity	[NM per hour]
$\tilde{v}$	Velocity at which fuel and operating cost are minimal	[NM per hour]

1

Introduction

This chapter provides background on the topic of dynamic vessel speed optimisation. In section 1.1, the topic is placed in context. In section 1.2, the problem statement is introduced, and the main stakeholders are identified. Through the problem statement, the case study is introduced. In section 1.3, relevant literature on the problem is set out. In section 1.4, the research gap is identified, and the scope of the paper is determined. In section 1.5, the main research question and the sub-questions used to answer the main research question are introduced. The outline for the report is shown in section 1.6.

1.1. Introduction

Global maritime shipping plays a crucial role in the world economy [29]. In 2024, 80% of the global volume of international trade was carried by maritime transport. The total volume of maritime transport is expected to increase by over 2% annually [29]. International shipping accounts for approximately 3% of global greenhouse gas emissions [13, 29]. The efficiency of port operations is crucial for maintaining a sustainable form of maritime transport.

Several factors contribute to efficient maritime operations. For a container terminal, high quay utilisation leads to efficient use of the available equipment. However, with high quay utilisation, the probability that incoming vessels have to wait increases. Waiting time decreases the port's service level. Through robust berth planning, quay utilisation and the terminals' service level can be balanced. An effective berth schedule leads to an increased port throughput and, in turn, increased port revenue [9].

However, the berth schedule has to be constantly updated, as ship arrival times are not deterministic. Container vessels visit ports in succession. If a delay occurs at the first port, the arrival times for all subsequent ports must be updated. Besides delays in the harbour, delays on route caused by technical issues or weather conditions also change the estimated arrival time of the vessel.

Container terminals often serve vessels based on a first-come, first-served (FCFS) basis. Each vessel departs from the previous harbour at its optimal speed, without considering other vessels or possible berthing slots. As it is possible that no berthing slot will be available upon arrival, the vessel will have to wait. This practice is also known as *Sail Fast, Then Wait* or *Rush-to-Wait* and is an energy-inefficient approach [28]. It would be more energy-efficient to sail at a slower speed and be serviced directly, rather than cruising at a higher speed to minimise sailing time and waiting for anchorage until a berthing slot becomes available.

In practice, berth schedules rely on uncertain or outdated Estimated Times of Arrival (ETAs). When a vessel arrives later or earlier than expected, it can disrupt the berth schedule, leading to increased port congestion, idle time, and higher fuel consumption. As a result vessels may spend 5 to 10% of their operational time waiting to get into port [14].

1.2. Problem statement

To ensure that vessels do not have to wait before being able to berth, the International Maritime Organisation (IMO) published the Just-in-Time (JIT) Arrival Guide in 2020 [14]. The goal of the Just in Time arrival guide is to allow the vessel to sail at an optimal velocity, such that it does not have to wait at anchorage. The Just-in-Time arrival guide states that waiting time at the anchorage point can be reduced through improved communication between vessels and the terminal.

When a vessel does not comply with the JIT arrival guide, it will set sail at a velocity that minimises operating and fuel costs over the sailing leg, without accounting for waiting costs. The operating costs of a vessel include crew costs and chartering/depreciation costs. These costs also occur when a vessel has to wait at the anchorage point, which is referred to as the waiting cost. If the waiting time is known in advance, a new optimal velocity of the vessel can be devised. This new optimal velocity will be lower, such that the vessel does not have to wait at the anchorage point. The lower velocity will, in turn, lead to lower fuel costs, as sailing at a lower velocity decreases the fuel consumption per distance travelled.

Besides benefiting vessel operators through lower service costs, vessel speed optimisation also benefits other stakeholders around the container terminal. Decreased waiting times lead to a higher service level at a container terminal. Vessel speed optimisation also leads to more accurate arrival predictions, enabling the container terminal to service more vessels.

Local governments would like to decrease the pollution in the water near the terminal. Vessels waiting at the anchorage point have to remain operational. By decreasing the number of vessels waiting at the anchorage point, pollution may decrease. Additionally, decreasing the vessel's speed on the last sailing leg reduces pollution during the voyage.

Port authorities of busy ports want to avoid congestion in their waterways. Deep-sea container vessels are hard to manoeuvre. Congestion in the port's waterways increases waiting times for all vessels and the risk of accidents. Knowing in advance when a vessel arrives allows the port authority to plan the vessel's manoeuvring without interfering with other vessels.

1.2.1. Introduction to the case study

The case study for this thesis will be performed at Portwise. Portwise has commissioned this study on dynamic vessel speed optimisation to allow vessels to steam efficiently. Portwise is a logistics consultancy and simulation company specialising in port and terminal processes. They have created a simulation software in eM-Plant called TRAFALQUAR. TRAFALQUAR is a discrete event simulation, where at each step a new situation is calculated. It simulates the arrival and berthing of vessels, the container handling, and the undocking and departure of vessels. Currently, TRAFALQUAR is used to test the possible throughput of a container terminal. In TRAFALQUAR, vessels arrive at an anchorage point of the terminal and wait until a berth becomes available. In this study, an extension for TRAFALQUAR will be made, which allows vessels to be simulated from the moment they depart their previous port.

AIS (Automatic Identification System) Data of vessels arriving from 1-1-2024 to 31-12-2024 to Rotterdam World Gateway (RWG) at the Port of Rotterdam has been provided by Teqplay. RWG has been selected as it has the largest selection of first calls in the Port of Rotterdam. Based on the provided AIS data, vessels that do not arrive at RWG as their first call typically have a capacity of less than 4000 Twenty-foot Equivalent Units (TEU).

The AIS data will be used to validate the extension made in TRAFALQUAR, by comparing the average handling time and waiting time of groups of vessels with these values from the simulation.

1.3. Literature review

The literature review on the problem will be conducted in two parts: berth allocation problem (BAP) under uncertainty, and vessel speed optimisation strategies. The berth allocation problem under uncertainty entails a situation where variables, such as time of arrival or handling time, are not known in advance. For instance, the time of arrival can vary due to delays at previous ports.

1.3.1. Berth allocation under uncertainty

On the berth allocation problem under uncertainty, extensive research has been conducted. Rodrigues and Agra reviewed the literature on this topic in 2022 [23]. In their paper, they discuss the methods for modelling uncertainty and the approaches to handling the berth allocation problem in practice. In their research, four main modelling approaches to uncertainty are identified;

- **Stochastic programming**: The uncertain parameters are defined by known probability functions. Stochastic programming optimises the expectations of functions involving random variables.
- **Robust optimisation**: The probability of uncertain parameters is not exactly known, or does not exist at all. Robust optimisation often tries to optimise the worst-case scenario. Robust optimisation is often used when a more risk-averse approach to the uncertainty is taken.
- **Fuzzy models**: The uncertain parameters belong to a fuzzy set. The most common fuzzy set is a triangular set, based on three numbers: a minimum, a maximum, and an expected value.
- **Deterministic models**: A Deterministic model plans the operations when a disruption from the schedule occurs. The objective is either to minimise the cost of rescheduling or to minimise the deviation from the previous schedule.

These four models can be implemented in a BAP. The use of a model depends on knowledge of uncertainty and the approach to risk. Stochastic programming is the most commonly used method, especially when the uncertainty is known. Robust optimisation is used when a more risk-averse strategy is taken. Fuzzy models are not often used.

- **Proactive**: In a proactive approach, all the decisions made for the planning are made in advance. Generally, a proactive approach is adopted, such that the solution minimises the average cost (under a cost function) while remaining feasible for all realisations of the uncertainty.
- **Reactive**: In a reactive approach, decisions are made when the realisation of uncertain parameters becomes certain. In the reactive approach, a simple baseline plan is made. When deviations from the baseline plan occur, the goal is to minimise the cost of rescheduling and to minimise the deviation from the previous schedule. Deterministic models are often used in combination with this approach.
- **Proactive/reactive**: In a proactive/reactive approach, a baseline plan is made, while taking uncertainty into account. Then, if an uncertain parameter becomes known and would disrupt the baseline plan, the planning is adjusted. A proactive reactive/approach tries to minimise both the cost of the baseline schedule and the cost of rescheduling the operations.

These approaches differ on three key attributes: the expected value of the uncertainty, the number of static decisions, and the number of recourse actions taken. Proactive approaches are suitable when uncertainty is low, allowing the baseline plan devised at the start to remain feasible. In this approach, only static decisions are made at the start, and no recourse actions are possible during the realisation of uncertainty. As the proactive approach only has to be calculated once, the computational complexity of the model can be larger. Reactive approaches are suitable when uncertainty is high, allowing for recourse actions to be taken and a better plan to be formulated once the uncertainty becomes known. In this approach, recourse decisions are taken when deviations occur. As these corrective actions need to be taken as soon as the deviations occur, the problem must be calculated as quickly as possible, so that the new berthing plan can be updated as soon as possible. Proactive/reactive approaches balance the advantages and disadvantages of both, as they can handle more uncertainty in parameters than proactive approaches, which allow for re-planning. However, since recourse actions still have to be taken after a disruption, the problem must be calculated as soon as possible. When dealing with uncertainties that become better understood over time, proactive/reactive and reactive approaches are more suitable.

Besides the two-stage method described in the list under proactive/reactive, a proactive/reactive model can also be solved using a rolling horizon. In a rolling horizon strategy, the problem is solved iteratively. The planning horizon is divided into decision cycles, for which the BAP will be planned. In this decision cycle, the first execution period is implemented. The schedule is updated after the first

execution period, and the decision cycle for the updated scenario is calculated. The rolling horizon enables continuous re-optimisation based on updated information. A rolling-horizon strategy is a good fit for the BAP, as the arrival time of a vessel becomes more accurate when updated regularly, particularly for vessels close to the harbour.

1.3.2. Berth allocation combined with vessel speed optimisation

Vessel speed optimisation (VSO) can also be integrated into the BAP. Through vessel speed optimisation, vessel arrival time can be optimised, resulting in decreased waiting time and fuel consumption.

Golias et al. have integrated VSO for the first time in the BAP [10]. They treat the arrival time of vessels as a decision variable, and optimise vessel arrival times. They do not directly minimise fuel consumption. By reducing waiting times, the fuel consumption is indirectly reduced, as the vessel no longer produces emissions in idle mode. Du et al. focus on minimising fuel consumption directly [8]. To capture the non-linear behaviour of fuel consumption, they formulate the BAP as a mixed-integer second-order cone programming (MISOCP) problem; in addition to fuel consumption, they analyse the emissions resulting from it. By directly minimising fuel consumption, the emission reduction over the sailing leg can be calculated precisely. Both for the sailing line and the terminal, this is a win-win situation, as the sailing line can save fuel on the sailing leg, while the terminal reduces overall waiting times. In the study by Du et al, the emissions during the waiting time at the anchorage point are not taken into account.

1.4. Research gap and scope

While berth allocation and vessel speed optimisation have been studied separately and in combination, limited research has investigated dynamic vessel speed optimisation for heterogeneous vessels arriving from ports at varying distances under realistic terminal operating conditions.

The scope of the research is to investigate the influence of vessel speed optimisation on vessel performance and on the berth occupancy and service levels at a container terminal. The model is dynamically updated based on information currently available at the terminal. This research focuses on deep-sea vessels arriving at a single deep-sea terminal in a multi-terminal port. Seaside operations in the container terminal are assumed to be unaffected by the yard operations.

1.5. Research question and sub-questions

The research gap mentioned in the previous section can be filled by answering the following main research question.

”What is the effect of dynamic vessel speed optimisation on the performance of container vessels and berth occupancy of a deep-sea container terminal?”

To help answer the main research question, the topic has been divided into the following sub-questions.

1. How do vessel types, vessel operations, and seaside operations influence waiting times at the terminal?
2. How can vessel speed optimisation be incorporated into the berth allocation problem?
3. Which vessel speed optimisation strategies can be identified and structured for use in simulation?
4. How can vessel speed optimisation be implemented in a discrete event simulation model that realistically represents real-life vessel operations?
5. How do different vessel speed optimisation strategies affect the performance of container vessels?
6. How can vessel speed optimisation influence the berth occupancy of a container terminal?

1.6. Outline of the report

To help answer these sub-questions, different research methods are used. For each of the mentioned research questions, the method is displayed.

Sub-question 1 will be answered by conducting a system analysis of the arrival of container vessels. This analysis will focus on three key aspects of the arrival of container vessels: first, the types of vessels, second, the operations of vessels, and finally, the operations at the seaport. The system analysis will help determine processes, performance indicators, and requirements of the model.

Sub-question 2 will be answered by conducting a literature review. Scientific insights into uncertain arrival times and vessel speed optimisation will be used to develop an overview of how the vessel speed optimisation can be implemented efficiently into the berth allocation problem.

Sub-question 3 will be answered by conceptual modelling. Based on the insights obtained from the literature review in sub-question 2, vessel speed optimisation strategies will be systematically defined. These strategies are conceptualised around three main factors: the moment of decision-making, the information available to the vessels at that moment, and the optimisation objective. The conceptualisation is used to translate the strategies into defined simulation scenarios.

Sub-question 4 will be answered through design and modelling in TRAFALQUAR. The model will be validated using AIS data from vessels arriving in 2024 at RWG, ensuring that the model's representation of waiting times and time spent at the quay is comparable to the results from the AIS data. TRAFALQUAR is chosen as the modelling environment for this research because it is an established and validated simulation model used for terminal performance analysis. Developing a new berth allocation or simulation tool would require extensive validation and lack the operational capabilities TRAFALQUAR already provides. Using the existing simulation software ensures the berth assignment results are realistic, allowing the research to focus on implementing vessel speed optimisation.

Sub-questions 5 and 6 will be answered through experiments conducted on the validated model. In sub-question 5, different speed optimisation strategies will be tested. The performance of the vessels will depend on their total costs and the departure delay. These strategies will be compared to the current Sail Fast, Then Wait situation. In sub-question 6, the best-performing speed optimisation strategy will be tested to see how the model reacts under different scenarios with additional vessels. By analysing the outcomes of these experiments, insights will be gained into how vessel speed optimisation strategies can contribute to more efficient container terminals.

2

System Analysis on the Arrival of Container Vessels

In this chapter, the sub-question, ”How do vessel types, vessel operations, and seaside operations influence waiting times at the terminal?” is answered. This chapter answers the sub-question by introducing the types of vessels and analysing the two subsystems separately. First, the system is defined. Then, the three main types of vessels considered in this research are introduced, with their roles in international trade, operational characteristics, and technical specifications. Then, the subsystems of seaside operations and vessel operations are analysed separately. These analyses are based on academic literature, technical reports, the Portwise Academy, and one year of AIS data on the Port of Rotterdam, provided by Teqplay. Using these separate analyses, the combined system is formed and analysed. Finally, the chapter concludes by discussing how both subsystems influence waiting time.

2.1. System definition

The main topic of the research is the influence of speed optimisation on vessels arriving at a container terminal. First, it is important to correctly define the overall system. The system mainly consists of two subsystems. Vessels arriving at the anchorage point of the container terminal and the loading and unloading of vessels at the seaside of the container terminal. The schematic side view of the operations on the container terminal by Steenken et al. is shown in Figure 2.1. In this overview, two main areas of operations are defined: quayside and landside. The quayside is where vessel interactions with the container terminal take place. At this location, vessels arrive at berth, are loaded and unloaded, and depart for the next port.

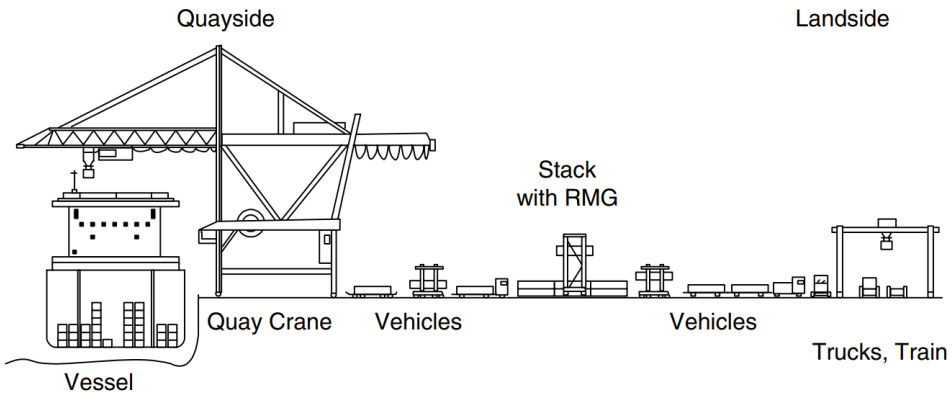


Figure 2.1: Schematic sideview of a container terminal [27]

Often, a third area of operations is defined. The quayside interacts with the landside via the stack, where containers are stored. This can also be called the yard of the container terminal. In Figure 2.2, the correct scale of these three different areas of operation is shown by an aerial overview of ECT Delta in the Port of Rotterdam. In this figure, the seaside operations are marked blue, the yard operations are marked green, and the landside operations are marked yellow. Quayside and seaside operations are used interchangeably.

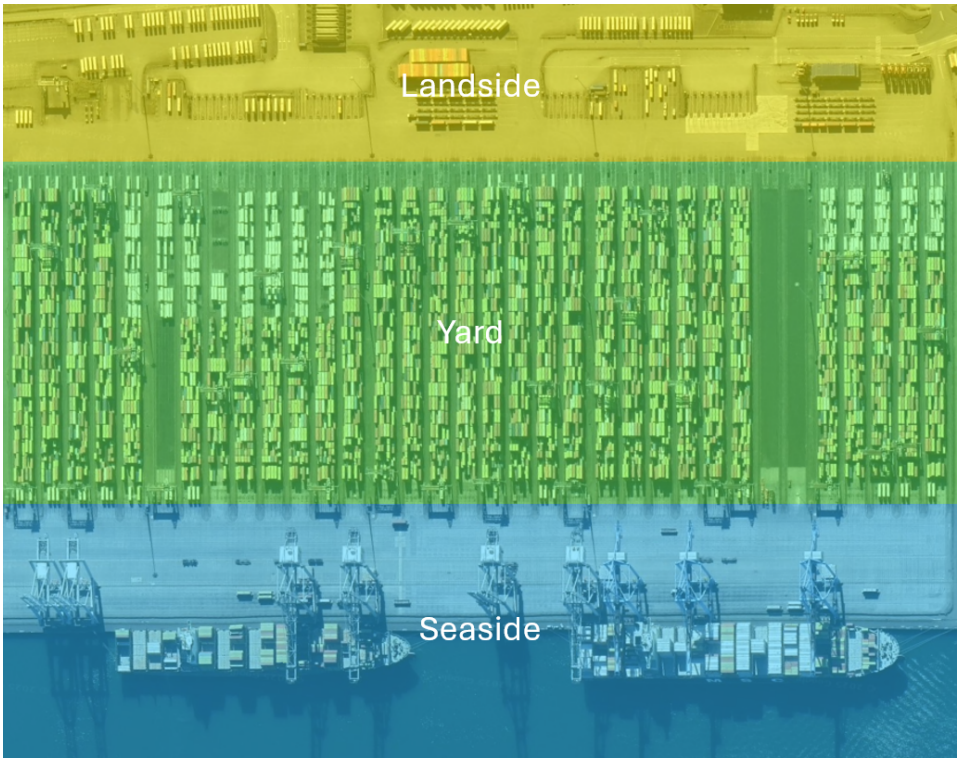


Figure 2.2: Areas of operation in a container terminal (ECT Delta, Port of Rotterdam)

The three areas of operations interact in the schematic way of Figure 2.1. At the quayside, quay cranes load and unload containers from vessels. The containers loaded and unloaded from the vessel are transported from and to the yard by terminal vehicles, such as automatic guided vehicles (AGVs). In the yard, these containers are stored while waiting for their pickup. Dependent on the size of the terminal, these containers may need special equipment such as rubber-tired gantry cranes (RTGs) or rail-mounted gantry cranes (RMGs) to be stacked on top of one another. From and to the yard, containers are also transported to the landside. At the landside, containers can enter and exit the container terminal via trucks and trains.

The main topic of this research is the interaction between the arriving vessels and the container terminal at the seaside of the terminal. According to Carlo et al. [6], the primary research topic regarding the seaside operations of a container terminal is the berth allocation problem (BAP), which aims to assign incoming vessels to berths. The berth allocation problem can be extended to include the quay crane scheduling problem (BACSP) and the quay crane assignment problem (BACAP), which, in turn, schedule and assign cranes to berthing vessels.

The type of vessel determines the number of containers it transports. The three main vessel types are deep-sea, short-sea, and inland barges. Deep-sea vessels transport a large number of containers intercontinentally, while short sea vessels transport a smaller number over relatively shorter distances over seas, and barges transport a smaller amount inland.

2.2. Vessel types

There are three vessel types, each serving different roles in the industry. From large to small, there are deep-sea vessels, short-sea/feeder vessels, and inland barges. Due to economies of scale, the larger a vessel is, the higher the total operating cost, but the lower the cost per container. Teqplay has gathered information on all vessels travelling to the Port of Rotterdam by sea from 2024 to 2025.

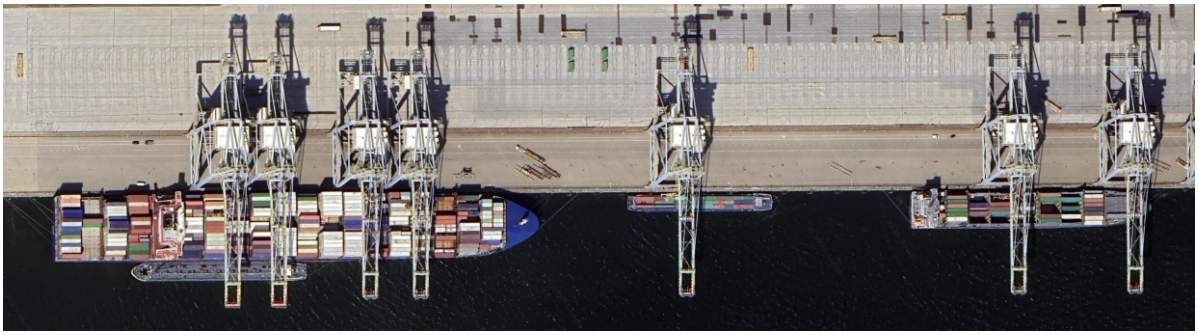


Figure 2.3: Three different vessel types at Rotterdam World Gateway (Google Earth 23-07-2024), Left to right, deep-sea vessel, inland vessel, feeder vessel

2.2.1. Deep-sea vessels

The largest vessels are the deep-sea vessels. In this research, they are defined as vessels capable of transporting 4000 or more Twenty-foot Equivalent Units (TEUs), the standard size of one container. According to the Port of Rotterdam, a vessel is considered to be sailing in a deep-sea service when it operates on a scheduled service, where the timetable is announced 28 days before arrival, and when the sailing area is not limited to Europe, the Mediterranean Sea, the Black Sea, Morocco, the Canary Islands, Madeira and the Cape Verde Islands [20]. Deep-sea vessels are used in liner shipping to connect major ports on different sides of the world. In liner shipping, vessels operate on fixed schedules on predefined ocean routes [18]. This can be compared to the operations of a bus. A bus can travel from one city to another, stopping at fixed points along a fixed schedule. Finally, it returns to its original starting point, and it begins the round again. In the context of liner shipping, a vessel departs from Asia and stops at a fixed schedule in multiple ports in Europe before going back to Asia and visiting multiple ports there. A shipping company often operates multiple shipping lines. On a single shipping line, multiple vessels of the company operate simultaneously to ensure that a port in the shipping line is served once a week.

There are different classes of deep-sea vessels. The largest deep-sea container vessels are 400 meters long and can carry over 20000 TEU. The vessel size is constrained by the route it follows. Common vessel classes include the Panamax and the Suezmax, which are constrained by their width, draft, height, and length, allowing them to still pass through the respective canal. The larger a vessel is, the larger its call size and the longer it takes to service. The call size of a vessel is the total number of containers loaded and unloaded at the terminal. The AIS data from Teqplay does not assign call sizes to vessels; it displays the time the vessel has spent at the quay. In the data, each vessel is categorised by TEU Class. It is important to note that the number of TEU corresponds to the number of containers a vessel can carry, not to its call size. For each class of the deep-sea vessels and the short-sea/feeder vessels, the average length is displayed in Table 2.1. According to Teqplay data, deep-sea vessels typically call only at one terminal in Rotterdam.

Due to their size, deep-sea vessels are the most expensive type of vessel to operate. Therefore, it is expensive to have them wait at an anchorage point. Terminals often give them priority over smaller vessels. Even with priority, these vessels still have to wait if other deep-sea vessels occupy the quay.

2.2.2. Feeder / short-sea vessels

The smaller seagoing vessels are the feeder/short-sea vessels. In this research, feeder vessels are defined as those carrying between 250 and 4000 TEU and travelling to other ports by sea. According to the Port of Rotterdam, a vessel is considered to be sailing in a feeder/short-sea service when it

Class	2	3	4	5	6	7
Number of TEU	250-1000	1000-4000	4000-8000	8000-12000	12000-16000	16000-30000
Role	Short-Sea	Short-Sea	Deep-Sea	Deep-Sea	Deep-Sea	Deep-Sea
Average Vessel Length (m)	135	165	285	345	365	400

Table 2.1: Table of class of sea-going vessels, with their average length based on AIS data of vessels travelling to Rotterdam World Gateway in 2024

sails on a scheduled service, and when the sailing area is limited to Europe, the Mediterranean Sea, the Black Sea, Morocco, the Canary Islands, Madeira, and the Cape Verde Islands [20]. Feeder vessels can distribute the containers of a deep-sea vessel to nearby container terminals. For the locations a feeder vessel serves, it is often not feasible to use a deep-sea vessel due to the smaller number of containers that it can carry. Feeder vessels are smaller than deep-sea vessels and have a lower draft, and can therefore enter terminals that are too shallow for deep-sea vessels. The operating cost of a single vessel is significantly lower than that of a deep-sea vessel. However, due to economies of scale, the cost per container is higher for feeder vessels. Although more feeder vessels arrive at a port than deep-sea vessels, the throughput (the total number of containers handled) of the container terminal is mainly influenced by the deep-sea vessels. The feeder vessels often call at multiple terminals in a multi-terminal port. So they can also redistribute containers from one terminal to another.

2.2.3. Inland barges

Ports are generally located at sea or near a river. The river can be used to transport containers to inland destinations. When a location inland requires a larger number of containers, it is generally more cost-effective to use inland barges to transport the containers than to move them by truck. The size of inland vessels can vary, as they may be pusher units handling multiple barges or dedicated inland vessels. Inland container terminals, such as the Port of Duisburg, rely on inland vessels travelling from seaports like Rotterdam.

In addition to transporting containers inland, barges are also used to transport containers from one terminal in the same port to another.

The most common vessel type in a terminal is a barge. Altogether, they move a similar number of containers as all short-sea/ feeder vessels combined. However, deep-sea vessels account for a larger number of containers than both combined.

Barges often visit multiple terminals in the same port. In the Port of Rotterdam, Nextlogic is used to plan berths for inland barges. Nextlogic plans a cycle of terminal visits for a barge operator, based on the terminals the barge must visit and the future berth capacity of each terminal, ensuring optimal use of barge capacity. From Nextlogic, vessel operators get the rotation and schedule of the terminals they have to visit [19].

In Figure 2.4, two barges are located at the quayside side by side. This allows them to be serviced simultaneously. Although this increases the service time of a single barge, it can increase the productivity of the crane. As this allows the crane to load more containers, it can be loaded before it has to move to switch to another bay of the barge. The higher productivity of the crane, in turn, leads to a higher overall productivity of the terminal. For larger vessels, double berthing is generally not possible. The combined width of the two vessels must be smaller than the length of the quay crane's beam.

2.3. Seaside operations

This research focuses specifically on the seaside operations of a container terminal. The seaside operations of a container terminal consist of a number of different processes. These processes are shown in Figure 2.5.

As shown in the figure, these processes interact with each other. For instance, the assignment of a berth to a vessel depends on the berth planning. The transportation of a container from the yard to

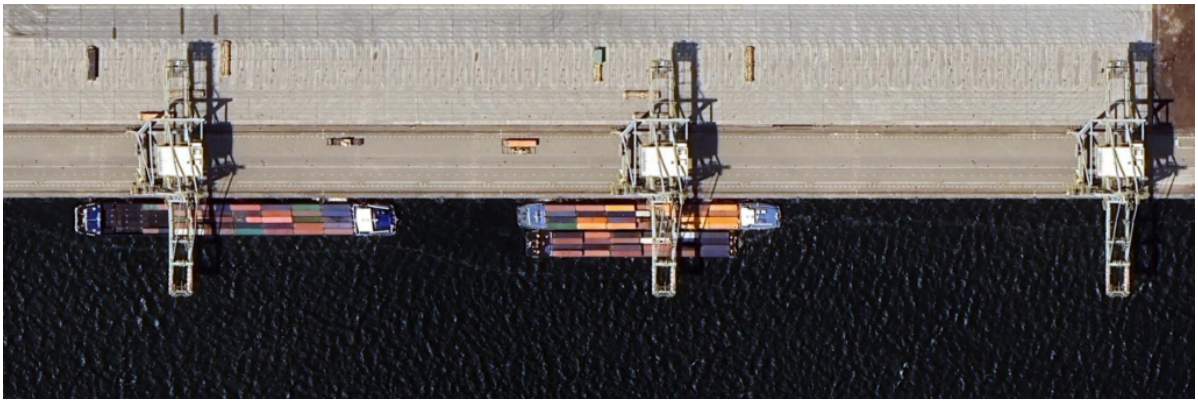


Figure 2.4: Double berthing of two barges at the barge quay of Rotterdam World Gateway (Google Earth 23-07-2024)

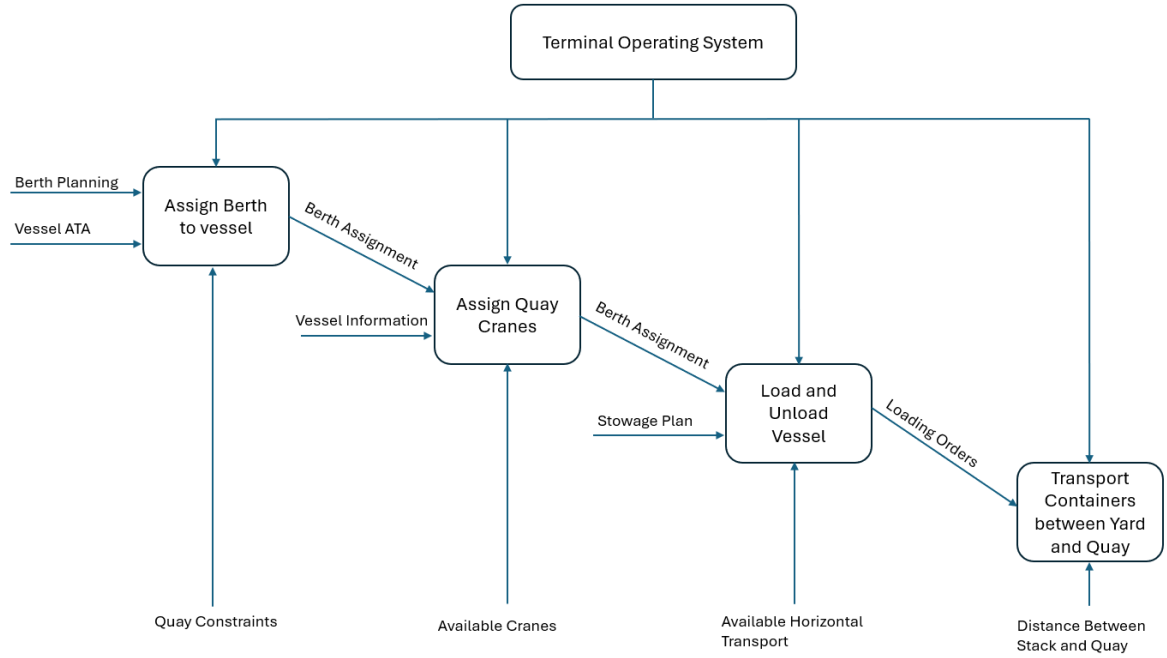


Figure 2.5: Seaside operations

the quay depends on the loading and unloading of the vessel.

The following performance measures for the seaside operations have been identified by Hartmann et al. [12].

1. Container Throughput
2. Throughput of Ships
3. Average and Maximum Ship Waiting Time
4. Percentage of Waiting Ships
5. Berth occupancy

Where the container throughput is the overall throughput in TEU of the container terminal, the throughput of ships, average and maximum ship waiting time, percentage of waiting ships is given per ship class, and the quay utilisation is the average percentage of the berth occupied during the simulation time.

2.3.1. Analysis of the seaside operations

The seaside operations of a container terminal are analysed in this paragraph using a blackbox. A blackbox is a tool in which a system transforms inputs into outputs. The requirements describe the system's physical constraints. Performance indicators are used to determine how well the system operates. In the blackbox displayed in Figure 2.6, the seaside operations are shown. In this blackbox, the seaside operations, described in Figure 2.5, 1-2 are performed, while operations 3 and 4 are combined. These are combined by using a single number for the combined productivity of a quay crane and its terminal tractors, rather than using a stock number of productivity for the crane and a separate number for the number of containers that can be transported to the yard. This is because transportation to the yard can depend on multiple factors, such as distance to the stack and terminal tractor congestion. To ensure this number is included, the quay crane productivity used is the average per-hour productivity during operation. Performance measures 1 and 2 are taken as inputs to the blackbox, while 3 and 5 are taken as performance measures. The entries for the inputs, outputs, requirements and performance indicators are discussed in the following paragraphs.

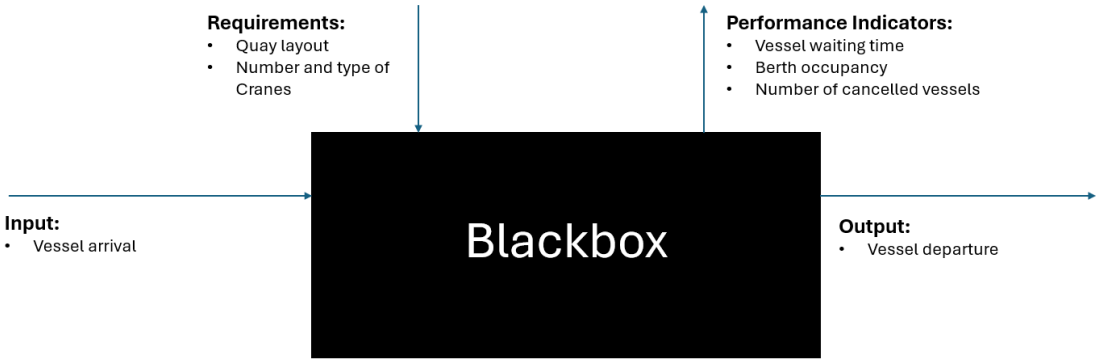


Figure 2.6: Blackbox of the seaside operations of a container terminal

Input

The only input of the blackbox is vessel arrival. For arriving vessels, a few parameters must be defined. The moment at which a vessel arrives is defined by its ETA. However, a vessel does not always arrive at its planned ETA. A deviation for each vessel type can be defined. This deviation can be based on historical data. Through this deviation, the actual time of arrival of a vessel can be stochastically defined. The call size per vessel is not fixed. Although vessels of roughly the same size can carry a similar amount of containers, the actual call size varies. For each vessel, a distribution of the number of containers can be provided. Furthermore, the number of containers per bay on a vessel must be defined, as this influences the number of cranes that can work on a vessel and the time it takes a crane to clear a bay. The vessel's call size is also dependent on the seasonal pattern; in some months, the terminal handles more containers than others. The actual call size of a vessel can be defined by combining the distribution and the seasonal pattern.

Based on the ETAs of all incoming vessels, the container terminal defines a vessel arrival schedule. In this schedule, the number of vessels expected to arrive each day of the week is defined. For each vessel arrival, the vessel lengths and the average call size are defined. A proforma berth schedule of a container terminal with 2 berths, 77 barges, 24 short-sea vessels, and 6 deep-sea vessels is shown in Figure 2.7. In this figure, the x-axis represents the location of the berth on the quay, and the y-axis represents time. Each vessel is represented as a block, and the number of cranes expected to work on a vessel is displayed. In the proforma berth schedule, vessels may overlap, indicating that two vessels arrive at the same time. This is not possible in the realised berth schedule, as two vessels cannot be served at the same time on the same quay.

Output

The only output of the black box is vessel departure. When the vessel leaves the system, its data is stored. The data can be summarised as the vessels' berth location. The berth location is

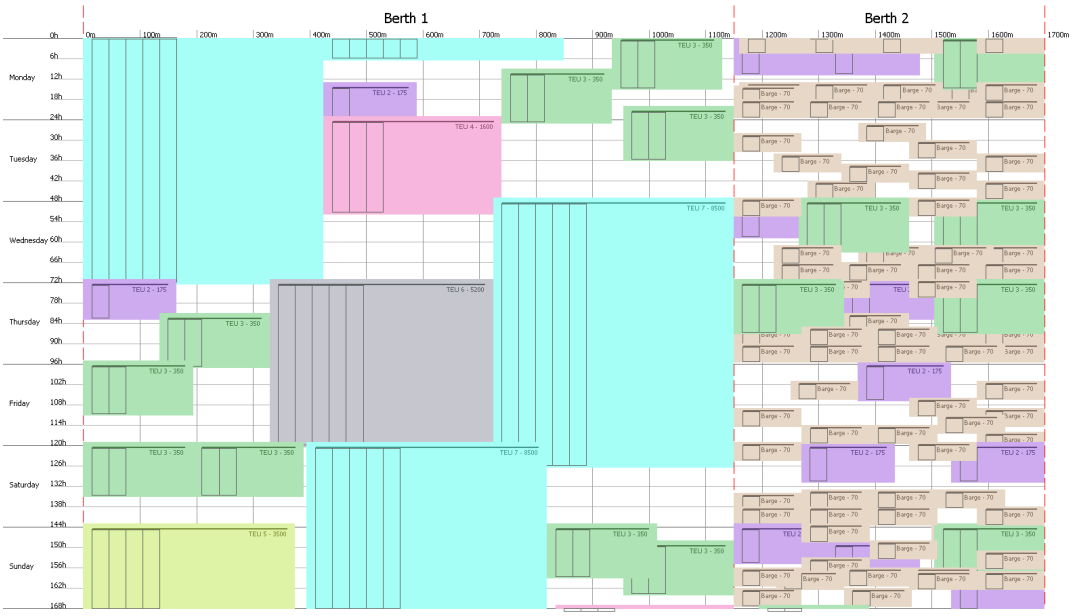


Figure 2.7: Proforma berth schedule

two-dimensional. On the time axis, and on the quay axis. The realised vessel berth location differs from the vessel berth location in the proforma schedule. Vessels can arrive later, leading to a deviation in the time axis. When arriving later, the position may also differ from the originally planned one. This is counterproductive for the container terminal, as the containers' yard locations depend on the planned positions of the vessels on the quay. Deviations from the planned position increase travel in horizontal transport at the container terminal, reducing the system's effective productivity. The realised berth schedule is shown in Figure 2.8. When a vessel cannot start its berth immediately, it has to wait; waiting is indicated by the red bar. In the realised schedule, the vessel's call size is also defined, which is the number at the top right of a vessel block.

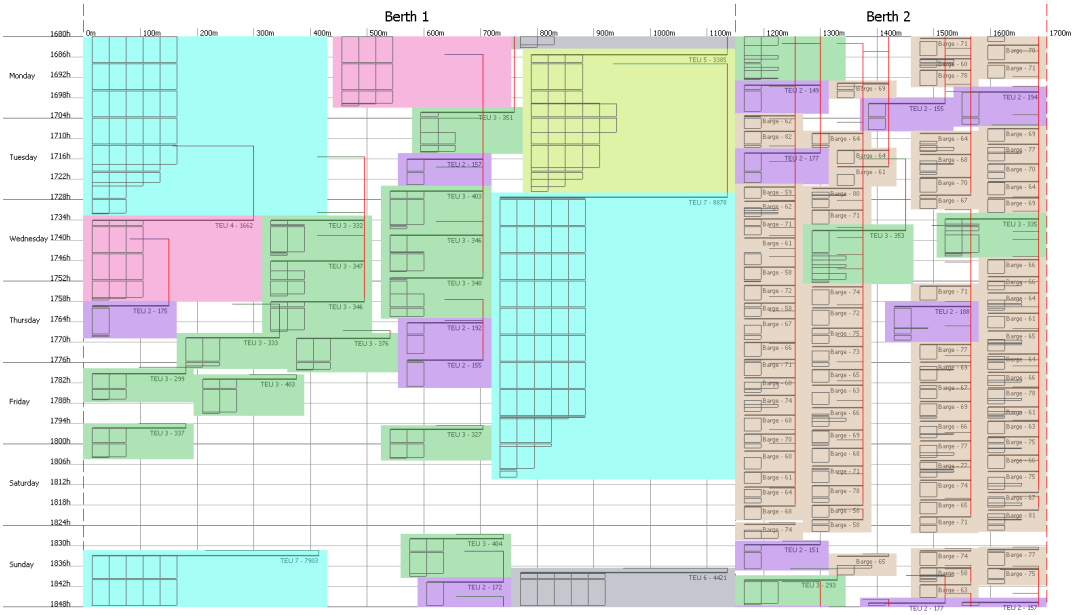


Figure 2.8: Realised berth schedule

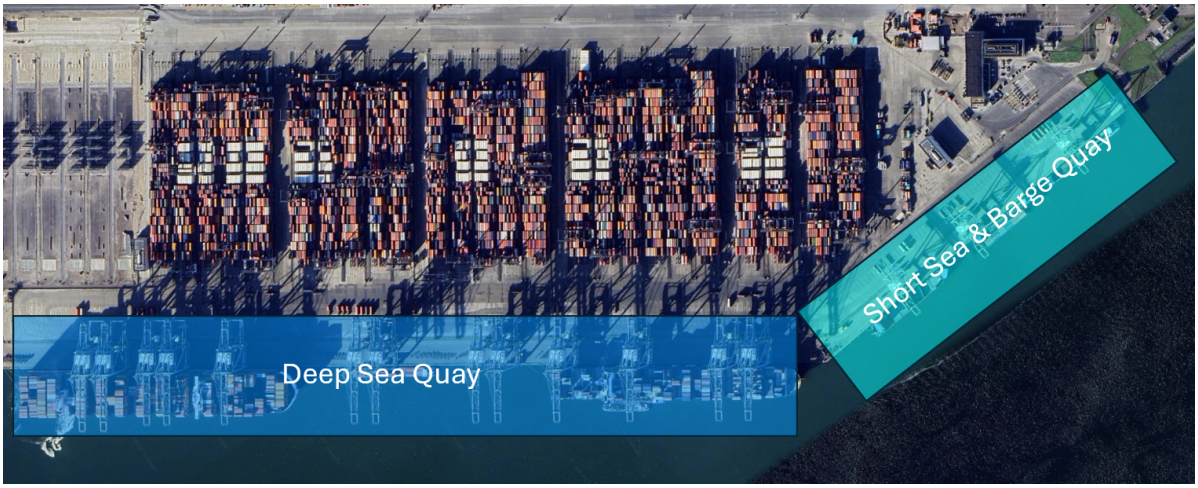


Figure 2.9: Deep-sea and feeder/short-sea quay at RWG

Requirements

Quay layout

The first requirement is knowledge of the quay layout of the simulated terminal. The quay layout consists of three parts: the number of different quays, the length of each quay, and the depth of each quay. The length of the quay determines which types of vessels can berth simultaneously at the terminal. While the depth of the quay constrains the type of vessels that can berth. A terminal can have several parts of quay, some of which are more shallow than others. These shallower waters are often only for inland barges or smaller feeder vessels. Rotterdam World Gateway (RWG), for instance, has one deep-sea quay and one short-sea/ barge quay. The main difference between the two is that the deep-sea quay has a depth of 20 meters, allowing for the largest vessels to be able to arrive, while the short-sea /barge quay is only 11 meters deep [7]. The distance from the short-sea quay to the yard is longer than that from the deep-sea quay, so transport vehicles travel farther, which could slow the cranes on the short-sea quay. When necessary, short-sea vessels and barges can also arrive at the deep-sea quay; however, deep-sea vessels can not arrive at the short-sea quay due to the above-mentioned draft restrictions.

Number and type of cranes

The second requirement is on the number and types of cranes at the terminal. The number of cranes affects a vessel's loading and unloading productivity. The more cranes that are available, the faster a vessel can be unloaded. The type of crane is also important, as quay cranes can have height and width restrictions. This can cause cranes to be unable to handle vessels with containers stacked above a certain height or vessels wider than a certain width. The larger the vessel, the larger the container crane has to be. The operations in the yard influence the quay, as the number of available transport equipment directly influences the speed of loading and unloading the vessel. As yard operations are not taken into account, crane productivity should be scaled to the average crane productivity during operations, not the maximum theoretical achievable productivity.

Performance Indicators

Vessel waiting time

The first performance indicator is the vessel waiting time. The vessel waiting time indicates the service level of a container terminal. The larger the expected vessel waiting time, the less likely a vessel is to visit a terminal.

Berth occupancy

The second performance indicator is the berth occupancy. The berth occupancy indicates the throughput of a container terminal. The higher the berth occupancy, the more effectively terminal resources are utilised. However, higher berth occupancy in turn leads to more vessel waiting time.

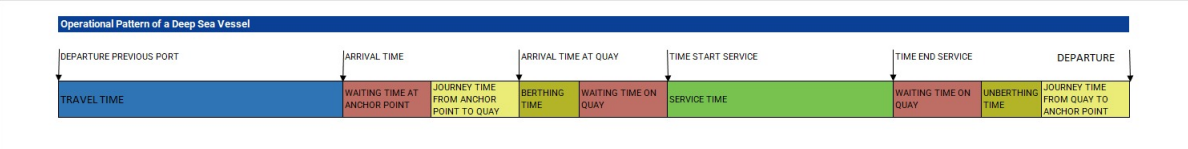


Figure 2.10: Schematically drawn vessel operations

Number of cancelled vessels

The final performance indicator is the number of vessels that have been cancelled. If vessels have to wait too long before being able to berth, they will deliver their containers to competing terminals in the region. To avoid vessels from cancelling, container terminals need to finish servicing a vessel before its due date.

2.3.2. Conclusion

The analysis of the seaside operations of a container terminal shows that waiting times at the terminal primarily result from deviations in arrival and service times, rather than physical capacity. The container terminal has sufficient berth length and crane capacity to handle all the incoming vessels. However, the variability in vessel arrivals and handling times creates peaks that exceed the available berth capacity at specific moments. When too many ships arrive at the same time, vessels will queue at the anchorage point. When too many cranes are required at once, the vessel's service time is extended.

The deviations in ETA reduce the effectiveness of the berth planning. Berth planning often includes buffers in the vessel's arrival time to reduce the impact of deviations. However, the buffers allow the terminal to utilise less of its full capacity. If uncertainties in the vessel ETA are reduced, more accurate berth planning can be made, which in turn reduces waiting times.

2.4. Operational pattern of deep-sea Vessels

In this section, the operational pattern of deep-sea vessels will be analysed. The analysis is done based on the problem defined in the JIT arrival guide [14]. The JIT arrival guide aims to reduce vessel operating costs and fuel consumption by decreasing the number of vessels waiting at the anchorage point and by optimising their velocity. This optimisation is only possible for vessels that arrive directly from another seaport to the next terminal. For vessels travelling from a terminal in the same port, no speed optimisation is possible. The JIT arrival guide states that container vessels are suitable for speed optimisation because they travel at high speed. In this section, a schematic representation of a vessel's operation will be given. From this operational pattern, a black box is defined.

2.4.1. Schematically defined vessel operational pattern

In Figure 2.10, the operational pattern of a deep-sea vessel is schematically drawn. This pattern follows a cycle, from leaving the previous port until departing the current port. As stated in subsection 2.2.1, deep-sea vessels often operate on a shipping line, following multiple ports in succession. The schematically drawn figure is only one leg of the shipping line. When travelling to the next port, the same vessel will follow the same cycle of operations, only with different sailing and/or waiting times. In the figure, all blocks represent a length of time. For each block in the figure, a description is provided. The sections displayed in Figure 2.10 will be described in more detail in the following sections.

Travel time

When the vessel departs from the previous port, as indicated by the arrow on the left side of Figure 2.10, it transmits an estimated time of arrival (ETA) to the next terminal. After the vessel departs, it starts traversing the sailing leg to the next terminal. During the sailing leg, the vessel sails at its own optimal velocity. The optimal velocity of the vessel is defined as the velocity that minimises its total operating costs over the sailing leg while still arriving at the terminal on time. The time this

section takes is defined as the time it takes to travel from the anchorage point of one terminal to another. The length of the sailing leg is dependent on the port, as sailing from Singapore to Rotterdam takes longer than sailing from Antwerp to Rotterdam. However, the length of the sailing leg between the two ports is not always the same, as for instance in 2024, vessels travelling from Singapore to Rotterdam due to the Red Sea Crisis had to travel around South Africa, instead of travelling through the Suez Canal, adding an additional 3500 nautical miles to their voyage [26].

Waiting at anchorage point
As soon as the vessel arrives at the anchorage point, it transmits its actual time of arrival(ATA). It is possible that no berth is available due to delays from other vessels, longer-than-expected service times, or a lack of pilots needed to guide the vessel to the berth. During the waiting time at the anchorage point, the vessel is not productive.

Journey time from anchorage point to quay
When a spot on the quay is available, the vessel can travel from the anchorage point to the quay. For this part of the journey, a pilot is on board to traverse through the relatively small entry points of the container terminal. The time it takes a vessel to travel from the anchorage point to the quay can be different per port, depending on the access route to the terminal.

Berthing time
As soon as a vessel arrives at the container terminal, it will start to berth. Using tugboats, the vessel is precisely manoeuvred to the designated spot on the quay. For this time to start, the designated berth of the vessel has to be empty.

Waiting time on quay for service
Once the vessel arrives at the terminal, loading and unloading the containers can begin. However, if no cranes are available and other vessels have higher service priorities, the vessel must wait until the cranes finish servicing the other vessels.

Service time
Service time refers to the time it takes to load and unload a vessel's containers. Service time depends on the number of containers being loaded and unloaded, as well as the number of cranes working on the vessel. Vessel operators can contractually obligate a terminal to use a minimum number of cranes on a vessel. The actual number of cranes working on a vessel is limited by the size of the vessel, as cranes have to be at least a minimum distance away from each other. The productivity of quay cranes can vary depending on the type of vessel. For instance, for smaller vessels with a low number of containers per bay, the cranes have to move more frequently. For larger vessels, the travel time of the trolley from the end of the beam of the crane to the quay is significantly longer.

Waiting time on quay for pilot
When a vessel's service time is complete, it must wait until a pilot is available to guide it back into open waters from the quay.

Unberthing time
When the pilot is available, it is manoeuvred from the quay back to the canal, facing the correct direction. During this time, the vessel's berth location is not available.

Journey time from quay to anchorage point
The vessel returns to the anchorage point and departs. The moment the vessel arrives at the anchorage point can be transmitted as the due date. Now the vessel can travel to another port, and the cycle repeats, with different information for the sailing leg.

2.4.2. Analysis of vessel operations
Vessel operations can be described as a system, with inputs, outputs, requirements and performance indicators. The operations of the system are shown in a blackbox in Figure 2.11. The inputs, outputs,

requirements and performance indicators shown in the blackbox are specified in the following paragraph.

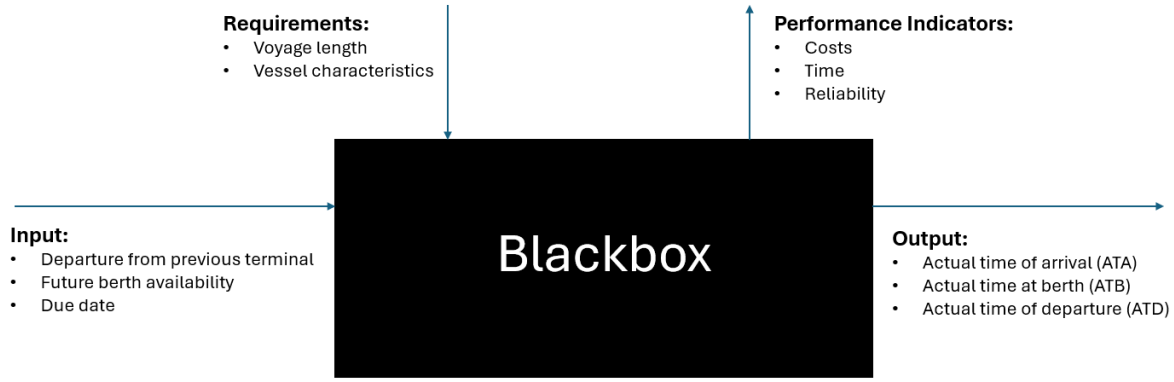


Figure 2.11: Blackbox of vessel operations

Input

Departure from previous terminal

The first input for a vessel's operations is the time it departs from a previous port. Starting at this moment, the vessel can start operations.

Future berth availability

The second input for a vessel's operations is the future berth availability. Based on the future berth availability, the vessel can determine the velocity over the incoming sailing leg.

Due date

The final input for a vessel's operation is its due date. The due date constrains the vessel's velocity, as its service must be completed before this moment. Providing a due date ensures the vessel does not sail at its lowest possible velocity.

Output

Actual time of arrival (ATA)

The first output of the blackbox is the actual time of arrival. The actual time of arrival is the moment at which the vessel arrives at the anchorage point.

Actual time of berth (ATB)

The second output of the blackbox is the actual time of berth. The actual time of berth is the moment at which the vessel starts the berthing process.

Actual time of departure (ATD)

The final output is the actual time of departure. The actual time of departure is the moment after the vessel has travelled from the quay back to the anchorage point and can begin to travel to its next destination.

Requirements

Voyage length

The first requirement of the blackbox is the voyage length. The voyage length is dependent on the port of departure and the port of arrival. Based on the voyage length, the vessel can determine the optimal velocity.

Vessel characteristics

The second requirement of the blackbox is the vessel characteristics. Such as its fuel consumption and average velocity. Based on the vessel's characteristics, the minimum and maximum velocities, and the vessel's operating costs can be determined.

Performance Indicators

Costs

The first performance indicator is the vessel's overall costs. This consists of variable costs, such as fuel consumption, and time-dependent costs, such as crew costs, vessel depreciation/chartering costs, and idle operating costs. The lower the overall cost, the higher the vessel operator's profit margin.

Time

The second performance indicator is the overall time it takes a vessel from the departure of one terminal to the departure of another. This time consists of all the different subsections of the schematically defined vessel operational pattern. If the total pattern takes less time, the vessel can berth at more terminals, increasing its potential revenue.

Reliability

The final performance indicator is the vessel's reliability. This is defined as the deviation of the originally transmitted ETA.

2.4.3. Conclusion

The analysis of deep-sea vessel operations reveals three points at which the vessel is not productive: waiting time at the anchor point, waiting time at the quay for service, and waiting time at the quay for the pilot. Of these three, only the first can be directly influenced by the vessel itself. This can be achieved through speed optimisation, in which the vessel adjusts its sailing speed to arrive at a berth when one becomes available.

However, the effectiveness of speed optimisation depends entirely on the accuracy of the berth availability forecasts. When the provided information is inaccurate or too uncertain, or no information is provided at all, the vessel cannot optimise its speed and therefore might have to wait at the anchorage point.

This demonstrates that deep-sea vessel operations influence waiting times at the anchorage point due to the lack of or inaccuracy of future berth availability.

2.5. Combined system analysis

The combination of vessel operations and seaside operations of a container terminal would allow for speed optimisation of arriving vessels. The advantage of combining both systems is that waiting times and operating costs can both be decreased, increasing performance indicators of both subsystems.

The blackbox for the combined system is shown in Figure 2.12. The system is similar to the blackbox shown in Figure 2.6. The main inputs are vessel departure from the previous terminal, voyage characteristics, and vessel characteristics. The previous terminal is the port of departure of the vessel, including the distance over seas in nautical miles. The requirements of the blackbox are dependent on the quay layout and the number and type of cranes. The main performance indicators are defined as vessel sailing time, vessel waiting time, and overall vessel costs. The output is the vessel departing the system.

2.6. Conclusion

This chapter answers the sub-question, "How do vessel types, vessel operations, and seaside operations influence waiting times at the terminal?" In the first part of the chapter, it is found that the terminal operations and vessel operations interact at the seaside of the container terminal. In the second part of the chapter, the three main vessel types are discussed. Deep-sea vessels are large vessels that handle the intercontinental transport of containers. Short-sea vessels are smaller

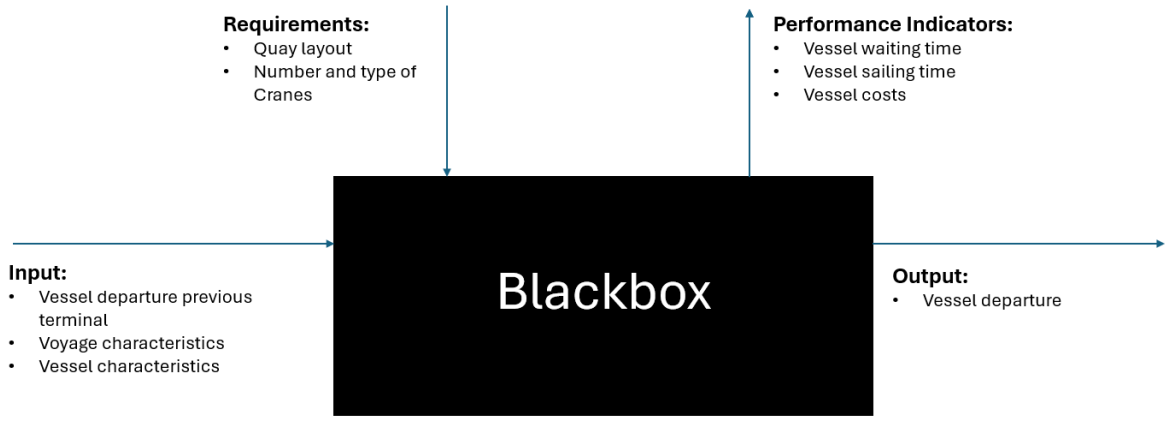


Figure 2.12: Combined blackbox of vessel operations and seaside operations

sea-going vessels that transport containers from deep-sea ports to smaller terminals in Europe. Barges can transport containers inland. Deep-sea vessels have the highest operating costs, and can thus gain the most by optimising their velocity. In the third part of the chapter, the system of vessel operations is analysed. The main cause of waiting times at the anchorage point is found to be a lack of knowledge about future berth planning, which prevents vessels from optimising their velocity. In the fourth part of this chapter, the system of seaside operations at a container terminal is analysed. Here, it is found that uncertainty in vessel arrival and service time contributes to the vessel's waiting time at the anchorage. In the final part, a combination of vessel operations and seaside operations is illustrated. The main causes of waiting times are uncertainties in the vessel's arrival time at the terminal and in future berth availability for the vessel. By combining the two models into a single model and having the vessels communicate with the terminal, waiting times can be decreased.

3

Literature Review

In this chapter, the sub-question ”How can vessel speed optimisation be incorporated into the berth allocation problem?” will be answered. To answer this sub-question, a literature review on the topics will be conducted. In section 3.1, the berth allocation problem will be introduced. In section 3.2, different methods and approaches to uncertainty are shown. In section 3.3, different applications of speed optimisation are shown. In section 3.4, the stakeholders and their roles in the berth allocation problem and vessel speed optimisation are identified. In section 3.5 the research question is answered.

3.1. Berth allocation problem

The berth allocation (and scheduling) problem is an operational problem that aims to plan vessel berths at container terminal quays. Different goals can be pursued in the berth allocation problem, for instance, minimising a vessel’s waiting time before reaching a quay or maximising resource utilisation. Over the past 25 years, various aspects of the berth allocation problem have been researched. These studies can be classified according to the scheme proposed by Bierwirth and Meisel [4]. The classification scheme is shown in Table 3.1.

Spatial	**Temporal**	**Handling Time**	**Performance Measures**
Discrete	Static	Fixed	Waiting time
Continuous	Dynamic	Positional	Handling time
Hybrid	Due Dates	Quay crane assignment	Completion time
Draft	Cyclic	Quay crane scheduling	Tardiness
	Stochastic	Stochastic	Speed up
			Rejection
			Resource utilisation
			Position
			Miscellaneous

Table 3.1: BAP classification according to Bierwirth and Meisel [4]

Spatial attribute

The spatial attribute defines how the terminal’s quay is divided into different berths. The berth is the location on the quay where the vessel is serviced. The following spatial attributes were identified by Bierwirth and Meisel [4].

(a) Discrete layout: When a discrete berth layout is chosen, the quay is divided into a number of sections. At each section, only one vessel can be served at a time. In reality, this is not an

accurate estimation of the quay, as ships have different lengths, so not all vessels can berth at the same location. This, however, is computationally the easiest, as there is an integer number of berths.

(b) Continuous layout: When a continuous layout is chosen, vessels can be served at all locations at a quay, without being hindered by predetermined berths. The only constraint is that the vessel has to be served within the physical boundaries of the quay. A continuous layout is more complicated than a discrete layout, but serves a more realistic representation of the quay.

(c) Hybrid layout: Similar to a discrete layout, the quay is divided into a number of berths. However, large vessels can be served at multiple berths, or multiple smaller vessels can be served at a single berth.

(d) Draft: Vessels can also be constrained to parts of the quay by their draft. This constraint ensures that vessels are berthed at positions with sufficient water depth.

Temporal attribute

The temporal attribute defines how vessel arrival and departure times are modelled. In some models, all the vessels are assumed to already be waiting for anchorage, while in other, more complex models, the arrival time of the vessel is uncertain, to account for delays that may occur at sea. The following temporal attributes were identified by Bierwirth and Meisel [4].

(a) Static arrival time: There are two types of static arrival, either there is no arrival time for incoming vessels, or it is assumed that they already wait at port, and can berth immediately. Alternatively, there is only a soft constraint on berthing times, and it is assumed that a vessel can speed up to arrive at a moment earlier than its expected arrival time.

(b) Dynamic arrival time: A fixed arrival time is given, and vessels cannot berth before the expected arrival time. Slow steaming could be integrated into the problem, such that vessels arrive later than their ETA, and do not have to wait at the terminal.

(c) Due dates: Besides the earliest time of arrival, the latest time of departure can be an additional constraint. This ensures that vessels are not delayed unreasonably long and can arrive at the next terminal at the expected time.

(d) Cyclic arrival time: Vessels repeatedly call at terminals in fixed intervals.

(e) Stochastic arrival time: The arrival time of vessels is a stochastic parameter defined by continuous random distributions or scenarios with discrete probability of occurrence. This adds uncertainty to vessel arrival times, making the simulation more realistic but increasing the problem's complexity.

Handling time attribute

The handling time attribute defines how the handling time is modelled in the berth allocation problem. The handling time can be fixed for a vessel based on its berthing position and crane assignment and scheduling, or defined as a stochastic function. The handling time attribute determines how the parameter handling time is described. The handling time is an interesting attribute to vary when combining the BAP with the Quay Crane Assignment and Scheduling Problem; in this case, it is also likely to serve as a performance measure. The following handling time attributes were identified by Bierwirth and Meisel [4].

(a) Fixed times: The handling time of vessels is known in advance and unchangeable. This simplifies the problem; however, a vessel does not always load and unload the same number of containers, so it is not an entirely accurate representation.

(b) Positional dependency: The vessels' handling time depends on the berth location of the vessel.

(c) Quay Crane assignment: The vessels' handling time depends on the number of quay cranes serving the vessel. This is specifically interesting when integrating the BAP with the Quay Crane Assignment Problem (QCAP). Integrations of the BAP with the Quay Crane Assignment Problem are also referred to as BACAP.

(d) Quay Crane scheduling: The handling time of vessels is dependent on the work schedule of the assigned quay cranes. This is specifically interesting when integrating the BAP with the Quay

Crane Scheduling Problem (QCSP). Integrations of the BAP with the Quay Crane Scheduling Problem are also referred to as BACASP.

(e) Stochastic: The handling time of the vessels is subject to stochastic parameters, either defined by continuous random distributions or by scenarios with discrete probability. This adds uncertainty to the number of containers loaded and unloaded from the ship, increasing the problem's complexity because the exact handling time of the vessel is unknown.

Performance measure

The performance measure describes what the model tries to minimise (or maximise in the case of resource utilisation). Multiple performance measures can be used simultaneously by assigning different cost weights to each. The following performance measures were identified by Bierwirth and Meisel [4].

(a) Waiting time: The waiting time represents the time a vessel has to wait, from arrival until serviced.

(b) Handling time: The time required to unload the vessel. This is a common performance measure in a QCAP or a QCSP integrated with the BAP. When handling time is a performance measure, it is not an input, but an output of the simulation.

(c) Completion time: When completion time is minimised, the model is optimised so vessels can depart at the earliest possible time.

(d) Tardiness: When tardiness is minimised, the number of vessels that leave the terminal later than their due date is minimised.

(e) Speed up: Vessels can speed up to arrive earlier at the quay; the amount by which a vessel has to speed up can be minimised to minimise extra fuel costs.

(f) Rejection: When a vessel can not be served within its desired time window, it can be rejected. The rejection rate can be minimised to avoid bringing opportunity to other terminals.

(g) Resource utilisation: Resource allocation can be optimised at the terminal side of the problem. This is common when resources are scarce.

(h) Position: The position on the quay where a vessel berths can be optimised to avoid large travel distances for horizontal transport.

(i) Miscellaneous: The term miscellaneous is added to the classification scheme if the performance measure does not fit in one of the listed performance measures above.

3.2. Berth allocation under uncertainty

The stochastic temporal and handling time attributes in Table 3.1 are one way to model uncertainties in the berth allocation problem. In the review paper of Rodrigues and Agra[23], multiple other approaches to handling the uncertainty and modelling methods are described.

3.2.1. Sources of uncertainty

In the research by Rodrigues and Agra [23], five sources of uncertainty are identified: arrival times, handling times, quay crane breakdowns, unscheduled vessels, and general disruptions. Most of the papers reviewed in Rodrigues and Agra discuss uncertainties in arrival time, handling time or both. In the following section, the causes for these two main uncertainties are discussed.

Arrival time

Vessel arrival time can be an uncertain parameter in the BAP. The actual time of arrival (ATA) of a vessel is not known in advance; only an estimated time of arrival (ETA) is known. Deviations from the ETA can occur due to weather conditions, delays in previous harbours or the vessel sailing at other speeds than expected. Deviations in the ATA from the ETA can be defined by historical port call data. Big deviations in the ATA from the ETA can lead to long waiting times at the quay and congestion at the harbour. Modelling the uncertainty in the ETA can create a more realistic BAP. If uncertainties in arrival time are modelled in the BAP, this leads to a stochastic classification in the temporal attribute of table 3.1.

Handling Time

Vessel handling time can also be modelled as an uncertain parameter in the BAP. Each vessel type has a maximum capacity of containers. However, the number of containers a vessel loads and unloads at a container terminal is not constant. Besides the number of containers, the time it takes for a vessel to unload also depends on the number and the efficiency of the quay cranes assigned to the vessel. From AIS data, the average time spent in port can be derived per vessel type. Uncertainty in the handling times can lead to uncertain turnaround times. Which, in turn, makes it harder to plan vessels back-to-back. An uncertain number of containers being loaded and unloaded also leads to a more difficult integration with the QCAP or QCSP. If uncertainties in handling time are modelled in the BAP, this leads to a stochastic classification in the handling time attribute of table 3.1.

3.2.2. Uncertainty optimisation approaches

The review paper by Rodrigues and Agra discusses three optimisation approaches to handle uncertainty in operations [23]. These approaches define how to use the BAP in practical situations.

Proactive approach

When the proactive approach is used to handle uncertainty, all decisions are made before the start of the time horizon. The solution with a proactive approach has to remain feasible for all possible realisations of the uncertain parameters. When the proactive approach is taken in the BAP, the original plan has to remain feasible when a vessel has been delayed. Therefore, buffers must be implemented to allow vessels to arrive at different times. Only one BAP has to be calculated, so computation time is not an issue.

Reactive approach

When the reactive approach is taken to handle uncertainty, actions are undertaken when the value of an uncertain parameter becomes known. The reactive approach starts with a baseline plan, which is determined using deterministic parameters, in which the cost for rescheduling operations is minimised. When the reactive approach is taken in the BAP, a baseline berth planning without buffers is made, which works for the expected time of arrival of all vessels. When the ETA of a vessel is changed, the berth planning is updated. When the berth planning has to be updated, a new solution to the BAP has to be calculated. This ensures that computation time is limited, as the new berth plan has to be ready as soon as possible. To allow for faster computation, the BAP has to be simplified.

Proactive/reactive approach

The proactive/reactive approach is a combination of the proactive and reactive approaches. Similar to the reactive approach, a baseline plan is made, and the cost of rescheduling operations is minimised when a disruption occurs. In contrast to the reactive approach, a form of uncertainty is taken into account when formulating the baseline plan. When the proactive/reactive approach is taken in the BAP, the first step is to establish the baseline berth planning. In baseline berth planning, the ETA uncertainty is already accounted for through buffers. When an event disrupts the baseline plan, the plan is updated. When this event occurs, a new solution to the BAP has to be calculated. The time available for the calculation is also limited.
In the papers reviewed by Rodrigues and Agra, a trend is visible in which the more recent papers tend to have a proactive/reactive approach to handle uncertainty.

3.2.3. Uncertainty modelling

The review paper by Rodrigues and Agra discusses four different classes of models dealing with uncertainty [23]. These four classes are stochastic programming, robust optimisation, fuzzy models and deterministic models.

Stochastic programming

In stochastic programming (SP), it is assumed that the uncertain parameters are defined by known probability distributions. Several different subsets of SP models are described in the paper: Risk measure-based models, chance-constrained models, queueing models, static classic stochastic models and adjustable classic stochastic models. Risk measure-based models are used to control

the deviations from expected values and control extreme deviations. Chance-constrained models ensure that the probability of large waiting times is less than a certain percentage. In queuing models, the BAP is treated as a queuing problem, from which policies can be derived for optimally allocating vessels. The classic stochastic models are the most common models in BAP. In the classic stochastic models, the objective is to optimise a function with uncertain parameters. In static classic stochastic models, the solution is found during the planning phase, whereas in adjustable classic stochastic models, the solution can change based on adjusted variables, such as berthing times and departure times.

For instance, Han Lu and Xi used a proactive approach to solve the berth allocation problem with uncertain arrival times [11]. In their research, the realisation of these stochastic parameters occurs during the model's planning phase, after which the parameters are no longer allowed to change.

Robust optimisation

In robust optimisation (RO), it is assumed that the probability distributions of uncertain parameters are not exactly known or do not exist. Because the uncertainty is not known, robust optimisation often depends on the worst-case scenario. The different solution procedures are divided into two groups: static models and adjustable models. As RO does not rely on probability distributions, it is suitable when no statistical data is present. A way to implement RO in the BAP is to use buffers in vessel arrivals or handling times. Using buffers decreases the possibility of vessels waiting before being able to berth; however, it also decreases the number of vessels that can be planned, as the quay is also claimed during a buffer period for a vessel. In the case of vessel arrivals, it is also possible to use time-dependent buffers, as the uncertainty in arrival decreases as the vessel gets closer to the port.

For instance, Rodriguez-Molins et al. use buffers to robustly plan the BACAP [24]. Their robust model is used in combination with a proactive approach to overcome uncertainties caused by incidents without previous knowledge. The robustness is introduced by using buffers on operational times.

Using buffers is generally useful for proactive planning to account for the uncertainty of all incoming vessels.

Fuzzy models

When fuzzy models are used, the unknown parameters are assumed to belong to a fuzzy set. The most common fuzzy set is a triangular set, defined by three numbers: the minimum, the maximum, and the expected value. Fuzzy models are useful when the value of a parameter is vague. Fuzzy models create a flexible plan that is adaptable within the range of the assumed parameters. Two different methods of fuzzy modelling are described in the paper: fully fuzzy programming and fuzzy programming. In fully fuzzy programming, the uncertain parameters can be represented by fuzzy numbers, and a fuzzy mixed-integer linear program is used. In fuzzy programming, only a fuzzy optimisation model is used for solving the problem.

Deterministic models

In a deterministic model, the focus is on replanning operations after a disruption. The deterministic models work with a baseline berth plan. When a disruption occurs, two approaches can be taken: either re-optimise the operation after the disruption or minimise the deviation of the updated schedule from the baseline plan.

3.2.4. Rolling horizon

A reactive approach to dealing with uncertainty is via a rolling-horizon heuristic. Andersson et al. propose a rolling-horizon heuristic for solving a BAP with vessel-speed optimisation for Roll-on/Roll-off (RoRo)shipping [2]. In a rolling horizon heuristic, the model is solved for the planning horizon, of which a small part, the sub-horizon, is executed. After the sub-horizon has passed, the model is solved for a new planning horizon. The planning is only made for vessels that arrive within the planning horizon. Ji et al. propose a hybrid rolling horizon heuristic [16]. In the hybrid rolling horizon heuristic, the problem is not only solved when the sub-horizon has passed, but also when an event is detected. Examples of events are unscheduled vessels.

3.3. Vessel speed optimisation

Vessel speed optimisation (VSO) can be used to achieve just-in-time sailing. For a vessel to optimise its speed, it needs to know the terminal's future berth availability. In the BAP, the future berths are planned. By combining these in the berth plan, the moment the berth becomes free can be communicated to the vessel, so it can arrive at the terminal just in time without waiting. Section 3.3.1 discusses the speed optimisation, while Section 3.3.2 focuses on the integration with terminal operations.

3.3.1. Vessel speed optimisation

Psarafitis et al. investigated ship speed optimisation and developed models for different routing scenarios [21]. They focus on optimising the velocity of a single vessel. The goal is to minimise the total cost of the trip, which is divided into fuel costs, time charter cost and cargo inventory cost. The paper concludes that the optimal environmental solution of slow steaming is not always the same as the optimal solution with minimal cost. However, waiting time is not considered, limiting applicability to congested terminals.

Andersson and Ivehammar examine the benefits of adjusting speed rather than anchoring [3]. They do not integrate it into a BAP, but they study the number of vessels at the anchorage points of all terminals in the Baltic Sea. They state the savings in fuel, emissions, and euros that could be achieved based on the moment vessels receive information about their berth and the reduction in speed. If the information could be provided 12 hours in advance, and vessels would reduce their velocity by 25 %, 27 million euros could be saved annually in the Baltic Sea. However, the assumption is that the information can be provided 12 hours in advance, subject to berth availability at the terminal and the vessel's previous port. While this is a simplification of the speed optimisation, the actual reduction in speed and the information available on the vessel's past destination, terminal location, and terminal quay utilisation depend on these factors. However, it indicates the potential to integrate VSO into the BAP.

Wang and Wang investigate the influence of heterogeneous vessels on the deployment, scheduling, and sequencing of vessels along a liner container shipping route [32]. In shipping, container vessels differ in capacity and in all costs associated with them.

3.3.2. Optimisation of vessel arrival times

Van Leeuwen and Dekker compare the vessel-speed optimisation of a stream of vessels arriving at a single-berth terminal with a *FCFS* strategy [31]. The velocity of the vessel is optimised at an adjustment point of 320 nautical miles out of the terminal. In the research, fuel and operating costs are defined so that if the vessel does not experience interference from other vessels, it will sail at v^{eco}. If it sails at a lower velocity, the fuel costs will be lower, but the operating costs will be higher; if it sails at a higher velocity, the operating costs will be lower, and the fuel costs will be higher. Another study by van Leeuwen and Dekker discusses the value of the vessel's information [30]. The better the information a vessel has, the more accurately it can predict when the vessel ahead of it will depart, and the more accurately it can determine a velocity at which no waiting time is minimised. These situations differ from 1, where the vessel has no information on vessels ahead or behind and sails at v^{eco}, to 7, where every vessel knows the arrival time and service time of all the other vessels.

Lang and Veenstra compare central and decentralised arrival planning strategies [17]. In the central strategy, the terminal has full visibility over the vessel attributes such as speed, call size, and remaining distance of incoming vessels. Vessels will always obey speed directions from the terminal. While in the decentralised strategy, vessels determine their velocity without knowledge of the actual terminal conditions. In their research, all vessels are assumed to depart from 5400 nautical miles. The research shows the potential savings in fuel costs from centralised planning.

Alvarez et al. compare three different situations in a discrete event simulation of a berth allocation problem [1]. In the discrete-event simulation, vessels arrive stochastically. The three situations are different berthing strategies: First Come First Serve (FCFS), Standardised Estimated Arrival Time (SETA), and Global Optimisation of Speed, Berth, and Equipment Allocation (GOSBEA). In policy FCFS, the vessel that arrives first is served first; in SETA, vessels are admitted based on their ETA.

In GOSBEA, vessels sail at a recommended speed when they are 500 nautical miles from the port, or at the moment of departure, when the vessel arrives from within 500 nautical miles. Every time a new vessel enters, it solves a cost function and chooses a velocity with minimal costs. Berthing policy GOSBEA scores better than the other policies.

Ji et al. study a multi-port berth allocation problem (MBAP) [15]. In an MBAP, vessels visit ports sequentially, and the problem is to minimise the total cost of sailing, waiting, delays and handling. In their research, they propose a modified adaptive large neighbourhood search (MALNS) algorithm to solve large-scale MBAPs. Their heuristic is combined with three different discrete speed selection strategies. It selects a discrete speed, calculates the arrival time, the earliest start of operation and determines its departure time. Speed is discretised into 11 levels, and the fuel consumption at each speed is calculated.

Yu et al. divide the berth allocation into a lower level (vessel cost minimisation) and upper level departure delay minimisation[33]. The lower level tries to minimise the cost for all the different shipping companies. These costs are minimised through minimising fuel cost, electricity cost and port fees. As the problem contains the BACAP, which is an NP-hard problem, it is challenging to solve the proposed mixed-integer programming model. To solve it, both levels are decomposed and solved with a Nested Non-dominated Sorting Genetic Algorithm II.

3.4. Stakeholders

The berth allocation problem and vessel speed optimisation have a number of key stakeholders. The main stakeholders are the vessel operators and the container terminal operators. But others can be identified as well, such as the local government and the port authority. Each stakeholder values different performance measures.

3.4.1. Vessel operators

Vessel operators would like to maximise the profit from their vessels. The profit a vessel generates depends on the operating costs and the number of ports it can visit. Operating costs can be minimised by reducing the fuel consumption of a vessel. For this, the velocity of the vessel has to be optimised such that it can arrive just in time. In theory, this should result in vessels that would otherwise have waited at the anchorage point now travelling at a lower velocity. The other way to maximise profits is by reducing the time a vessel spends in port. This can be achieved by minimising completion time or tardiness performance measures of the berth allocation problem.

3.4.2. Container terminal operators

The container terminal operator seeks high resource utilisation, minimal horizontal transport distance, and a robust schedule. The high resource utilisation can be achieved by maximising the resource utilisation performance measure of the berth allocation problem. The minimum distance can be achieved by minimising the position performance measure of the berth allocation problem. A robust schedule, which makes it easier for the terminal operator to plan shifts in advance, can be achieved by minimising vessel waiting time. Minimising vessel waiting time can be achieved using vessel speed optimisation, to allow vessels to arrive just in time.

3.4.3. Local government

The local government would like to minimise pollution in the coastal waters while still maintaining the country's maritime import and export. Lower greenhouse gas emissions can be achieved by arriving just in time [14]. When arriving just in time, vessels slow down to avoid waiting for anchorage. Vessel speed optimisation is key to achieving this minimisation. This reduces pollution in nearby waters in two ways: fewer vessels have to run their engines while waiting for a berth, and vessels can travel at lower speeds, using less fuel per unit of distance travelled.

3.4.4. Port authority

The port authority aims to prevent harbour congestion. By minimising the vessels' waiting time, the model reduces the time vessels spend, on average, at anchorage, thereby reducing congestion in the port.

3.5. Conclusion

The goal of this chapter is to answer the sub-question: ”How can vessel speed optimisation be incorporated into the berth allocation problem?” The sub-question is answered by conducting separate literature reviews of the berth allocation problem under uncertain arrivals and of vessel speed optimisation. The hybrid rolling-horizon strategy is shown to be a good solution for handling uncertainty about incoming vessels, and the uncertain parameters, such as arrival times, can be modelled using stochastic variables, since there is data on vessel arrivals. Different approaches to integrating berth allocation and vessel speed optimisation are presented in section 3.3. These integrations aim to achieve the lowest costs, either at the vessel/operator level or globally. The costs then consist of vessel speed-optimisation objectives, such as minimising fuel consumption, and terminal objectives, such as minimising waiting time. The reviewed literature shows that while vessel speed optimisation can reduce costs and emissions, its effectiveness depends on the level of information sharing and its integration with berth allocation.

4

Speed Optimisation Strategies

In this chapter, the following sub-question is answered: ”Which vessel speed optimisation strategies can be identified and structured for use in simulation?” First, based on the literature study, three different speed optimisation strategies are proposed in section 4.1. Then the different models are described in section 4.2. In section 4.3, an overview of all four models is provided based on the optimisation objective, the moment of decision-making and the coordination level. Finally, in section 4.4 an answer is given to the sub-question.

4.1. Speed optimisation strategies in literature

Vessel speed optimisation can be used in different ways. One way to implement it could be similar to the models by Alvarez et al. [1] or van Leeuwen and Dekker [31]. In their models, vessels are assigned a velocity once, and that velocity is assumed to be fixed from that point onward. In the model of Alvarez et al., the velocity is recommended when the vessel is 500 nautical miles away from the port, or at the moment of departure from the origin port if this is less than 500 nautical miles. The distance from the focal port is uniformly distributed between 400 and 2400 nautical miles. In the van Leeuwen and Dekker model, the velocity is determined as soon as a vessel reaches the adjustment point at 320 nautical miles [31]. Both are forms of slot booking, in which the terminal books the vessel a slot that is best according to the cost function. To simplify solving the problem, it can be decomposed into a lower level (vessel speed optimisation) and upper level (berth planning) as proposed by Yu et al [33], at which the upper level produces an optimal berth planning based on the knowledge of all current vessels, and the lower level determines the best moment of arrival based on the produced berth planning of the upper level. This strategy is called single-speed optimisation.

Vessels arriving at the Port of Rotterdam often have different origins, with varying distances. For instance, a vessel can depart from Antwerp, at around 140 nautical miles, or from Singapore. When a slot has already been booked by a vessel further away, there may not be any possible slots for vessels departing from Antwerp. To combat this, the speed of all incoming vessels has to be re-optimised. This is done by first removing all booked slots from every sailing vessel. The vessels are then sorted based on remaining distance. The terminal books a new slot for all incoming vessels. This strategy is called re-optimise at every new departure.

In addition to a slot-booking approach, vessels can also be assigned using a more complex berth-allocation algorithm. Similar to the problem discussed by Ji et al. [15], the total cost of sailing, waiting, and handling can be minimised for all vessels. In this case, either when a vessel departs its previous port or after a fixed time has passed, velocities are selected such that the lowest possible costs of all sailing vessels combined are reached. The problem can be solved using a hybrid rolling horizon strategy [16], in which the vessel speeds are optimised for all sailing vessels at every planning interval, and when a new vessel departs the previous terminal. The final strategy is called global speed optimisation.

These three strategies differ in what they aim to impact. The first one tries to find the optimal solution

for each vessel at the moment of departure and favours vessels from further away. The second one tries to find the optimal solution for every vessel based on the remaining distance. In the first two strategies, the terminal seeks an optimal solution for each vessel based on current information; however, as it seeks an optimal solution for just the single vessel, the vessel will not speed up. The final strategy tries to minimise the overall cost, sacrificing the performance of some vessels for the benefit of others.

These three strategies also differ in their planning frequencies. In the first strategy, the velocity of each vessel is optimised only once. In the second strategy, the vessel's velocity is optimised at each new vessel arrival. In the third strategy, the velocity of the vessels is optimised at a fixed time interval, and at every vessel arrival. The advantage of optimising velocity more often is that the vessel has more accurate information about the moment of departure of vessels currently at the quay and about the arrival of other vessels. The first strategy assumes fixed berth slots once a vessel has departed the previous terminal, while the second and third strategies allow for re-optimising of already scheduled vessels.

These three strategies are selected as they represent increasing levels of coordination between vessels and the terminal.

4.2. Description of strategies

In the following section, the models for the baseline strategy and the three speed optimisation strategies are described. In all speed-optimisation strategies, the terminal acts as the decision-maker, determining vessel speeds based on available information. In the baseline strategy, vessel speeds are determined by the vessel operator. Each strategy follows the same process: vessel departure, sailing, and arrival at the terminal. The strategies differ in how and when vessel speeds are determined, the information available at the time of velocity determination, and the optimisation objective.

The different speed optimisation strategies have three main actors: the vessel operator, who is responsible for executing the sailing leg and adhering to speed changes. The terminal speed optimiser, which determines vessel speeds when speed optimisation is applied. The terminal berth planner is responsible for maintaining the berth schedule and estimating current operation times. Communication between these actors occurs at vessel departure and, depending on the strategy, at re-optimisation points. The berth planner provides the berth planning information to the speed optimiser, which communicates recommended vessel speeds to the vessel operator. The vessel operator subsequently executes the sailing leg using the assigned speed.

4.2.1. Baseline

In the baseline strategy, vessels depart from the previous terminal at a fixed velocity. This velocity is chosen to minimise operating and fuel costs over the sailing leg, without accounting for waiting time at the terminal. The baseline strategy represents the current "Sail fast, Then Wait" practice. The process of the baseline strategy consists of three main phases: departure, sailing, and arrival.

Departure

Each vessel departs from the previous terminal at a given departure time. The vessel operator saves its departure time. At the same time, it calculates its time of arrival. The time of arrival is the distance to the port (D) divided by the optimal velocity ($\tilde{v}$). $\tilde{v}$ is defined as the operating velocity of the vessel at which the lowest combined operating and fuel costs over the sailing leg are incurred. The expected time of arrival is then communicated to the terminal berth planner.

The terminal berth planner registers the vessel as an incoming vessel. The berth planner calculates the optimal berth allocation for all sailing vessels, waiting vessels and vessels at the quay, without influencing their velocity.

Sailing

After determining the velocity, the vessel sets sail to the anchorage point of the terminal. The time the method takes is the distance to the port (D) divided by the optimal velocity ($\tilde{v}$). The cost of sailing is

defined using the cost function (Equation 4.1) from the research of van Leeuwen and Dekker [31].

$$C(v) = a \cdot D \cdot v^{b-1} \cdot c_f + t_w \cdot c_w + \frac{D}{v} \cdot c_o \tag{4.1}$$

Where $C(v)$ is the cost of sailing at velocity v, a and b are vessel parameters, based on the size of the vessel, D is the distance from the previous terminal to the new terminal in nautical miles, v is the velocity in nautical miles per hour, t_w is the expected waiting time, and c_f, c_w and c_o are the cost coefficients of fuel, waiting and operating. Since the vessel has no knowledge of the expected waiting time, t_w is set to 0. The cost coefficients are linked to each other by the following equations:

$$c_o = a(b-1) \cdot c_f \cdot (\tilde{v})^b \tag{4.2}$$

$$c_w = c_o \cdot 1.1 \tag{4.3}$$

The cost of waiting exceeds the cost of operating, thereby penalising vessels for waiting. The vessel parameters are linked to each other by the following function:

$$a = \frac{\tilde{f}}{\tilde{v}^b} \tag{4.4}$$

Where $\tilde{f}$ is the fuel consumption in metric tons per day.

In the baseline scenario, the expected waiting times for all vessels (t_w) are not communicated between the terminal and the vessels, and thus are set to zero. So the vessel minimises fuel and operating costs, setting sail at $\tilde{v}$.

Arrival

After the travel time, the vessel arrives at the anchorage point. At the anchorage point, the vessel joins the waiting vessels and waits until the terminal berth planner assigns a quay. As soon as the vessel arrives, the sailing leg is over. The sailing time is the total time between the moment of departure and the arrival at the anchorage point.

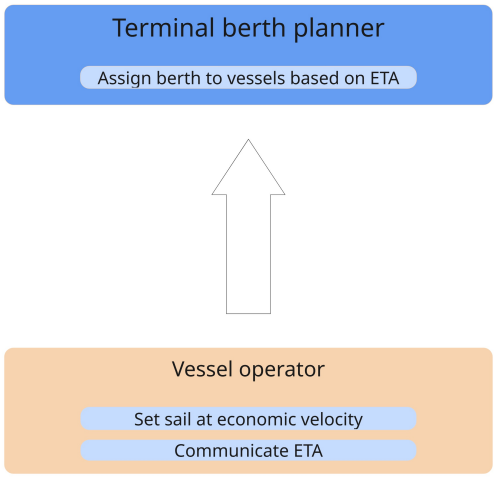


Figure 4.1: A schematic drawing of communications between stakeholders in the baseline strategy

In Figure 4.1, the communications between the vessel operator and the terminal berth planner are shown. The vessel only communicates its expected time of arrival at its departure from the previous terminal to the terminal planner and sets sail at its economic velocity.

4.2.2. Single-speed optimisation

The single-speed optimisation strategy extends the baseline by allowing the terminal speed optimiser to optimise the vessel's speed at the moment of departure. Compared to the baseline, the main difference is who is responsible for the decision-making. Now, the terminal can decide the speed of the incoming vessel.

Instead of the vessel determining its own velocity, the terminal speed optimiser selects the velocity for the vessel. The speed optimiser selects the velocity based on the current knowledge of all sailing vessels, waiting vessels and vessels at the quay. The terminal calculates the cost of arrival using the function displayed in Equation 4.1, for all the velocities of the vessel. The objective of the terminal speed optimiser is to minimise the Equation 4.5. Where V is the set of discretised velocities at which the vessel can sail. The terminal speed optimiser chooses the velocity at which the vessel has a minimum cost. The velocity with the lowest expected cost is communicated to the terminal berth planner and the vessel operator. This consequently influences the sailing and arrival phases, as the vessel's sailing time is now dependent on the chosen velocity.

$$\min_{v \in V} C(v) \tag{4.5}$$

This strategy introduces limited coordination into the system. The vessel can now sail more slowly, avoiding the need to wait at the anchorage point. The coordination level is, however, limited, as the velocity of the vessels is only optimised at the moment of departure from the previous terminal.

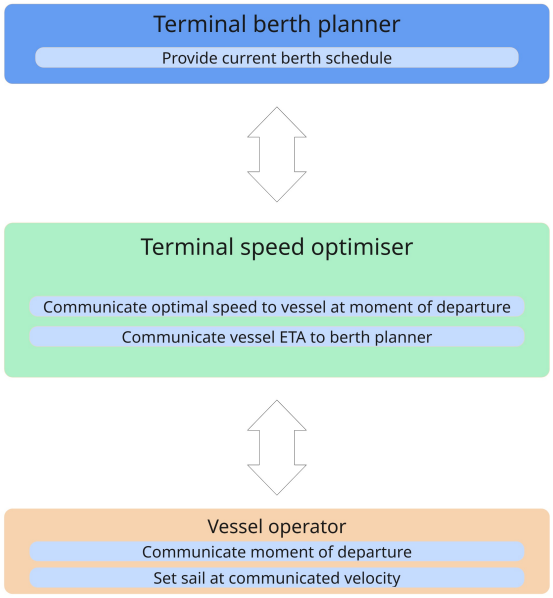


Figure 4.2: A schematic drawing of communications between stakeholders in the single speed optimisation strategy

Figure 4.2 illustrates the information exchange required for single-speed optimisation. Communication occurs only at vessel departure, after which the assigned speed remains fixed throughout the sailing leg.

4.2.3. Re-optimise at every new departure

In the single-speed optimisation strategy, when a vessel departs a terminal close to Rotterdam, all berth slots of the terminal can already be claimed by vessels that have departed from terminals further away. To address this issue, the strategy, re-optimise at every new departure, has been formulated. The key difference with the previous strategy is that vessel speeds can be re-optimised during sailing.

Whenever a new vessel departs from the terminal, the set of sailing vessels is modified. Vessels with a larger remaining distance than the newly departed vessel are reconsidered for speed optimisation and assigned updated velocities.

As the sailing leg of the vessels now consists of segments, the cost function and minimisation objective are now updated:

$$RC(v) = a \cdot R \cdot v^{b-1} \cdot c_f + t_w \cdot c_w + \frac{R}{v} \cdot c_o \tag{4.6}$$

The cost function changes over time as the vessel progresses on its sailing leg, with $C(v)$ becoming $RC(v)$ and D becoming R. In this case, R is the remaining distance to the port and $RC(v)$ is the cost over the remaining distance at velocity v.

The cost function used to calculate the final cost for a vessel is also different, as the fuel cost is now a sum of fuel costs. The total operating costs of a vessel would now look like this:

$$C = t_w \cdot c_w + t_o \cdot c_o + \sum_{i=1}^{n} a.S_i \cdot v_i^{b-1} \cdot c_f \tag{4.7}$$

$$t_o = \sum_{i=1}^{n} \frac{S_i}{v_i} \tag{4.8}$$

Where n is the number of segments, S_i is the length of a segment i in nautical miles, v_i is the velocity sailed at segment i in nautical miles per hour, and t_o is the operating time.

The goal of the method is now to solve the following equation:

$$\min_{v \in V} RC(v) \tag{4.9}$$

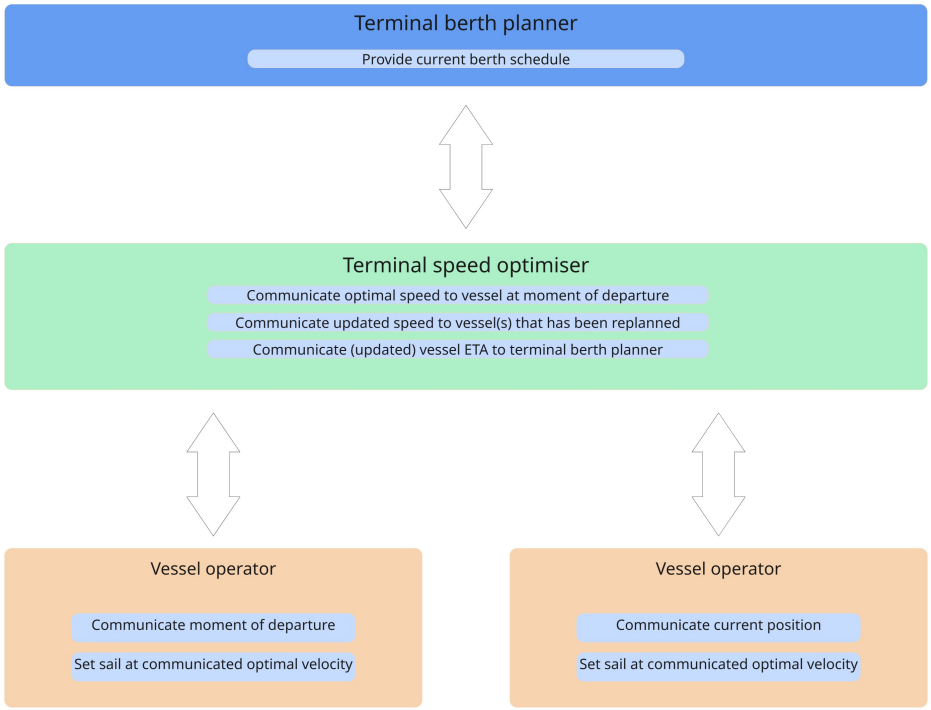


Figure 4.3: A schematic drawing of communications between stakeholders in the re-optimise at new departures strategy

This strategy allows the system to react to new information. It does, however, favour vessels departing from nearby. The strategy requires a higher level of coordination between vessels and the terminal, as a decision for one vessel can influence the velocity of others.

Figure 4.3 illustrates that vessels only receive updated speed recommendations whenever a new vessel departs. Compared to single-speed optimisation, this strategy therefore requires additional coordination between the terminal and sailing vessels.

4.2.4. Global speed optimisation

The global speed optimisation strategy extends the level of coordination by optimising the velocity of all sailing vessels together to minimise the overall cost. In this strategy, speed decisions are made both at new departures and at fixed planning intervals. By re-optimising at both new departures and at fixed planning intervals, the model devises a new schedule at new arrivals, and can update the schedule when the end time of vessels at the quay is updated.

At each planning moment, all vessels' speeds (except those within 50 nautical miles) are considered simultaneously. The goal of the function is no longer to optimise individual vessel velocities, but to minimise the cost of all sailing vessels:

$$\min_{v_j \in V_{j \in J}} \sum_{j \in J} RC_j(v_j) \tag{4.10}$$

In which J is the set of sailing vessels.

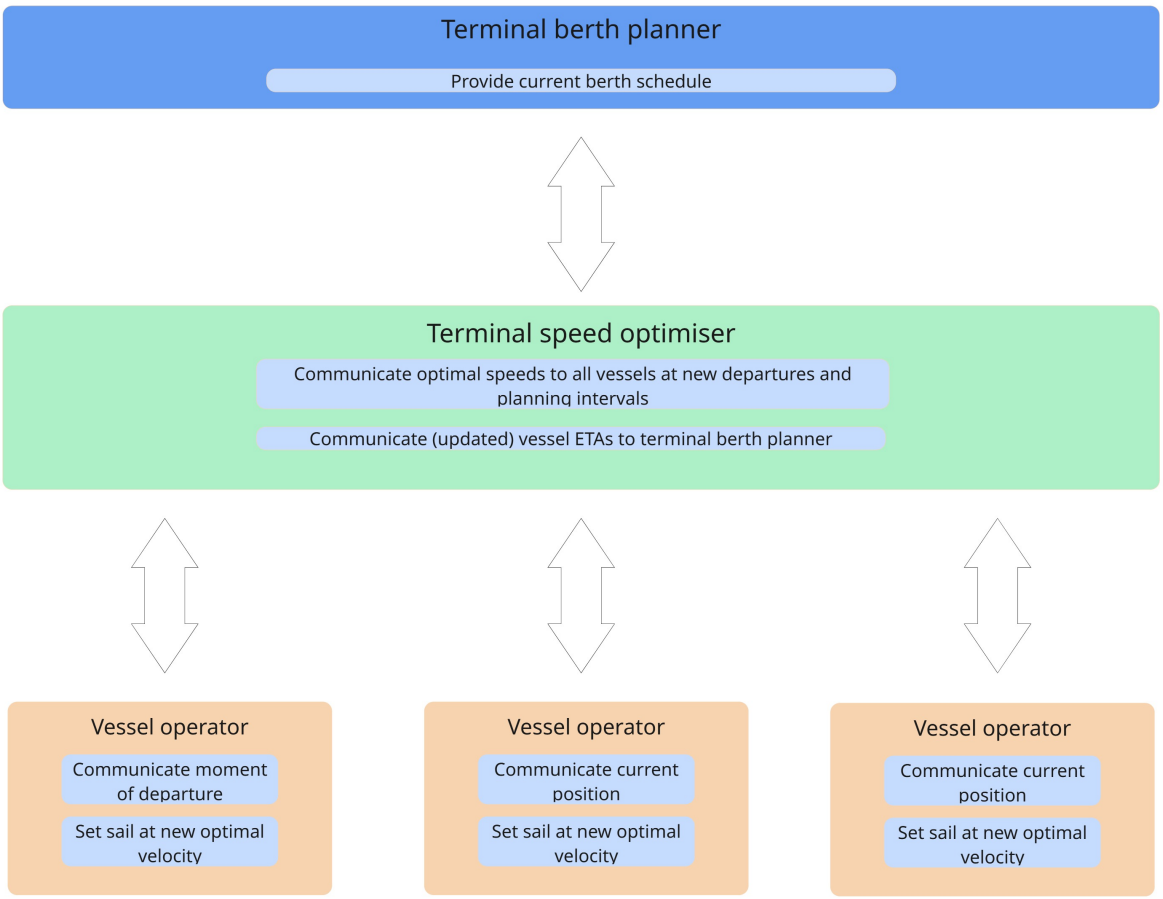


Figure 4.4: A schematic drawing of communications between stakeholders in the global speed optimisation strategy

To solve this problem, the optimisation is performed in two stages. First, the berth sequence is

determined by evaluating all possible arrival orders. This leads to J! number of orders. The order with the lowest costs is returned.

Given the arrival sequence, the vessels' speeds are optimised. From this sequence, all speed-up possibilities from every vessel (except those that arrive at the end of a sequence) are evaluated. The solution with the lowest overall cost is then selected.

This strategy requires the highest level of coordination, as individual vessel performance may be sacrificed to improve the overall system's performance.

Figure 4.4 shows the highest level of coordination. Information is exchanged both at vessel departures and at fixed planning intervals, allowing the terminal to continuously update vessel speeds and arrival sequences.

4.3. Overview

In Table 4.1, the baseline strategy and the three speed optimisation strategies are presented. In the overview, the strategies are compared with respect to the decision maker, the optimised parameters, the moments at which the vessel communicates with the terminal and the coordination level. The coordination level is the degree of interaction between the terminal and the sailing vessels. For the baseline, there is no interaction. For the single-speed optimisation, there is a low level of interaction. For the re-optimisation at new arrivals, there is a medium level of interaction, as the departure of a new vessel also influences the velocity of other sailing vessels. And in global speed optimisation, there is a high level of coordination, as the performance of a single vessel can be sacrificed for the overall system performance. This highlights the trade-off between individual vessel performance and overall system efficiency.

Strategy	**Baseline**	**Single-speed optimisation**	**Re-optimise at every new departure**	**Global speed optimisation**
Decision maker	Vessel operator	Terminal speed optimiser	Terminal speed optimiser	Terminal speed optimiser
Optimisation objective	Minimise cost of fuel consumption and operating costs	Minimise cost of fuel consumption, operating costs and waiting costs	Minimise cost of fuel consumption, operating costs and waiting costs	Minimise total costs of all sailing vessels
Communication moments	Departure only: ETA to berth planner	Departure only: recommend speed	New departures: updated speed advice	At new departures and planning intervals: updated speed advice
Coordination level	None	Low	Medium	High

Table 4.1: Overview of the baseline strategy and the different speed optimisation strategies

4.4. Conclusion

In this chapter, the following sub-question is answered: "Which vessel speed optimisation strategies can be identified and structured for use in simulation?" The sub-question is answered by first examining possible speed optimisation strategies in section 4.1. These speed optimisation strategies are then structured in section 4.2. And finally, an overview of the three strategies is given in section 4.3. The three identified vessel speed optimisation strategies are single-speed optimisation, re-optimise at every new departure, and global speed optimisation.

5

Model

In this chapter, the sub-question ”How can vessel speed optimisation be implemented in a discrete event simulation model that realistically represents real-life vessel operations?” is answered. The sub-question is answered by describing the requirements, inputs, outputs, performance indicators and operations of both subsystems defined in chapter 2. The model for the seaside operations is described in section 5.1, and the model for the sailing leg is described in section 5.2. The integration of both models is then described in section 5.3. The validation and verification of the combined model is described in section 5.4.

5.1. Model of the seaside operations of a container terminal

The first system modelled is the seaside operations at the container terminal. The TRAFALQUAR model of the seaside operations of a container terminal was developed and validated by Portwise. The terminal displayed in TRAFALQUAR is designed to model the operations of the deep-sea quay at Rotterdam World Gateway, as AIS data for vessels arriving at RWG is supplied, and the model's inputs can be validated against this data. A visualisation of the model is shown in Figure 5.1. The model does not include hinterland transport, yard operations or crane interference at a detailed level.

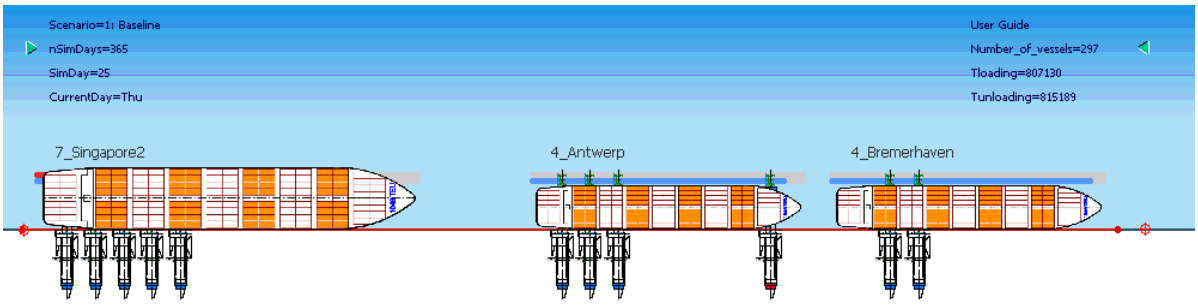


Figure 5.1: Visualisation of the seaside operation model

5.1.1. Model scope and goal

The scope of the seaside operations model is limited to the processes that take place from the moment a vessel arrives at the container terminal's anchorage until it departs after being serviced at the quay. The system includes vessels arriving at the anchorage point, waiting for a berth, being assigned to a quay, and undergoing loading and unloading operations.

The model starts when a vessel arrives at the anchorage point; it already has information on the vessel's time of arrival as soon as it departs the previous port. The model ends when a vessel has departed the quay. Individual containers are only modelled when they are loaded or unloaded from

the vessel at the quay. Horizontal transport of the container from the quay to the yard is not included in the model.

Based on these boundaries, the model assigns berths to vessels using information on all sailing vessels, waiting vessels, and vessels currently at the quay. This allows the model to represent berth planning decisions in a realistic environment. Sailing vessels are included in the berth planning as soon as they have departed their previous terminal and are within 1200 nautical miles (NM) of the terminal.

The goal of the model is to represent the seaside operations of Rotterdam World Gateway, so that the influence of different vessel speed optimisation strategies can be tested in a realistic environment.

Requirements
In order to test this, the model should have the following capabilities:

- The model should be able to generate vessels.
- The model should be able to assign a berth to a vessel based on time of arrival and priority or planned start of service.
- The model should be able to capture vessel service operations.
- The model should be able to estimate the end of operations for a vessel.
- The model should be able to determine service start and end time.
- The model should be able to calculate waiting time, service time and departure delay.
- The model should be able to calculate the berth occupancy.

Assumptions
The seaside operations of the container terminal have been simplified using the following assumptions:

- Quay crane operations are simplified to a single productivity parameter.
- Quay cranes can be reassigned between vessels, dependent on operational necessity.
- Vessels can not be repositioned on the quay as soon as they have arrived.
- Tidal restrictions are not modelled.
- The journey time from anchor point to quay and back takes a constant amount of time, during which travelling vessels do not interfere with each other.
- Vessels will wait until they are serviced.

5.1.2. Model inputs

The model has three types of inputs: vessel arrival information, vessel characteristics, and quay information.

Vessel arrival

- Seasonal call-size pattern
- Vessel arrival schedule

Vessel characteristics

- Vessel length l_j
- Priority
- Arrival delay distribution
- Call-size
- Call-size distribution
- Number of bays
- Expected productivity

Quay

- Quay length l_q
- Number of cranes
- Productivity of cranes

5.1.3. Outputs and performance indicators

The model has the following outputs and performance indicators:

- Departure delay d
- Waiting time t_w
- Service time t_s
- Berth occupancy Q

Of which the departure delay, waiting time and service time influence vessel performance and the service level of the terminal. While the berth occupancy influences the terminal's performance.

5.1.4. Operations

The seaside operations at a container terminal are divided into two main tasks: vessel berth assignment and vessel service. In this section, the operations of both tasks are described using swimlane diagrams. In the swimlane diagram, the inputs and outputs are linked. These swimlane diagrams are linked to the swimlane diagrams presented in chapter 4. As the final method of the swimlane of the sailing leg is the first method of the swimlane of the berth assignment. And the final process of the swimlane of the berth assignment is linked to the first method of the swimlane of the service.

Berth assignment

The process of the berth assignment is displayed in Figure 5.2. In the swimlane, the first method is "Arrive at anchorage point", which is the final method of the swimlane of the sailing leg in Figure 5.4. The terminal berth planner assigns berths to all sailing vessels, waiting vessels, and vessels at the quay. If, according to the calculated schedule, a berth is available for a waiting vessel, the waiting vessel is sent to the quay.

In the model, two different berth assignment methods are used. The first method is *first-come, first-served*, which is used for the baseline strategy. A vessel arrives at the anchorage point of the terminal. If there are other vessels at the anchorage point, the vessel is placed in the queue. The queue is sorted first by time of arrival and then by priority. Vessels with a higher priority are put in front of the queue. When a vessel arrives at the anchorage point or when another vessel departs the quay, a check is executed to determine whether the first vessel in the queue can berth. To berth, a slot must be available that is larger than the vessel's length, plus a safety margin times 2, to ensure that vessels do not sit directly next to another. If the first vessel in the queue cannot berth, it is checked whether another vessel in the queue can. However, a vessel with lower priority is not allowed to berth if doing so would increase the waiting time of a vessel with higher priority. If the vessel can berth, a signal is sent to the vessel in the queue, which then travels to the terminal's quay.

The other berth assignment method is dependent on the estimated time of service, which is used for all three speed optimisation strategies. The estimated time of service is determined by the terminal speed optimiser. The estimated time of service is the time, at the moment a velocity is determined, when the vessel is expected to begin berthing. By sorting vessels in the queue at the anchorage point by estimated start of service, vessels that arrived later but have an earlier expected time of service will be serviced before those that arrived earlier. This ensures that vessels are berthed in the order that is determined by the terminal speed optimiser.

The first performance indicator for the seaside model is the waiting time, which is the time a vessel spends at the anchorage.

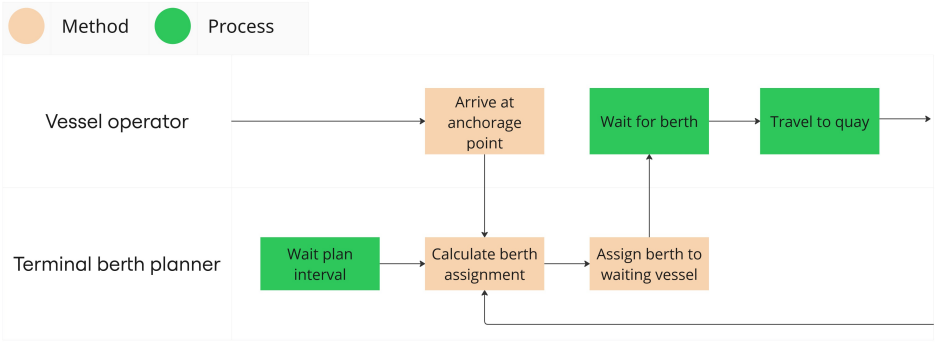


Figure 5.2: Swimlane diagram of the berth assignment

Service

The process of the service is displayed in Figure 5.3. During the service, all containers are loaded and unloaded from a vessel. These containers are divided over a number of bays. Only one crane can work on a bay at a time. One crane works at a fixed productivity and moves a fixed number of containers per hour. The crane continues operating until a scheduled break is taken. The number of cranes assigned to a vessel depends on their availability and the expected productivity.

When all containers are removed from the vessel, the vessel can depart the quay, and the second and third performance indicators for the seaside operations are achieved. The service time is the time a vessel spends at the quay. The departure delay is defined as the actual departure time minus the planned time of departure from the proforma schedule. The planned departure time is determined from the estimated time of arrival and the expected time of service, without the realisation of stochastic parameters.

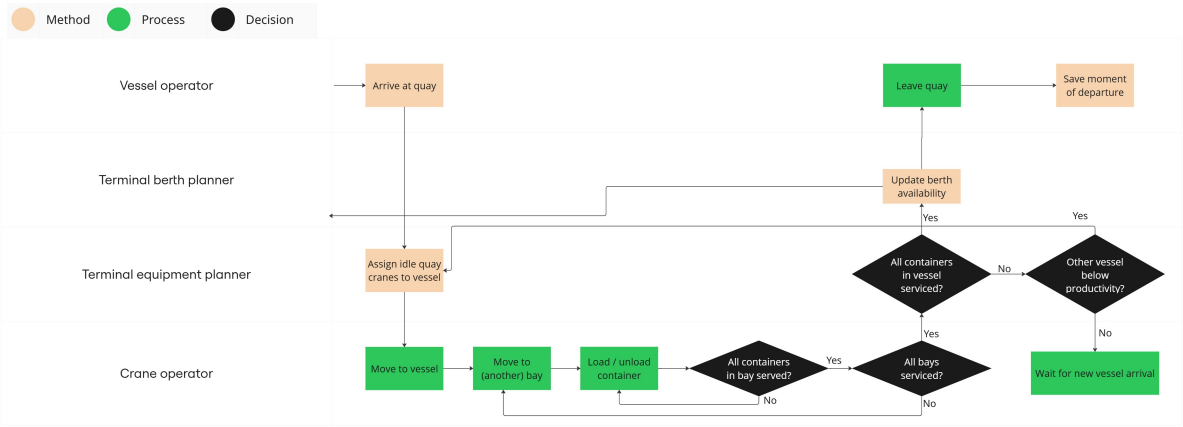


Figure 5.3: Swimlane diagram of the loading and unloading of containers at the quay

The final performance indicator for the seaside operations, the berth occupancy, is determined after a year of simulation. It is defined using Equation 5.1.

$$Q = \frac{\sum_{j \in J} t_{s,j} l_j}{l_q T} \tag{5.1}$$

In which T is the total simulation length, and J is the set of all vessels.

5.2. Model of the sailing leg

The second system that is modelled is the sailing leg of vessels. The sailing leg is defined as the moment at which a vessel departs a previous terminal, to the moment it arrives at the anchorage

point of the new terminal. During the sailing leg, different strategies can be taken. The steps during the strategies are displayed in operations. For each strategy, the model uses the same input, allowing the influence of each strategy to be determined.

5.2.1. Model scope and goal

The scope of the sailing leg model is limited to the movement of vessels from departure at the previous terminal to arrival at the anchorage point of the destination terminal. The model includes speed decisions, which influence sailing time and fuel consumption, but excludes the effects of routing decisions, weather, and interactions with other vessels at sea. Terminal operations, such as berthing and cargo handling, are not included in this model; however, they are indirectly influenced by the arrival times generated by the sailing leg model.

The goal of the model is to represent the sailing time and fuel consumption of a sailing vessel. The sailing leg model determines the arrival time of vessels by selecting sailing velocities that minimise a cost function under current terminal conditions.

The previous ports are selected to form an accurate representation of the ports of vessels arriving at RWG.

Requirements

To test the different strategies, the model should have the following capabilities.

- The model should be able to calculate the moment of arrival of a vessel.
- The model should be able to calculate the cost for the moment of arrival of a vessel.
- The model should be able to let vessels sail at an average velocity.
- The model should be able to select the velocity with the lowest possible arrival cost for a single vessel.
- The model should select the velocities for all sailing vessels that minimise the combined cost.

Assumptions

The sailing leg of the vessels is simplified using the following assumptions.

- Vessels can sail at exactly the correct velocity.
- The velocity will remain constant until changed by the model.
- Weather conditions are not modelled.
- Engine breakdowns are not modelled.
- Vessels do not find interference from other vessels.
- Transfer of information between the terminal and the vessel is instantaneous and accurate.

5.2.2. Model inputs

The model has the following inputs:

- Information on all sailing vessels, waiting vessels and vessels at quay
- Vessel parameter b
- Distance D
- Remaining distance R
- Fuel cost c_f
- Discretised velocities V_j
- Optimal velocity $\tilde{v}$
- Fuel consumption at average velocity $\tilde{f}$

5.2.3. Outputs and performance indicators

The model has the following output and performance indicators:

- Total costs C_t
- Fuel costs C_f
- Operating costs C_o
- Waiting costs C_w
- Mean velocity
- Number of velocity changes

5.2.4. Operations

The operations of the model are dependent on the chosen speed optimisation strategy. These strategies are described in section 4.2. In the model, the sailing leg of a vessel is modelled. The time a vessel has to sail depends on the chosen speed optimisation strategy. The speed optimisation strategy determines which vessels are included in the optimisation and at which moments speed decisions are updated. An overview of the optimisation strategies is shown in Table 4.1. A swimlane diagram of the baseline operations is shown in Figure 5.4. In baseline operations, the vessel selects the velocity that minimises costs over the sailing leg, without accounting for waiting time.

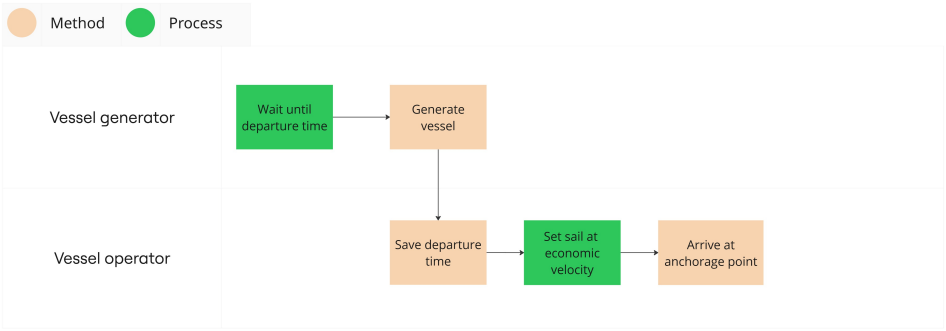


Figure 5.4: Swimlane diagram of the sailing operations of the vessel under the baseline speed optimisation strategy

In the swimlane diagram displayed in Figure 5.5, the operations of a vessel are displayed under the single speed optimisation strategy. Compared to the baseline strategy, the vessel no longer selects a velocity; instead, it is assigned a velocity by the terminal speed optimiser. This velocity is optimised to minimise waiting, operating, and fuel costs over the sailing leg.

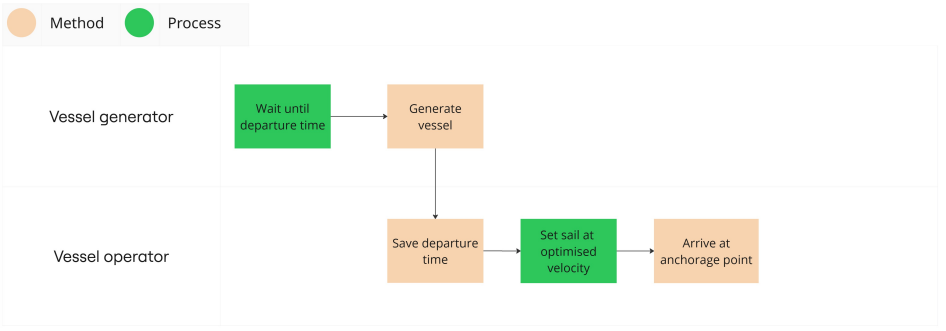


Figure 5.5: Swimlane diagram of the sailing operations of the vessel under the single-speed optimisation strategy

In the swimlane diagram displayed in Figure 5.6, the operations of a vessel are displayed under the re-optimise at new departures and global speed optimisation strategies. When compared to the swimlane diagram in Figure 5.5, the decision, re-optimisation triggered, is now added to the diagram.

The difference between the strategies lies in the trigger for re-optimisation and the objective of the optimisation. In re-optimise at every new departure, re-optimisation is triggered whenever a new vessel departs and has fewer nautical miles remaining. In global speed optimisation, re-optimisation is triggered both by new vessel departures and by planning intervals. If re-optimisation is triggered, the vessel will have to set sail at the newly optimised velocity.

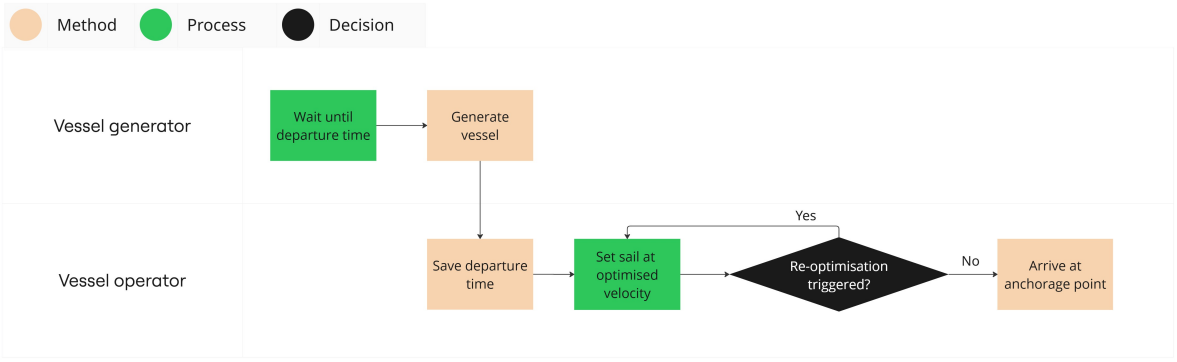


Figure 5.6: Swimlane diagram of the sailing operations of the vessel under the re-optimise at new departures, and global speed optimisation strategy

5.3. Model integration

The models described in section 5.1 and section 5.2 are combined. In this section, the integration of the model and the planning frequency are discussed.

5.3.1. Integration

The model integration consists of three parts. The determination of when a vessel departs the previous terminal, the shared information between vessels and the terminal and the moment of arrival at the anchorage point.

Vessel departure

At the start of the simulation in TRAFALQUAR, all vessel arrival times are determined. The ATA is determined by adding the arrival delay to the ETA. The departure time at the previous terminal is determined by dividing the distance to the previous port by the vessel's average speed and subtracting the result from the determined ATA. For vessels that would depart before the simulation starts, the departure time is set to 0:01:00.

As soon as a vessel has departed the previous terminal, it determines its velocity, based on one of the speed optimisation strategies described in section 5.2.

Shared information

The terminal has knowledge of the velocity, remaining distance, call size and length of all incoming vessels. The information is updated at every optimisation moment.

Vessel arrival

When a vessel arrives at the terminal, after having set sail at the velocity determined by the speed optimisation strategy, it arrives at the anchorage point. As soon as it has set sail, the vessel will be in the planning horizon of the berth planner of TRAFALQUAR. Based on the chosen method of operation (FCFS with priority for the baseline and sorted by planned time of service for speed optimisation), the method will determine a berth order. Based on the berth order, the model will assign berths to the vessel, and normal loading and unloading operations will begin.

5.3.2. Planning frequency

Every two hours, TRAFALQUAR's berth planning is updated. During this update, the estimated departure times of vessels currently at the quay are recalculated based on the number of remaining

containers. The estimated service start times of vessels waiting at the anchorage point and of sailing vessels are recalculated. The estimated service start times cannot be earlier than the vessel's arrival time.

The planning frequency of the speed optimisation model is dependent on the chosen model. For speed optimisation at departure, the speed is determined only at a vessel's departure. For replanning, a vessel's speed can be updated when another vessel departs from a terminal with a shorter remaining distance. For global optimisation, the vessel's speed can be updated at every new departure or at a fixed replanning interval.

5.4. Verification and validation

In this section, the model is verified and validated. According to Robinson [22], the aim of verification and validation is not to prove that a model is correct; rather, it is to prove that the model is incorrect. If this cannot be proved, then verification and validation increase confidence in the model and the results.

Verification ensures the model is built correctly, while validation ensures it is sufficiently accurate. TRAFALQUAR is a verified and validated model which simulates the seaside operations of vessels arriving at the quay. To verify the integrated vessel speed optimisation model, the decision logic and event flow of the combined model are verified. The model's input is then validated by ensuring that the baseline model produces waiting and service times that are similar to those observed in the AIS data.

5.4.1. Verification of the speed optimisation function

According to Robinson, the model can be verified by checking the code, performing a visual check, and inspecting the output of reports [22].

Several visual checks have been built into the code to ensure that the model behaves as expected. An example of built-in checks is displayed below:

1. RemainingDistance in SailingVessels < -0.1
2. |NumberOfVessels - QuayClaimed| > 0
3. ChosenVelocity > MaxVelocity
4. CostSpeedUp - CostOrder < 1
5. |CostCheck - CostSpeedUp| > 1

The function of check one is to break the model, if there are any vessels with a negative remaining distance; check two breaks the model, if the total number of vessels is not equal to all vessels in the planning horizon in QuayClaimed; check three breaks the model, if the chosen velocity of a vessel is larger then the maximum velocity of the vessel; check four breaks the model, if the lowest entry of speeding up, is higher then the lowest entry of the order; and check 5 breaks the model, if the realised costs, are higher then the cost of speeding up. Check 4 and 5 indicate that the wrong order has been chosen.

Besides code checks, other visual checks from Robinson [22], such as stepping through the model event by event and tracing the progress of an item (vessel) through the model, have also been performed.

In addition to visual checks, the output reports are also inspected and compared with the expected results. Examples of inspected output reports include the expected waiting and fuel costs, which need to be identical to the expected waiting and fuel costs, as the model directly influences them. The actual waiting costs may differ from the expected waiting costs, since the waiting time at the moment of optimising the velocity is not fully determined, as a vessel can still depart earlier or later than expected.

Besides costs, the sailing time of a vessel is also evaluated alongside the velocity. As the sailing time of a vessel is between a minimum and maximum value, it is dependent on the minimum and maximum velocity of the vessel.

While the examples presented in this section illustrate the verification approach, a more extensive verification of the model is provided in Appendix C. In this verification chapter, examples of checks on output reports are provided for the realised berth schedule, the optimisation outcomes from both single-speed and global-speed optimisation, and berth occupancy calculations are examined in detail.

5.4.2. Validation

The model is validated using black-box validation. According to Robinson [22], confidence should be placed in the model when the outputs of the model are sufficiently similar to those of the real world, if it is run under the same conditions as the real world system.

The TRAFALQUAR model without speed optimisation is validated using AIS data from deep-sea and short-sea vessels. The validation focuses on whether the model accurately reproduces the observed waiting and service times of deep-sea vessels at RWG. The AIS data is used to represent the inputs and the outputs of the real-world system. In the AIS data, information is available on all vessels arriving at RWG in 2024. For each vessel, the following data can be extracted:

1. Previous port
2. Vessel type (TEU Class)
3. Vessel length
4. Waiting time
5. Service time

In the AIS data, no information was provided on the number or service time of the incoming barges. Information on the quay length and the number of cranes was obtained from RWG's website [7]. The number of cranes on the deep-sea quay has been reduced from 14 to 11, as not all are in operation at the same time.

From the AIS data, a set of vessel arrivals was derived every two weeks. Using the average service time of the deep-sea vessels, the seasonal pattern is also derived from the AIS data.

The unknown parameters, such as crane productivity, expected vessel productivity, expected arrival delay, call-size distribution, the number of barges per day, and the average service time of barges, were calibrated using parameter values from previous Portwise simulation studies of container terminals.

The parameters that have been changed are vessel priority, vessel call size, and the vessel arrival schedule.

Vessel priorities were assigned such that vessel classes with lower observed waiting times in the AIS data received higher priority levels.

The vessel's call size was calibrated such that the simulated service times match those observed in the AIS data. By iteratively increasing the call size, the service time per vessel reached comparable values to those in the AIS data. The average vessel service time per vessel class from the AIS data, the simulation with deep-sea and short-sea vessels, and the simulation with only deep-sea vessels are displayed in Figure 5.7.

The total waiting time of all deep-sea vessels is selected as the main validation metric of the input schedule, as it directly reflects the interaction between vessel arrivals and berth allocation. After the calibration, the model's outputs are compared to AIS data for validation.

In order to determine the influence of the short sea vessels on the waiting time of the deep-sea vessels, two sets of 10 simulations were conducted. One that included all vessels, and another that included only the deep-sea vessels (TEU 4, TEU 5, TEU 6, and TEU 7). In the second set of simulations, the second quay was not simulated; further, all parameters remained the same. For each vessel category, the waiting time for both sets of simulations is shown in Figure 5.8, and the service time is shown in Figure 5.7.

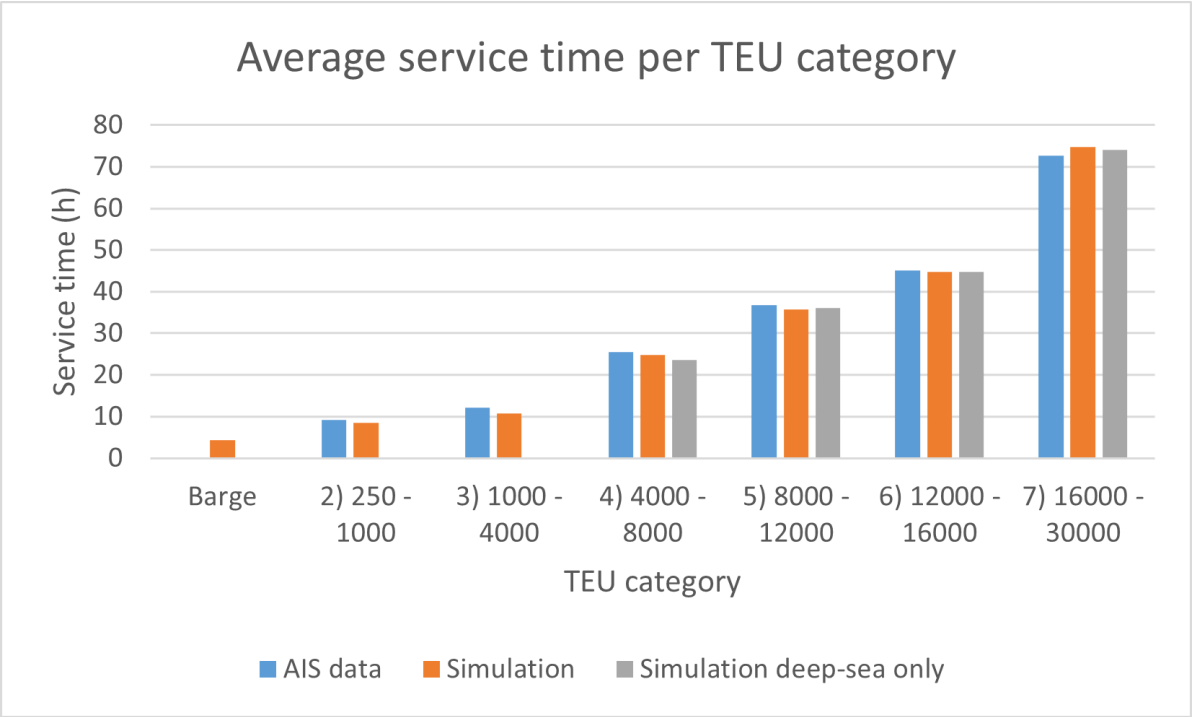


Figure 5.7: Average vessel service time per TEU category

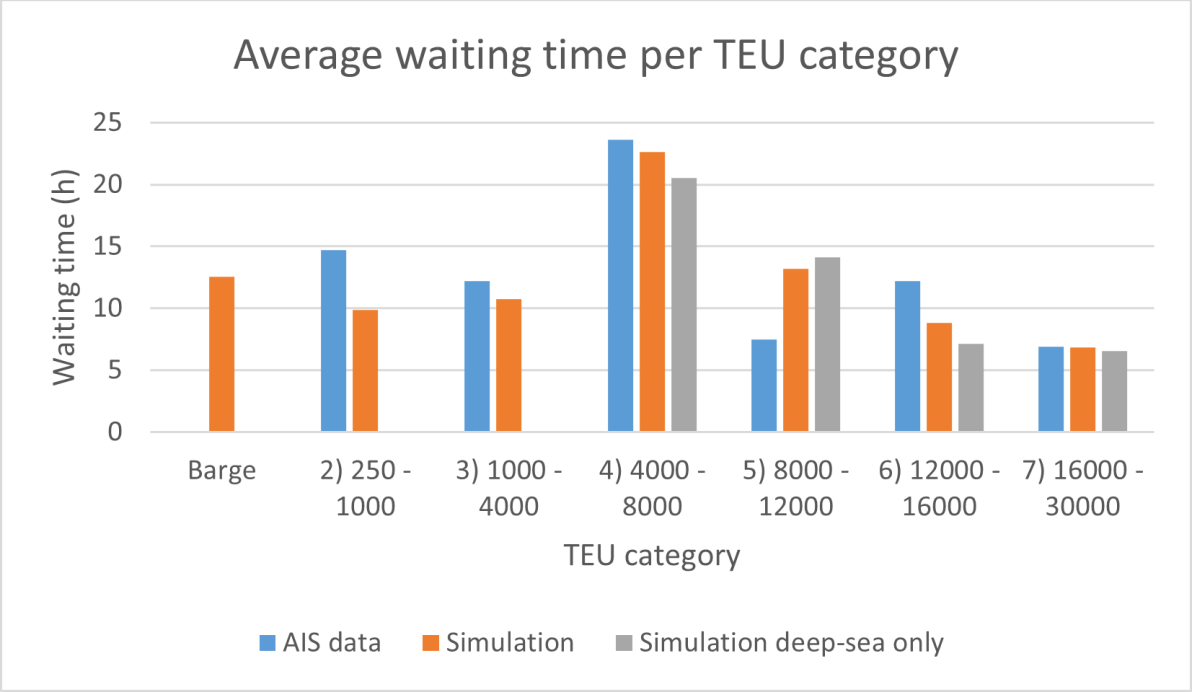


Figure 5.8: Average vessel waiting time per TEU category

Vessel class	AIS data	Simulation input
TEU 4	69	78
TEU 5	34	26
TEU 6	45	52
TEU 7	131	130

Table 5.1: Number of vessel arrivals per class from AIS data and simulation

	Total waiting time of all deep-sea vessels (h)
AIS Data	3322
Simulation with barge and short sea	3447
Simulation without barge and short sea	3189

Table 5.2: Total waiting time of all deep-sea vessels

The number of vessels in the simulation is not exactly the same as in the AIS data, because when dividing the arrivals per vessel class by the number of weeks (Table 5.1), the result is not an integer. This is why, for total waiting time, not the waiting time per vessel class, but the overall waiting time of all deep-sea vessels is taken as the indicator to which the simulation is tuned. The total waiting time is used as a performance measure and is compared for the AIS data and the two simulation sets in Table 5.2.

The total waiting time across both sets of simulations is similar, indicating that the waiting time of deep-sea vessels is not significantly affected by adding short-sea vessels and barges to the simulation. However, the simulation with barges and short-sea vessels overshoots the overall waiting time, while the simulation without barges undershoots the overall waiting time. This is mainly due to the difference in waiting time in the TEU 4 vessel class, as it experiences a drop of 2.5 hours in waiting time, while the other vessel classes remain similar. As the addition of short-sea vessels and barges significantly affects the simulation's computation time, these vessels are not included in the speed-up scenarios. The total waiting time and service time for deep-sea vessels of the scenario without barges and short sea vessels closely match the waiting time and service time of the AIS data, indicating that the model adequately represents real-world terminal operations. Based on the comparison, the model is considered sufficiently accurate for analysing the impact of vessel speed optimisation strategies.

5.5. Conclusion

This chapter answers the sub-question "How can vessel speed optimisation be implemented in a discrete event simulation model that realistically represents real-life vessel operations?" by first introducing the model for the seaside operations of a container terminal, and secondly introducing the model for vessel speed optimisation. The combination of these two models is described in section 5.3. In section 5.4, the decision logic of the model is verified, and the model is validated using the AIS data. The verified and validated model implements vessel speed optimisation in a realistic scenario, thus answering the research question.

6

Experimental Plan

The experimental plan defines a set of experiments to answer sub-question 5, ”How do different vessel speed optimisation strategies affect the performance of container vessels?” and sub-question 6: ”How can vessel speed optimisation influence berth occupancy of a container terminal?” First, the model input is defined in section 6.1. Then the key performance indicators of the experiments are discussed in section 6.2. Next, the experiments to answer sub-question 5 are presented in section 6.3, followed by those for sub-question 6 in section 6.4. Finally, a set of sensitivity experiments is introduced in section 6.5, in which different input parameters are tested.

6.1. Input

The model’s input is divided into two parts: scenario inputs, which describe the external arrival processes, and configuration inputs, which describe the system and decision logic.

6.1.1. Scenarios

The input for the base scenario is validated in section 5.4 using AIS data. The additional input for the vessel speed optimisation model is described below. The input is divided into two parts: the arrival schedule and the fuel consumption.

Arrival schedule

The main input for the scenarios is the vessel’s arrival. Three scenarios are defined. The first is a scenario based on AIS data for deep-sea vessel arrivals at RWG in 2024. The scenario is based on a biweekly schedule. The two other scenarios being tested are 1 additional vessel arriving every 2 weeks and 2 additional vessels arriving every 2 weeks. Information about the vessel arriving is presented in Table 6.1. Vessels of which the day of arrival is marked with one star (*) are added in scenario 2, while vessels marked with one star (*) and two stars (**) are included in scenario 3. Leading to 11 vessels arriving every two weeks in scenario 1, 12 vessels arriving every two weeks in scenario 2, and 13 vessels arriving every two weeks in scenario 3.

The previous ports of the vessels are derived from the AIS data. The corresponding distances between the previous port and RWG are obtained from Sea-distance.org [26], with a maximum of 1200 NM. When vessels are sailing at 16 NM/h, 1200 NM is travelled in 75 hours. The vessel-specific input parameters are summarised in Table 6.1.

Vessel class	Previous port	Vessel length (m)	Call size (TEU)	Expected productivity (Box/h)	Distance to previous port (NM)	Day of arrival
TEU 7	Singapore	400	8500	110	1200	1
TEU 4	Le Havre	285	1650	60	250	1*
TEU 4	Antwerp	285	1650	60	140	3
TEU 7	Tanger	400	8500	110	1200	4
TEU 6	Algeciras	365	5000	110	1200	6
TEU 7	Dunkerque	400	8500	110	1400	7
TEU 4	Bremerhaven	285	1650	60	250	8
TEU 7	Singapore	400	8500	110	1200	9
TEU 5	Antwerpen	345	3300	90	140	10**
TEU 6	Gdansk	365	5000	110	850	11
TEU 7	Algeciras	400	8500	110	1200	12
TEU 5	Southampton	345	3300	90	250	13
TEU 4	Hamburg	285	1650	60	300	14

Table 6.1: Vessel input parameters; the days of arrival marked with a star (*) are part of scenario 2 and scenario 3

Fuel consumption

The fuel consumption parameters are based on the research by Brouer et al [5]. A summary of the original parameters used in their research is reported in Table 6.2.

Vessel Class	Vessel capacity (FFE)	**s*** (NM/h)	**Fs*** (mt/day)
Feeder_450	450	12	18.8
Feeder_800	800	14	23.7
Panamax_1200	1200	18	52.5
Panamax_2400	2400	16	57.4
Post Panamax	4200	16.5	82.2
Super Panamax	7500	17	126.9

Table 6.2: Fuel consumption parameters used by Brouer et al. [5]. FFE: Forty-foot equivalent unit, s*: vessel design speed in NM/h, Fs*: bunker consumption at design speed in metric tonnes per day

Note that the vessel capacities in Brouer et al [5] are expressed in FFE (forty foot equivalent units), rather than in TEU. Using these parameters, the fuel consumption values are scaled to a design speed of 16 NM/h ($\tilde{v}$) using Equation 6.1. To reflect the larger modern vessels, the parameters of the TEU 6 and TEU 7 classes are scaled up accordingly. The fuel parameter is b, which is set to 3 for large deep-sea vessels. The resulting fuel consumption parameters used in this study are presented in Table 6.3.

$$\tilde{f} = \frac{\tilde{v}^b}{s*^b} \cdot f_{s*} \tag{6.1}$$

Vessel class	Maximum capacity (TEU)	$\tilde{v}$ (NM/h)	$\tilde{f}$ (mt/day)
TEU 4	8000	16	75
TEU 5	12000	16	100
TEU 6	16000	16	125
TEU 7	30000	16	200

Table 6.3: Fuel consumption parameters used in this research.

The cost of one metric ton of fuel is taken from Ship & Bunker. The average cost of one metric tonne of Very Low Sulphur Fuel Oil in Rotterdam during 2024 was $555.5 dollar [25], which is used as c_f in the simulation.

6.1.2. Configurations

The configurations consist of two parts: the speed optimisation model and the quay configurations.

Speed optimisation configurations

The model is tested across 4 different configurations. These configurations are described in chapter 4. These four different speed optimisation configurations are:

1. Baseline strategy
2. Single-speed optimisation
3. Re-optimise at every new departure
4. Global speed optimisation

Quay configurations

The quay is supposed to represent the deep-sea quay of RWG. The quay is 1150 meters long. On the quay, 11 deep-sea cranes are stationed, each capable of moving 26 containers per hour. These cranes work continuously in shifts for 7:30 hours, then take a 0:30-hour break. Each crane can only work on 1 bay of the vessel. Each vessel can berth anywhere on the quay; however, once berthed, it cannot be moved.

6.2. Key performance indicators

Not all of the performance indicators produced by the model, as described in subsection 5.1.3 and subsection 5.2.3, are used as key performance indicators in the experimental analysis. The model's key performance indicators are total vessel costs (divided into fuel, operating, and waiting costs), departure delay, and, for the different scenarios, the quay occupancy. The vessel's costs and departure delay combined serve as a benchmark for a container vessel's performance. The vessel's costs indicate the operator's willingness to adhere to the optimised velocity, and the departure delay indicates whether vessels depart later than in the current situation.

The total vessel costs consist of the sum of fuel costs, operating costs and waiting costs over the entire operation period, as defined in Equation 4.7. Costs are reported both per vessel and averaged across all vessels.

The departure delay is defined as the actual time of departure minus the estimated time of departure. The departure delay is reported per vessel and is also averaged across all vessels. The estimated time of departure is derived from the proforma schedule and equals the planned arrival time plus the expected service time, without any realised stochastic delays. This definition is chosen because the vessel speed optimisation directly influences arrival times. Using the estimated arrival plus the realised delay would therefore partly remove the effect that the optimisation strategies intend to achieve. By comparing the actual departure times with the planned departure time from the schedule, the KPI measures the total delay to the planned operation.

The final performance indicator for the different scenarios is the berth occupancy. The berth occupancy is defined as the fraction of total quay length occupied over time, and is calculated using Equation 5.1. The berth occupancy is only calculated once per simulation. It is used to indicate the increase in berth occupancy caused by the additional vessels arriving in scenarios 2 and 3. The berth occupancy would remain the same in a given scenario, since the same number of vessels have arrived. However, slight deviations may occur, as a vessel delayed by stochastic factors may not begin service before the end of the simulation.

6.3. Influence of speed optimisation strategies

In this section, a set of experiments is defined to address sub-question 5: "How do different vessel speed optimisation strategies affect the performance of container vessels?" To answer the sub-question, the four different configurations of subsection 6.1.2 will be compared in scenario 1 of subsection 6.1.1. They will be compared in scenario 1, as it serves as an accurate representation of the vessels currently arriving at RWG. Leading to the following set of simulations:

1. Baseline configuration, scenario 1

2. Single-speed optimisation, scenario 1
3. Re-optimise at new departures, scenario 1
4. Global speed optimisation, scenario 1

The performance of the vessel speed optimisation strategies is described by the total vessel costs. The two KPIs used to test the model are average vessel costs (waiting, operating, and fuel) and departure delay. These KPIs represent vessel performance and the terminal's service level. These KPIs are displayed as averages over the entire scenario and per previous port. The average across an entire scenario indicates improvement in the entire system, while the average across a previous port indicates local improvement or deterioration. Every configuration is simulated 10 times to account for the system's stochastic variability and to reduce the influence of outliers on the results. Based on 10 simulation runs, 95% confidence intervals are constructed for each KPI. Pairwise Welch's t-tests (two-sample, unequal variance) are performed on the simulation averages of a KPI to determine whether differences between configurations are statistically significant.

6.4. The performance of global speed optimisation under scenarios 1, 2 and 3

In this section, a set of experiments is defined to address sub-question 6: "How can vessel speed optimisation influence berth occupancy of a container terminal?" To test the influence of vessel speed optimisation on berth occupancy, the two additional scenarios (2 and 3) defined in subsection 6.1.1 will be applied to the global speed optimisation configuration described in subsection 6.1.2. Based on the conceptual design of the strategies, the global speed optimisation configuration is expected to provide the highest level of coordination between vessels and the terminal. Therefore, the influence of the different scenarios is only evaluated for this configuration. The baseline configuration in scenario 1 serves as the benchmark for current operations at the container terminal. The results of the three scenarios under the global speed optimisation configuration will be compared to investigate how the system performs when additional vessels need to be handled, leading to the following set of simulations, of which 1 and 2 are not repeated, as their results have already been produced by the experiments in section 6.3:

1. Baseline configuration, scenario 1
2. Global speed optimisation, scenario 1
3. Global speed optimisation, scenario 2
4. Global speed optimisation, scenario 3

The two new scenarios are simulated 10 times. These four combinations of scenarios and configurations are compared on the same KPIs as the vessel speed optimisation strategies. However, the additional vessels are excluded from the average vessel cost calculation as their short sailing distances and relatively low costs would bias the comparison with scenario 1. In addition to these KPIs, the berth occupancy is also measured.

6.5. Sensitivity analysis

Besides experiments to answer the main research question and the sub-questions, a set of sensitivity tests is devised. The sensitivity tests are displayed in Table 6.4. The first two sensitivity tests focus on the fuel consumption parameters displayed in Table 6.3. These parameters are found by using values from earlier research [5]. However, no values were provided for vessels up to 30000 TEU, so the fuel consumption parameters had to be scaled up for the larger vessels. As no actual parameters were provided, it is relevant to assess how sensitive the model results are to these assumptions. In sensitivity tests 1 and 2, the fuel consumption parameter $\tilde{f}$ is changed. In test 1, the difference in $\tilde{f}$ between vessel classes is relatively larger, while in test 2, the difference is relatively smaller. Another interesting parameter is in Equation 4.3, where the waiting costs for all vessels are 1.1 times their operating costs. In sensitivity 3, this value is set to 1.5 for the largest vessels (TEU 6 and TEU 7) to assess the influence on departure delay.

Every sensitivity test is simulated 10 times, under scenario 1, with the global speed optimisation

configuration. These sensitivity tests are only compared on the average departure delay of the vessel, as changes in cost parameters no longer allow comparison on costs. The departure delay is presented for the overall system, the smaller vessels (TEU 4 and TEU 5), and the larger vessels (TEU 6 and TEU 7). The overall average reflects system-level effects, while the per-vessel captures the results per class.

Vessel class	Global speed optimisation	Sensitivity 1	Sensitivity 2	Sensitivity 3
TEU 4	$c_w = 1.1 \cdot c_o$, $\tilde{f} = 75$	$c_w = 1.1 \cdot c_o$, $\tilde{f} = 75$	$c_w = 1.1 \cdot c_o$, $\tilde{f} = 75$	$c_w = 1.1 \cdot c_o$, $\tilde{f} = 75$
TEU 5	$c_w = 1.1 \cdot c_o$, $\tilde{f} = 100$	$c_w = 1.1 \cdot c_o$, $\tilde{f} = 125$	$c_w = 1.1 \cdot c_o$, $\tilde{f} = 85$	$c_w = 1.1 \cdot c_o$, $\tilde{f} = 100$
TEU 6	$c_w = 1.1 \cdot c_o$, $\tilde{f} = 125$	$c_w = 1.1 \cdot c_o$, $\tilde{f} = 175$	$c_w = 1.1 \cdot c_o$, $\tilde{f} = 100$	$c_w = 1.5 \cdot c_o$, $\tilde{f} = 125$
TEU 7	$c_w = 1.1 \cdot c_o$, $\tilde{f} = 200$	$c_w = 1.1 \cdot c_o$, $\tilde{f} = 300$	$c_w = 1.1 \cdot c_o$, $\tilde{f} = 125$	$c_w = 1.5 \cdot c_o$, $\tilde{f} = 200$

Table 6.4: Changed variables for sensitivity test

7

Results

In this chapter, the results of the experiments described in chapter 6 are presented. Using these results, a conclusion is drawn on sub-question 5, ”How do different vessel speed optimisation strategies affect the performance of container vessels?” and sub-question 6: ”How can vessel speed optimisation influence berth occupancy of a container terminal?” In section 7.1, the results of the different speed optimisation strategies are presented, and a conclusion is drawn on sub-question 5. In section 7.2, the results of additional vessels are discussed, and a conclusion is drawn on sub-question 6. In section 7.3, the results of the sensitivity tests are discussed. Finally, in section 7.4, the conclusions on the results are summarised.

7.1. Influence of speed optimisation strategies

First, the three different speed optimisation configurations are compared to the baseline configuration at scenario 1. Every set of simulations is simulated 10 times. Based on these runs, 95% confidence intervals are constructed around the average vessel cost and the departure delay.

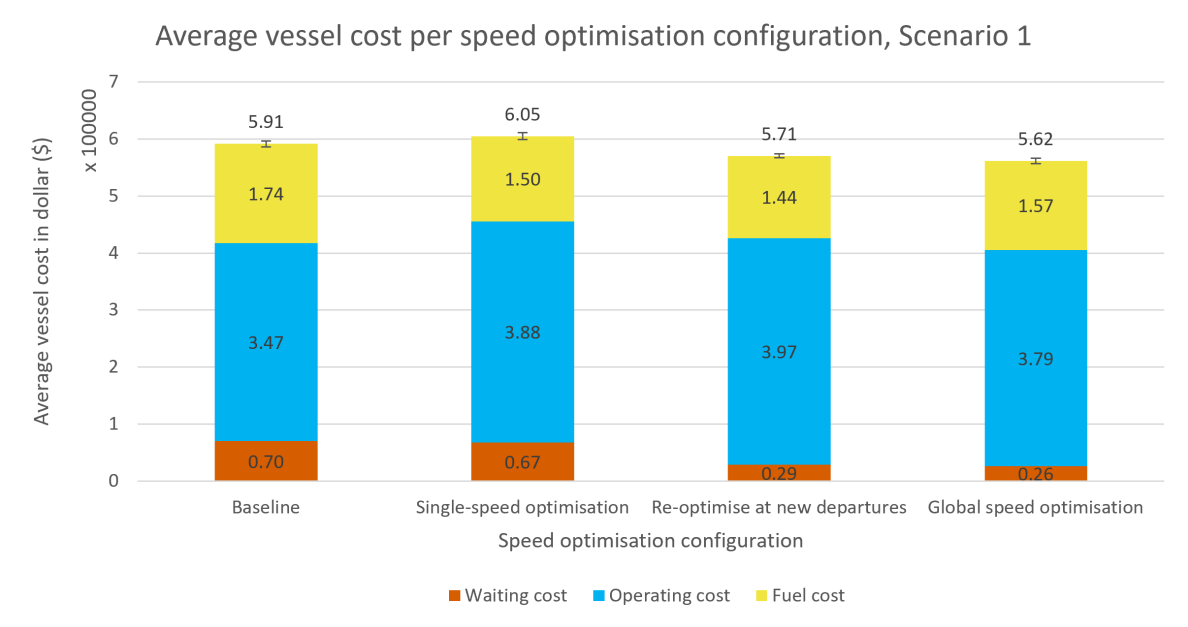


Figure 7.1: Average vessel cost per configuration

	Baseline	Single speed optimisation	Re-optimise at new departures
Single speed optimisation	0.0011	x	x
Re-optimise at new departures	$1.60 \cdot 10^{-6}$	$1.58 \cdot 10^{-8}$	x
Global speed optimisation	$2.51 \cdot 10^{-8}$	$3.20 \cdot 10^{-10}$	0.0037

Table 7.1: Pairwise t-test p-values for the different vessel speed configurations in scenario 1 for the average vessel cost

In Figure 7.1, the average vessel cost per strategy is displayed. The costs are divided into waiting, operating, and fuel costs. Single-speed optimisation has lower fuel costs than the baseline. However, the average total cost is still higher. This is because of higher operating costs from slower sailing, and the waiting cost does not decrease significantly. In re-optimise at new departures, both fuel and the waiting costs decrease compared to the baseline, performing better than both the baseline and the single-speed optimisation. Finally, in global speed optimisation, the overall costs are lowest. Although there is an increase in fuel costs compared to re-optimise at new departures, the decrease in operating costs significantly lowers the overall cost. The increase in fuel cost is caused by vessels occasionally increasing their velocity to improve the overall system's performance. Based on total costs, the global speed optimisation performs better. However, based on fuel costs, re-optimise at new departures performs better. In total, the global speed optimisation configuration in scenario 1, on average, saves $29616 per vessel when compared to the average of the baseline strategy.

The t-test p-values for average costs of different vessel speed optimisation configurations in scenario 1, displayed in Table 7.1, show that the costs are significantly different from one another.

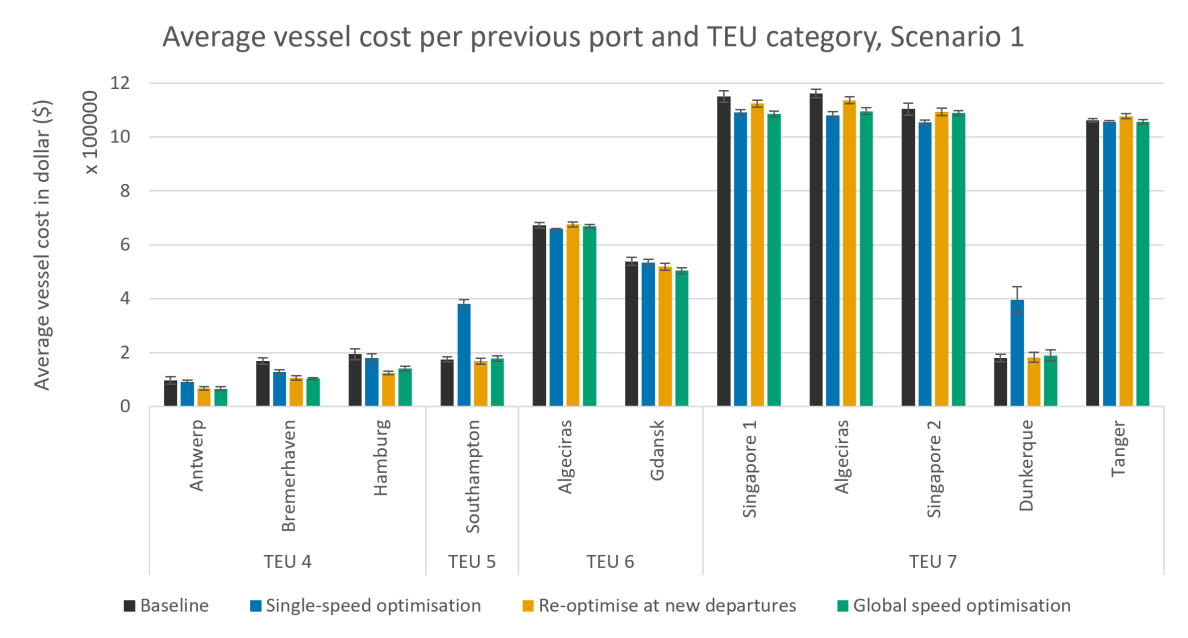


Figure 7.2: Average vessel cost per port and TEU category per configuration

In Figure 7.2, the average vessel cost per port and TEU category per strategy is displayed. The figure shows why single-speed optimisation, on average, does not perform well. As both the TEU 5 vessel departing from Southampton and the TEU 7 vessel departing from Dunkerque have significantly higher costs than the baseline and the other strategies. This is because these previous ports are close to RWG, and when these vessels depart, there is often waiting involved at every velocity at which they can depart. Although the vessels departing from Antwerp, Bremerhaven, and Hamburg also do not have much distance to travel to Rotterdam, this is not the case for them, as they are smaller and more often have berthing space available. The other thing to note is that global speed optimisation performs better for larger vessels (TEU 6 and TEU 7) while Re-optimise at new departures performs better for smaller vessels (TEU 4 and TEU 5).

In Figure 7.3, the average departure delay for all four strategies is displayed. Single-speed optimisation performs worse than the baseline, resulting in a higher average departure delay. The average departure delay for re-optimise at new departures is slightly lower than the average departure delay at the baseline. The global speed optimisation clearly performs best in terms of average departure delay.

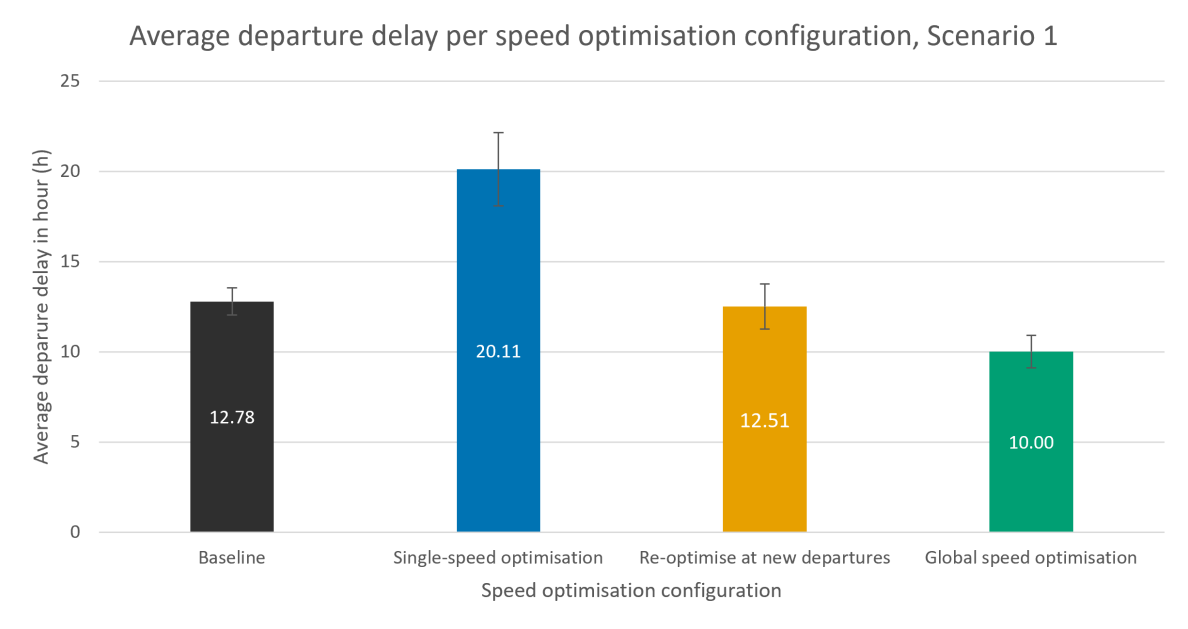


Figure 7.3: Average departure delay per configuration

	Baseline	Single speed optimisation	Re-optimise at new departures
Single speed optimisation	$7.96 \cdot 10^{-6}$	x	x
Re-optimise at new departures	0.6738	$3.06 \cdot 10^{-6}$	x
Global speed optimisation	$4.65 \cdot 10^{-5}$	$1.97 \cdot 10^{-7}$	0.0019

Table 7.2: Pairwise t-test p-values for the different vessel speed configurations in scenario 1 for the average departure delay

The t-test p-values for the average departure delay of different vessel speed optimisation configurations in scenario 1, displayed in Table 7.2, show that re-optimise at new departures has a significant overlap with the baseline configuration. For all the other combinations of strategies, $p < 0.05$, so no clear statistical similarity is visible.

In Figure 7.4, the average departure delay per previous port of the vessel is displayed per strategy. From the graph, it becomes clear why single-speed optimisation performs the worst in Figure 7.3, as the TEU 5 vessel from Southampton, and the TEU 7 vessel from Dunkerque show significant peaks in the departure delay. This also indicates why their costs in Figure 7.2 perform as they do. This graph shows that re-optimise at new departures for larger vessels (TEU 6 and TEU 7) performs worse on departure delay than the baseline, while global speed optimisation performs similarly to the baseline.

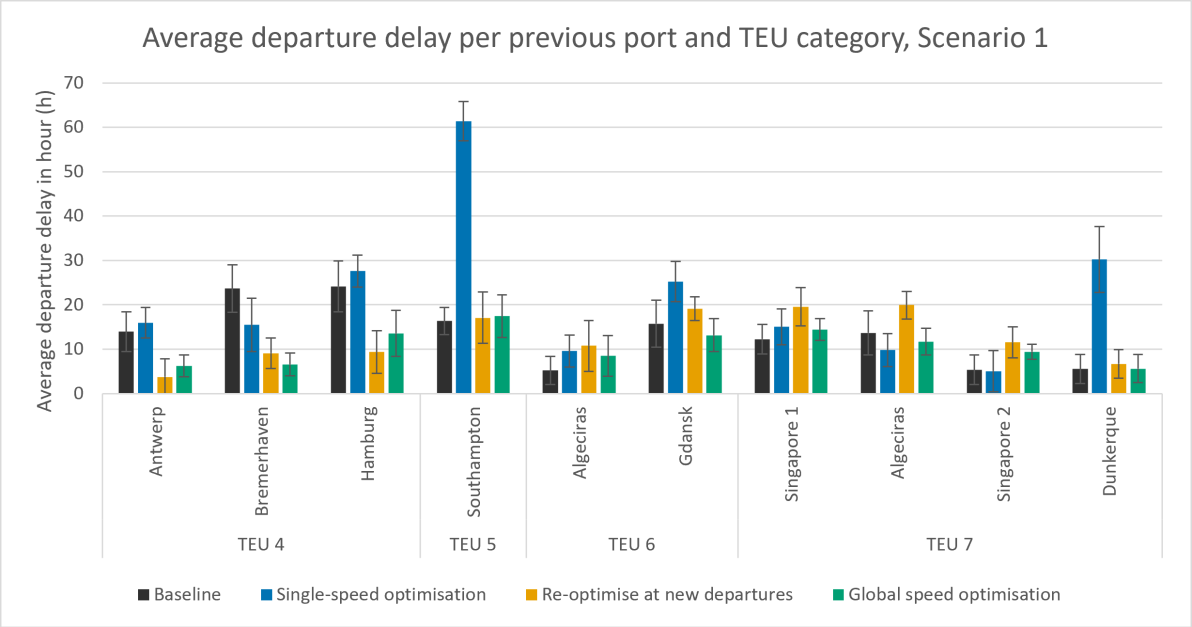


Figure 7.4: Average departure delay per port and TEU category per configuration

Based on these results, a conclusion can be drawn for the sub-question 5, ”How do different vessel speed optimisation strategies affect the performance of container vessels?” The effect that different vessel speed optimisation strategies have on performance depends on the strategy. The single-speed optimisation strategy performs worse than the baseline, both in total costs and departure delay. The re-optimise at new departures performs better than the baseline, both in average cost per vessel and in departure delay. However, the strategy does not outperform the baseline strategy across all previous ports of the vessel and favours smaller vessels. Global speed optimisation outperforms both the baseline and the other strategies, resulting in lower vessel costs for all previous ports and lower departure delays for most of them compared to the baseline. Resulting in an average cost save of $29616 per vessel compared to the baseline in scenario 1.

The half-width of the 95% confidence intervals in Figure 7.1 is relatively small compared to the total vessel costs, ranging from $6000 for the single-speed optimisation configuration, to $3300 for the re-optimise at new departures. However, these confidence intervals are relatively large compared to the differences in average costs between the configurations. However, the t-test p-values indicate that the difference in costs across all configurations is statistically significant. In Figure 7.2, the half-width of the 95% confidence intervals generally become larger per previous port, going up to $49000 dollar for the TEU 7 Dunkerque vessel, in the single-speed optimisation configuration. A similar trend can be seen in Figure 7.3 and Figure 7.4, as the half-width of the 95% confidence intervals of the departure delay for individual vessels is also larger than that of the average departure delay over the entire simulation. This can be explained by the stochastic variables in the model of the arrival time. This indicates that the performance of individual vessels is highly sensitive to the stochastic variations in arrival times. In contrast, the performance of the overall system is more stable, as stochastic effects average out over multiple vessels.

7.2. Influence of scenarios

Based on the results of section 7.1, the global speed optimisation strategy is selected to test the influence of additional vessels on the system. In this section, the results of the experiments devised in section 6.4 are discussed. 4 sets of 10 simulations were conducted: baseline configuration, scenario 1, global speed optimisation, scenario 1, global speed optimisation, scenario 2, global speed optimisation, scenario 3. The goal of the simulations is to find the influence of the speed optimisation strategy on the berth occupancy. The total number of containers handled increases by 44000 TEU on a yearly basis in scenario 2, and by 130000 TEU in scenario 3, compared to scenario

1. The berth occupancy becomes higher in scenario 2, as one additional vessel is added to the biweekly schedule, and in scenario 3, as two additional vessels are added to the biweekly schedule. The results are analysed in terms of vessel costs, departure delay and berth occupancy.

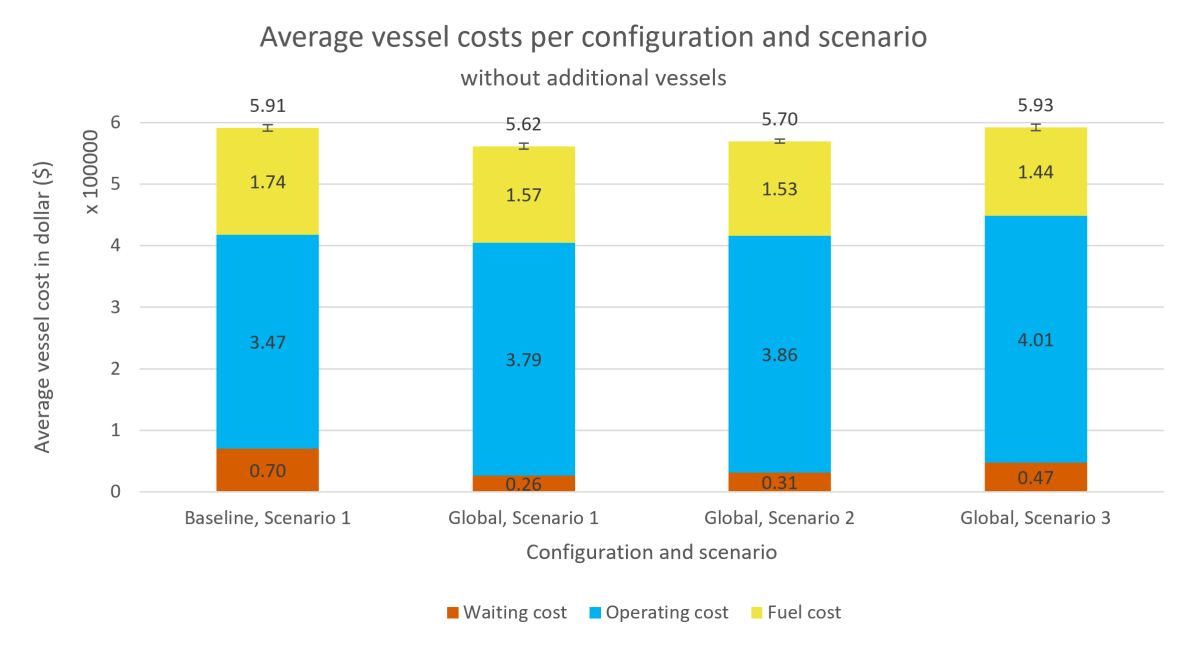


Figure 7.5: Average vessel cost per configuration and scenario without additional vessels

	Baseline, Scenario 1	Global speed optimisation, Scenario 1	Global speed optimisation, Scenario 2
Global speed optimisation, Scenario 1	$2.51 \cdot 10^{-8}$	x	x
Global speed optimisation, Scenario 2	$1.07 \cdot 10^{-6}$	0.0065	x
Global speed optimisation, Scenario 3	0.7286	$1.91 \cdot 10^{-8}$	$8.25 \cdot 10^{-7}$

Table 7.3: Pairwise t-test p-values for the different vessel speed configurations and scenarios for the average vessel costs

In Figure 7.5, the average vessel cost per configuration and scenario is displayed. The average vessel cost excludes the additional vessels (TEU 4 from Le Havre and TEU 5 from Antwerp), as they are not included in scenario 1 and would skew the average cost. The costs are thus compared to the cost of the original set of vessels, to ensure comparability with the other scenarios.

By comparing the three different scenarios, on the global speed optimisation configuration in Figure 7.5, an increase in waiting cost and a decrease in fuel cost are shown. This indicates that vessels adapt by sailing more slowly to avoid congestion, thereby reducing fuel costs and increasing operating costs. When comparing the global speed optimisation configuration, scenario 2, to the baseline, scenario 1, it performs at a lower average cost than the baseline, whereas the global speed optimisation configuration, scenario 3, results in a higher average cost, indicating worse performance.

The t-test p-values for the average vessel costs per configuration and scenario displayed in Table 7.3 show that the baseline scenario 1 has a similar average cost to the global speed optimisation scenario 3. The costs of other scenarios have statistically significant differences.

In Figure 7.6, the average vessel costs per previous port are displayed. It is evident that the average vessel costs for the largest vessel classes (TEU 6 and TEU 7) in the global speed optimisation configuration in scenario 3, exceed those of the baseline strategy, while the smaller vessel classes remain below it.

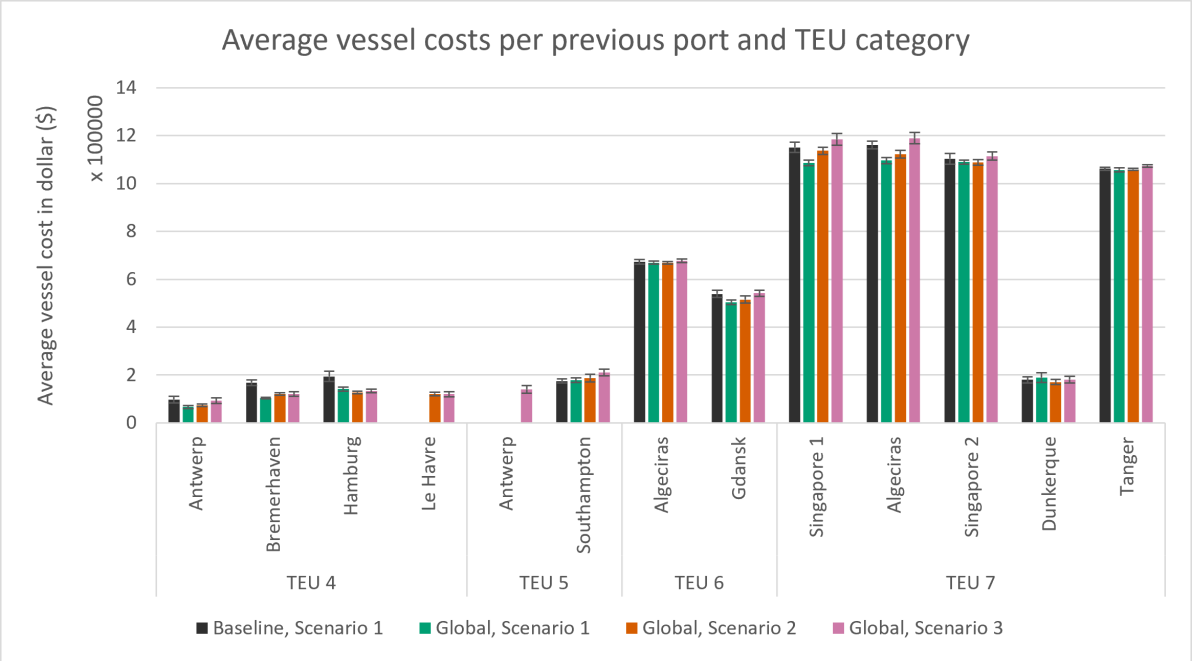


Figure 7.6: Average vessel cost per port and TEU category per configuration and scenario

By comparing the three different scenarios, on the global speed optimisation configuration in Figure 7.7, an increase in average departure delay is shown for scenario 2 and scenario 3. However, the global speed optimisation configuration in scenario 2 still performs better than the baseline configuration in scenario 1. In contrast, the global speed optimisation configuration, under scenario 3, performs worse than the baseline in terms of departure delay.

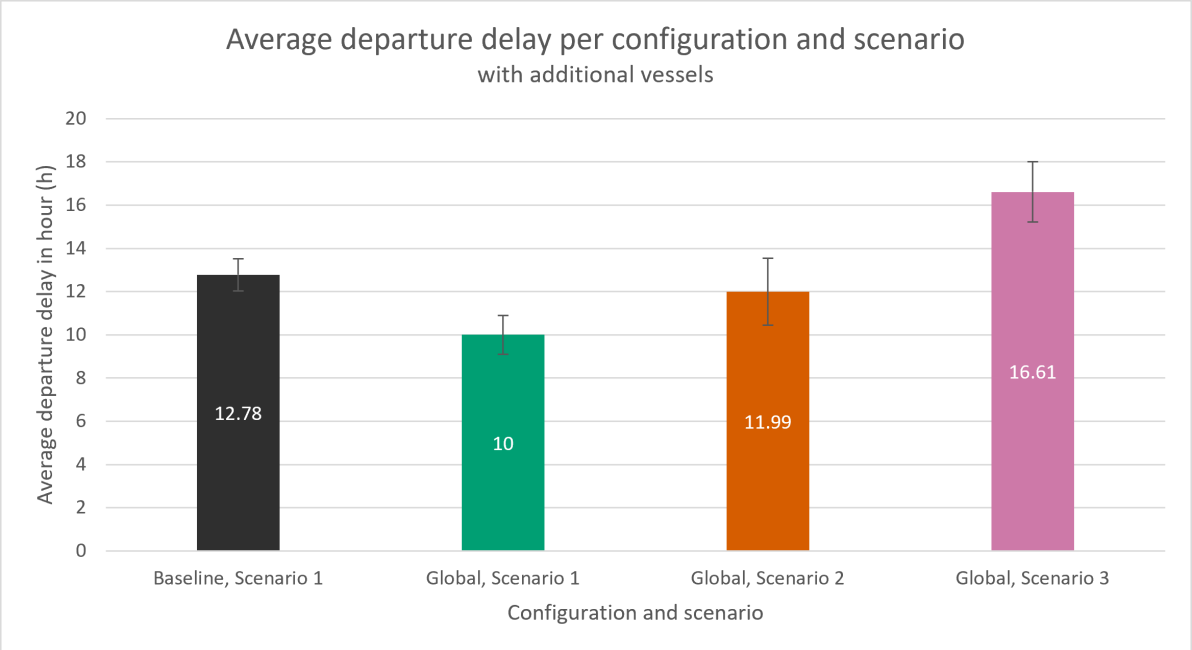


Figure 7.7: Average departure delay per configuration and scenario with additional vessels

	Baseline, Scenario 1	Global speed optimisation, Scenario 1	Global speed optimisation, Scenario 2
Global speed optimisation, Scenario 1	$4.65 \cdot 10^{-5}$	x	x
Global speed optimisation, Scenario 2	0.2225	0.0387	x
Global speed optimisation, Scenario 3	$6.84 \cdot 10^{-5}$	$8.31 \cdot 10^{-8}$	$6.95 \cdot 10^{-5}$

Table 7.4: Pairwise t-test p-values for the different vessel speed configurations and scenarios for the average departure delay

In Figure 7.8, the average departure delay per port and TEU category is displayed. These results demonstrate that in the global speed optimisation scenario 2 and scenario 3, the departure delays of some of the larger vessels increase compared to the baseline scenario 1, while the departure delay of the smaller vessels remains at or below the baseline scenario 1.

The t-test p-values for the average departure delay per configuration and scenario are displayed in Table 7.4. The average departure delay in baseline scenario 1 is not significantly different from global speed optimisation scenario 2 (p = 0.2225), whereas all other pairwise comparisons show statistically significant differences.

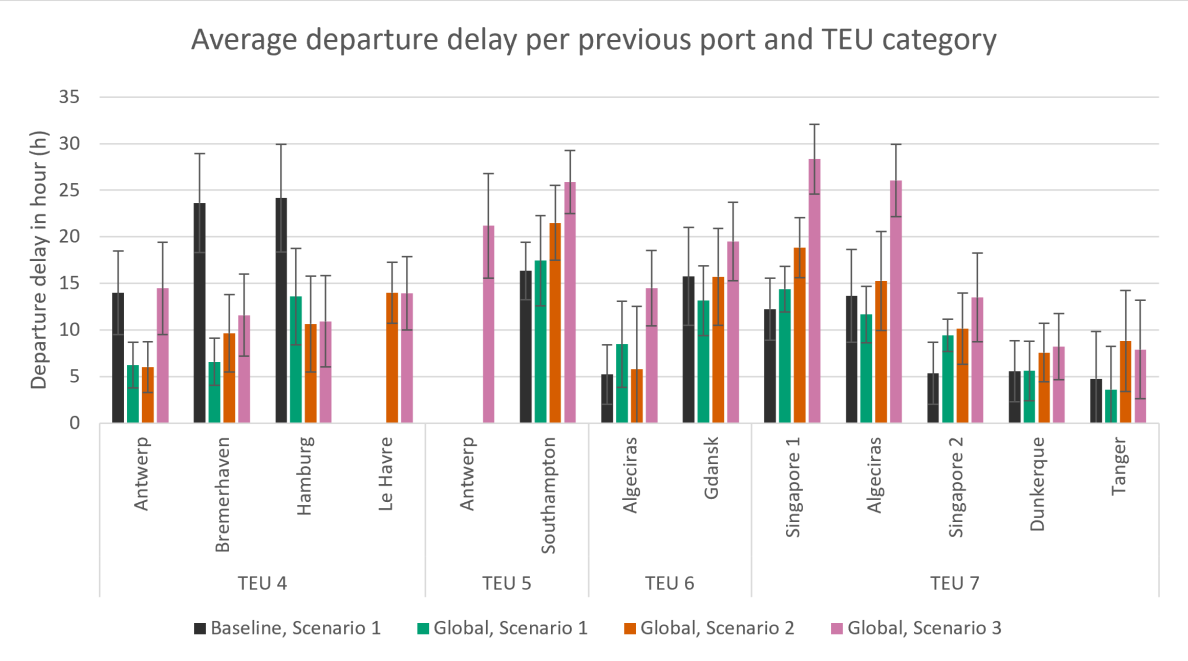


Figure 7.8: Average departure delay per port and TEU category per configuration and scenario

In Figure 7.9, the berth occupancy of all four scenarios and configurations is displayed. The berth occupancy of scenario 2 increases by 2.2%, and the berth occupancy of scenario 3 increases by another 3.5%, as a direct result of the extra vessel arrivals. In Appendix section C.4, the reason for the confidence intervals in the figure is explained, and the increase in berth occupancy is calculated by hand, as additional verification.

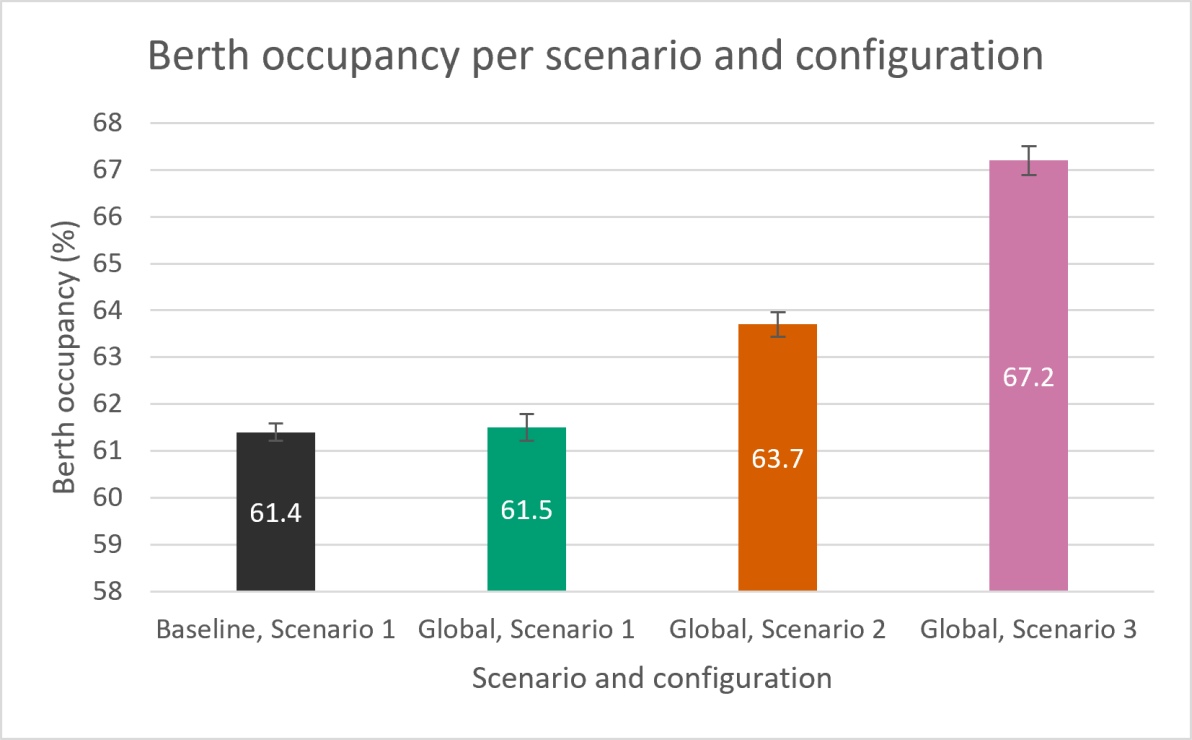


Figure 7.9: Berth occupancy per configuration and scenario

Based on these results, a conclusion can be drawn for the sub-question 6, ”How can vessel speed optimisation influence berth occupancy utilisation of a container terminal?” As Figure 7.9 indicates, adding additional vessels in the case of scenario 2 and scenario 3 increases the terminal’s berth occupancy. However, increasing the number of vessels reduces the performance of both the terminal and the vessels, as average vessel costs and average vessel departure delays increase. In the global speed optimisation scenario 2, the average vessel cost and average departure delay remain lower than in the baseline. However, for individual vessels, the departure delay can increase. Therefore, concluding that vessel speed optimisation can increase the terminal’s overall berth occupancy; however, the effect is limited, as the terminal’s service level decreases when too many vessels are added. Within the tested scenarios, scenario 2 represents the highest increase in vessel arrivals that can be accommodated without deteriorating average vessel performance below the baseline. This indicates that the terminal is approaching the capacity limit, beyond which additional vessels lead to disproportionate increases in delay and costs. Scenario 2 would lead to the terminal handling an additional 44.000 TEU per year, when compared to scenario 1. The confidence intervals show a similar pattern to that in section 7.1, indicating that while system-level trends are robust, variability at the individual vessel level across simulations remains high.

7.3. Sensitivity analysis

This section evaluates the model’s sensitivity to key input parameters. Three sets of sensitivity tests are compared to the base input used in all previous simulations. All scenarios are run using the global speed optimisation configuration, scenario 1. Using these results, the sensitivity of several parameters can be determined from the departure delays of the vessels. In sensitivity tests 1 and 2, the fuel consumption parameter $\tilde{f}$ is changed. In test 1, the difference in $\tilde{f}$ between vessel classes is relatively larger, while in test 2, the difference is relatively smaller. In sensitivity test 3, the waiting costs for the TEU 6 and TEU 7 vessels are set to 1.5 times their operating costs, instead of the usual 1.1. The parameters used for the sensitivity tests are shown in Table 6.4.

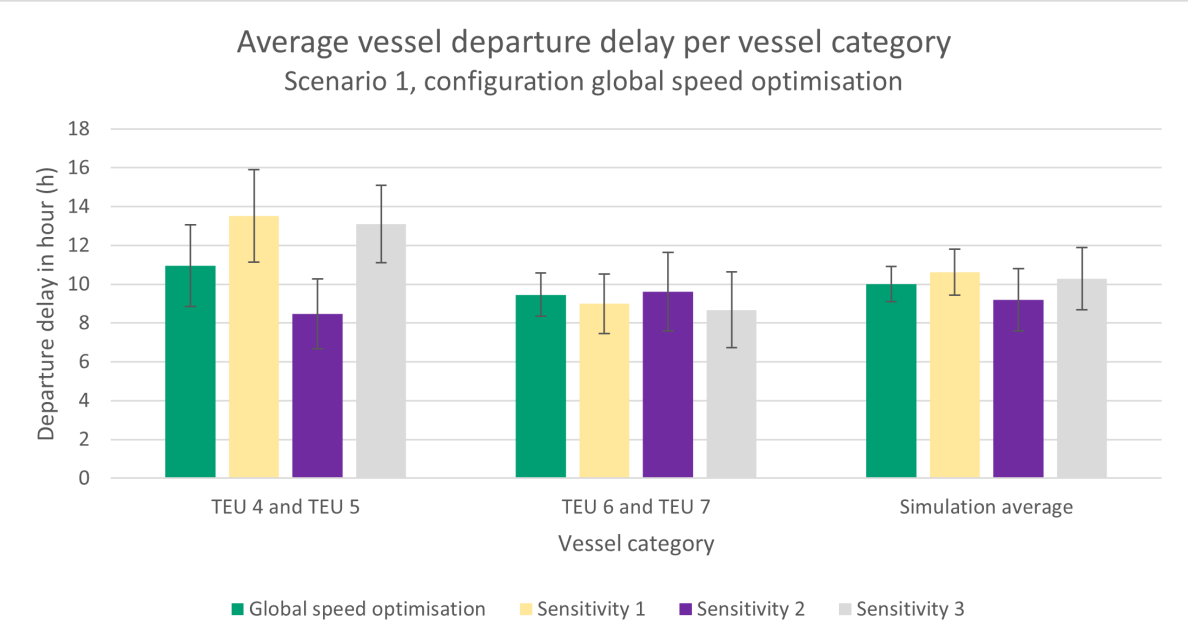


Figure 7.10: Sensitivity per class

The results are portrayed in Figure 7.10. The figure shows that increasing the relative cost difference between smaller (TEU 4 and TEU 5) and larger vessels (TEU 6 and TEU 7) results in higher departure delays for the smaller vessels and lower departure delays for the larger vessels. In total, the average departure delay over the entire simulation increases compared to the standard global speed optimisation scenario. Decreasing the relative costs yields the opposite results: a lower departure delay for the smaller vessels and a slight increase for the larger vessels, leading to a lower simulation-average departure delay. Increasing the waiting cost multiplication factor for the larger vessels in sensitivity 3 yields results similar to sensitivity 1, causing a larger departure delay than the global speed optimisation strategy for the smaller vessels and a lower departure delay for the larger vessels, resulting in a slightly higher simulation average.

These sensitivity tests show that the relative size of the parameter $\tilde{f}$ influences the departure delay of vessels. With a larger difference in $\tilde{f}$ between vessels, the model prefers to decrease the delay of larger vessels; with a smaller difference in $\tilde{f}$, it prefers to decrease the delay of smaller vessels. Increasing the waiting cost multiplier also results in a preference for larger vessels. This behaviour indicates that the model prioritises vessels with higher cost impacts, shifting delays towards the smaller, less expensive vessels.

However, the confidence intervals in Figure 7.10 show significant overlap between the sensitivity tests and the overall system average, indicating that the observed differences may not be statistically significant to the actual difference in changing the parameters. The change in smaller vessels, however, is much more significant. This is also visible in Appendix Table B.1, Table B.2 and Table B.3, where the t-test p-values of the figure are displayed. When comparing global speed optimisation to other sensitivities, $p > 0.05$, indicating no statistical differences in departure delay. However, sensitivity 1 and 2, and sensitivity 2 and 3, show statistically significant differences in the departure delay of the TEU 4 and TEU 5 vessels.

7.4. Conclusion

In this chapter, the simulation model results are discussed. In section 7.1, a conclusion is drawn to sub-question 5: ”How do different vessel speed optimisation strategies affect the performance of container vessels?” The conclusion is that the vessel speed optimisation strategies improve the vessel performance, and can reduce the total vessel costs up to $29616 per vessel. The strategies global speed optimisation and re-optimise at new departures both have significantly lower average

vessel costs than the baseline. Pairwise t-tests confirm that the differences in average vessel costs between all four configurations are statistically significant. The global speed optimisation strategy also results in lower departure delays. Compared to the re-optimise at new departures, the global speed optimisation strategy has slightly higher fuel costs, due to the possibility of some vessels sailing faster, to lower the costs for other vessels. An observation on the results is that the simulation's confidence intervals are relatively small, while those per previous port of the vessel are generally larger. This indicates that the behaviour of a single vessel can vary with the realisation of stochastic parameters, whereas the results of an entire simulation are generally more stable.

In section 7.2, a conclusion is drawn to sub-question 6: "How can vessel speed optimisation influence berth occupancy of a container terminal?" The conclusion is that with the global vessel speed optimisation strategy, the berth occupancy can be increased by 2.2% in scenario 2 with the global speed optimisation configuration, while maintaining significantly lower average vessel costs and statistically similar average departure delays when compared to the baseline configuration in scenario 1. However, for some specific vessels, departure delays and vessel costs will increase relative to the baseline. In scenario 2, the container terminal would handle an additional 44.000 TEU. Increasing berth occupancy further is not recommended, as it leads to congestion that deteriorates the system's overall performance.

In section 7.3, the results from the sensitivity tests are presented. The sensitivity tests indicate that the model prioritises vessels with higher cost impacts, shifting delays towards the smaller, less expensive vessels. By making it more expensive for the larger vessels to wait, either by increasing $\tilde{f}$ (sensitivity 1) or by increasing the waiting cost multiplication factor (sensitivity 2), the departure delay of the larger vessels decreases. However, an observation on the results is that the confidence intervals of the different sensitivity tests largely overlap in the TEU 6 and TEU 7 vessel category and the simulation average category.

Overall, the results in this chapter demonstrate that vessel speed optimisation can positively influence the performance of a container terminal and the container vessels. However, under a larger number of scenarios with a larger number of incoming vessels, the benefits diminish.

8

Conclusion and Discussion

In this chapter, the overall conclusion of this thesis is drawn in section 8.1. In the conclusion, first, the conclusions to the sub-questions are drawn, and finally, a conclusion is drawn on the main research question: ”What is the effect of dynamic vessel speed optimisation on the performance of container vessels and berth occupancy of a deep-sea container terminal?” In section 8.2, a few discussion points on the model are introduced, and in section 8.3, recommendations for future research and more accurate data collection are proposed.

8.1. Conclusion

In this section, the conclusion to the overall thesis is drawn. This is done by answering the main research question: ”What is the effect of dynamic vessel speed optimisation on the performance of container vessels and berth occupancy of a deep-sea container terminal?” In order to answer the main research question, the answers to the sub-questions will first be summarised.

8.1.1. Answer to sub-question 1

Sub-question 1: *How do current vessel types, vessel operations, and seaside operations influence waiting times at the terminal?*

The sub-question is answered in chapter 2 by conducting a system analysis into three separate topics of the sub-question. In the chapter, the interactions between the terminal and vessel operations are discussed. The analysis shows that the waiting time at the terminal is influenced by both the terminal’s sea-side operations and the vessel’s operations. A combined model of vessel and terminal operations is proposed that can analyse waiting times.

8.1.2. Answer to sub-question 2

Sub-question 2: *How can vessel speed optimisation be incorporated into the berth allocation problem?*

The sub-question is answered in chapter 3 by conducting a literature review of the berth allocation problem and vessel speed optimisation. The rolling-horizon strategy is selected to address uncertainty about the arrival of incoming vessels. In the chapter, different integrations of the berth allocation problem with vessel speed optimisation are presented. Most of these integrations rely on minimising one objective. The common objective is to minimise the overall cost. The costs consist of vessel speed-optimisation objectives, such as minimising fuel consumption, and terminal objectives, such as minimising waiting time.

8.1.3. Answer to sub-question 3

Sub-question 3: *Which vessel speed optimisation strategies can be identified and structured for use in simulation?*

The sub-question is answered in chapter 4 by first determining which speed optimisation strategies have been used in literature, and then changing those to be used for the simulation. In total four strategies have been identified: (1) the baseline strategy, in which vessels depart at a fixed velocity, (2) the single speed optimisation strategy, at which the terminal determines the velocity of the vessel at the moment of departure, (3) re-optimise at every new departure, at which at every new departure, the terminal will determine new velocities for all sailing vessels with a larger remaining distance, (4) global speed optimisation, at which the terminal will decide the best velocity for all sailing vessels at every planning interval, sometimes sacrificing the performance of single vessel for the gain of others.

8.1.4. Answer to sub-question 4

Sub-question 4: *How can vessel speed optimisation be implemented in a discrete event simulation model that realistically represents real-life vessel operations?*

The sub-question is answered in chapter 5. In the chapter, a simulation model with inputs and outputs is proposed for both the seaside operations of a container terminal and the vessel speed optimisation. These models are combined. The decision logic is verified based on visual checks and output reports. The model's input is validated using AIS data on vessels arriving at RWG in 2024. The verified and validated model implements vessel speed optimisation in a realistic scenario, thus answering the research question.

8.1.5. Answer to sub-question 5

Sub-question 5: *How do different vessel speed optimisation strategies affect the performance of container vessels?*

In section 7.1, the results from the simulation of the baseline configuration and the 3 different speed optimisation configurations under scenario 1 are shown. The conclusion to the research question is that the vessel speed optimisation strategies improve the vessel performance, and can reduce the total vessel costs of approximately $30000 per vessel compared to the baseline strategy. However, not all speed optimisation strategies perform better than the baseline. The single-speed optimisation strategy performs worse in both average vessel costs and average vessel departure delay. The strategies global speed optimisation and re-optimise at new departures both have significantly lower average vessel costs than the baseline. The global speed optimisation strategy also results in lower departure delays. Compared to re-optimise at new departures, the global speed optimisation strategy has slightly higher fuel costs, due to the possibility of some vessels sailing faster, to lower the costs for other vessels.

8.1.6. Answer to sub-question 6

Sub-question 6: *How can vessel speed optimisation influence the berth occupancy of a container terminal?*

In section 7.2, the results from baseline scenario 1, global speed optimisation scenario 1, global speed optimisation scenario 2, and global speed optimisation scenario 3 are discussed. In scenario 2, the berth occupancy increases by 2.2%, and in scenario 3, the berth occupancy increases by 5.7% compared to the baseline. However, in scenario 3, the vessels perform worse on departure delays and average vessel costs than in the baseline of scenario 1. However, the average vessel cost in the configuration global speed optimisation with scenario 2 is lower than that in the configuration baseline with scenario 1, while the departure delay remains similar. Although for some specific vessels, the departure delay and vessel costs will increase compared to the baseline. Concluding that, with the global vessel speed optimisation strategy, berth occupancy can be increased by 2.2% without decreasing the average vessel performance. In scenario 2, the container terminal would handle an additional 44.000 TEU. Increasing berth occupancy further is not recommended, as it leads to congestion that deteriorates the overall system's performance.

8.1.7. Answer to the main research question

Main research question: *What is the effect of dynamic vessel speed optimisation on the performance of container vessels and berth occupancy of a deep-sea container terminal?*

The model used to answer the main research question has been constructed in chapter 5 using the knowledge gained from answering sub-questions 1, 2, and 3. The main research question consists of two parts, both of which are addressed separately in sub-questions 5 and 6. Dynamic vessel speed optimisation has a positive effect on the performance of vessels, leading to cost savings of approximately $30000 on average per vessel, in total costs, while using the global speed optimisation strategy. The berth occupancy can also be affected by vessel speed optimisation. The berth occupancy can be increased from 61.5% to 63.7%, without deteriorating system performance compared to the baseline scenario. Concluding that the global speed optimisation strategy can improve the performance of container vessels while increasing the berth occupancy of a deep-sea container terminal.

8.2. Discussion

The thesis's conclusion relies on several simplifications and assumptions. Some of those simplifications or assumptions can affect the research if they are unrealistic.

The main assumptions concern vessels departing at a fixed velocity. In the model, it is assumed that, in the baseline configuration, vessels depart from their previous port at a fixed velocity of 16 NM/h. However, this assumes that there is no contact between vessel operators and terminal operators, who may already influence the velocity of incoming vessels.

Besides the fixed velocity, the willingness of incoming vessels to change their velocity can also be up for debate. In one of the results, a vessel had to change its velocity up to 23 times. These velocity changes were mainly between 15 NM/h and 16 NM/h, and back, depending on the best time of arrival at the calculation point. An image displaying the velocity changes of a vessel is found in Appendix Figure B.5. Although 23 velocity changes are an excess, if more vessels are added to the biweekly schedule, or if the maximum distance is further than 1200 NM away from the focal port, this number may even increase. The behaviour of changing velocity 23 or more times during a voyage may not be feasible in practice.

Besides the velocity of the incoming vessel, the waiting time of the vessels is also a point of discussion. In section 5.4, the input of scenario 1 is validated. However, the scenario with short-sea vessels and barges has slightly longer waiting times for the deep-sea vessels than the scenario without short-sea vessels and barges (3447 hours vs 3189 hours over the entire year-long simulation), indicating that they do affect the system. The model without short-sea vessels and barges, however, provides proof of concept: when simulating only the deep-sea vessels, the vessel speed optimisation strategies affect the overall deep-sea vessel's performance. However, a simulation with short-sea vessels and barges can provide a more accurate assessment of overall effectiveness.

The exogenic input of vessel arrivals is also simplified into a schedule that repeats every 2 weeks for an entire year. However, from the AIS data, vessel arrivals were not as consistent, especially at the vessels' previous ports. For instance, 16 out of the 19 vessels that departed from Tanger arrived before 01-06-2024, while 28 out of the 33 that departed from Dunkerque arrived after 01-06-2024. These changes can be caused by the liner taking a different route through Europe. However, the remaining distance to the new port heavily influences the time of arrival of incoming vessels, and thus influences the reliability of the system.

8.3. Recommendations

8.3.1. Recommendations for future research

The first recommendation ties into one of the discussion points. If it is possible to accurately predict when vessels depart their previous port, they can be scheduled into berth planning at the new port much earlier. With more knowledge of the vessel departing from nearby, vessels from further out no longer have to change their velocity as harshly or as often.

Another recommendation is to research the influence of speed optimisation for deep-sea vessels on all arriving vessels at the container terminal. In this case, short-sea vessels and barges must be included in the simulation. This could be done in two ways: either by letting the vessels arrive directly

at the terminal without accounting for their sailing leg, or by simulating their sailing leg and optimising their velocity. However, by also optimising the velocity of the short-sea vessels, another 50 vessels must be simulated every 2 weeks, and by optimising the barges, another 154 vessels must be simulated. The model's current solution mechanism is not suited to this, so heuristics must be implemented to solve the problem.

In this research, every deep-sea vessel has the same $\tilde{v}$ at 16 NM/h. This velocity has been chosen so that the fuel consumption and costs of each vessel within the same class remain the same. However, the vessels within a class do not always set sail at the same velocity, so $\tilde{v}$ can be made a variable to range from 14 NM/h to 18 NM/h, with the maximum velocity being $\tilde{v}$ + 6 NM/h and the minimum velocity being $\tilde{v}$ - 6 NM/h. By varying $\tilde{v}$ across vessel classes and individual vessels, the model would yield more realistic results.

8.3.2. Recommendations for data collection

To make the model more accurate, more realistic data is needed.

The first data set must be the vessel's average velocity over the sailing leg. Using the data, the $\tilde{v}$ of a vessel can be determined. Also, the deviations from $\tilde{v}$, to avoid waiting at the quay, can already be determined.

The other main data set is an estimate of the vessel's call size. In this research, the call size was reverse-engineered based on the average time a vessel spent on the quay relative to the average productivity of vessels of the same size. However, using more accurate data can eliminate the estimates made during reverse engineering. Besides the call size at the current port, an estimate of the call size at the previous port is also interesting to know. Using this value, an estimate can be made of when a vessel will depart the previous terminal, and the vessel's arrival can be planned as soon as it is being serviced there, rather than in this model, after it has departed.

In this research, the ETA deviation was based on TRAFALQUAR's basic values for each vessel class. In the AIS data, no ETA was given, only the arrival at the anchorage point and the arrival at the terminal. If the realised deviation from the ETA was given in the AIS data, TRAFALQUAR's ETA deviation can be more realistic.

References

[1] J. Alvarez, Tore Longva, and Erna Engebrethsen. “A methodology to assess vessel berthing and speed optimization policies”. In: *Maritime Economics and Logistics* 12 (Dec. 2010), pp. 327–346. DOI: 10.1057/mel.2010.11.

[2] Henrik Andersson, Kjetil Fagerholt, and Kirsti Hobbesland. “Integrated maritime fleet deployment and speed optimization: Case study from RoRo shipping”. In: *Computers and Operations Research* 55 (2015), pp. 233–240. DOI: 10.1016/j.cor.2014.03.017.

[3] Peter Andersson and Pernilla Ivehammar. “Green approaches at sea – The benefits of adjusting speed instead of anchoring”. In: *Transportation Research Part D: Transport and Environment* 51 (2017), pp. 240–249. DOI: 10.1016/j.trd.2017.01.010.

[4] Christian Bierwirth and Frank Meisel. “A follow-up survey of berth allocation and quay crane scheduling problems in container terminals”. In: *European Journal of Operational Research* 244.3 (2015), pp. 675–689. DOI: 10.1016/j.ejor.2014.12.030.

[5] Berit D. Brouer et al. “A base integer programming model and benchmark suite for liner-shipping network design”. In: *Transportation Science* 48.2 (2014), pp. 281–312. DOI: 10.1287/trsc.2013.0471.

[6] Héctor J. Carlo, Iris F. A. Vis, and Kees Jan Roodbergen. “Seaside operations in container terminals: literature overview, trends, and research directions”. In: *Flexible Services and Manufacturing Journal* 27.2-3 (June 2013), pp. 224–262. DOI: 10.1007/s10696-013-9178-3.

[7] *Diensten | Rotterdam World Gateway*. URL: https://www.rwg.nl/nl/diensten.

[8] Yuquan Du et al. “Berth allocation considering fuel consumption and vessel emissions”. In: *Transportation Research Part E: Logistics and Transportation Review* 47.6 (2011), pp. 1021–1037. DOI: 10.1016/j.tre.2011.05.011.

[9] Mihalis M. Golias, Hercules E. Haralambides, and Michiel Nienhuis. “Berth scheduling by customer service differentiation: A multi-objective approach”. In: *Transportation Research Part E: Logistics and Transportation Review* 45.6 (2009), pp. 878–892. DOI: 10.1016/j.tre.2009.05.006.

[10] Mihalis M. Golias et al. “The berth allocation problem: Optimizing vessel arrival time”. In: *Maritime Economics and Logistics* 11.4 (2009), pp. 358–377. DOI: 10.1057/mel.2009.12.

[11] Xiao-le Han, Zhi-qiang Lu, and Li-feng Xi. “A proactive approach for simultaneous berth and quay crane scheduling problem with stochastic arrival and handling time”. In: *European Journal of Operational Research* 207.3 (2010), pp. 1327–1340. DOI: https://doi.org/10.1016/j.ejor.2010.07.018. URL: https://www.sciencedirect.com/science/article/pii/S037722171000528X.

[12] Sönke Hartmann, Jennifer Pohlmann, and Axel Schoenknecht. “Simulation of Container Ship Arrivals and Quay Occupation”. In: vol. 49. Feb. 2011, pp. 135–154. ISBN: 9781441984074. DOI: 10.1007/978-1-4419-8408-1_8.

[13] International Maritime Organization. *Fourth Greenhouse Gas Study 2020*. 2020. URL: https://www.imo.org/en/ourwork/environment/pages/fourth-imo-greenhouse-gas-study-2020.aspx.

[14] International Maritime Organization. *Just In Time Arrival Guide*. 2020. URL: https://wwwcdn.imo.org/localresources/en/OurWork/PartnershipsProjects/Documents/GIA-just-in-time-hires.pdf.

[15] Bin Ji et al. “A modified adaptive large neighborhood search algorithm for solving the multi-port continuous berth allocation problem with vessel speed optimization”. In: *Computers & Industrial Engineering* 198 (2024), p. 110699. ISSN: 0360-8352. DOI: 10.1016/j.cie.2024.110699.

[16] Bin Ji et al. “Hybrid rolling-horizon optimization for berth allocation and quay crane assignment with unscheduled vessels”. In: *Advanced Engineering Informatics* 54 (2022). DOI: 10.1016/j.aei.2022.101733.

[17] Niels Lang and Albert Veenstra. “A quantitative analysis of container vessel arrival planning strategies”. In: *OR Spectrum* 32.3 (2010), pp. 477–499. DOI: 10.1007/s00291-009-0186-3.

[18] *Liner Shipping: Fixed Routes for Efficient Global Trade*. URL: https://www.lmitac.com/articles/liner-shipping-global-trade-backbone.

[19] Nextlogic. *Schipper - NextLogic*. Aug. 2023. URL: https://www.nextlogic.nl/zowerkthet/schipper/.

[20] HbR NV. *GENERAL TERMS AND CONDITIONS INCLUDING PORT TARIFFS*. Tech. rep. 2025. URL: https://www.portofrotterdam.com/sites/default/files/2024-12/General%20Terms%20and%20Condition%202025.pdf.

[21] Harilaos N. Psaraftis and Christos A. Kontovas. “Ship speed optimization: Concepts, models and combined speed-routing scenarios”. In: *Transportation Research Part C: Emerging Technologies* 44 (2014), pp. 52–69.

[22] Stewart Robinson. “Simulation model verification and validation: Increasing the users’ confidence”. In: 1997, pp. 53–59.

[23] Filipe Rodrigues and Agostinho Agra. “Berth allocation and quay crane assignment/scheduling problem under uncertainty: A survey”. In: *European Journal of Operational Research* 303.2 (2022), pp. 501–524. DOI: 10.1016/j.ejor.2021.12.040.

[24] Mario Rodríguez-Molins, Miguel Salido, and F. Barber. “Robust Scheduling for Berth Allocation and Quay Crane Assignment Problem”. In: *Mathematical Problems in Engineering* 2014 (Dec. 2014), pp. 1–17. DOI: 10.1155/2014/834927.

[25] *Rotterdam bunker prices*. 2024. URL: https://shipandbunker.com/prices/emea/nwe/nl-rtm-rotterdam.

[26] *SEA-DISTANCES.ORG - distances*. URL: https://sea-distances.org/.

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

[28] Inkyung Sung, Haris Zografakis, and Peter Nielsen. “Multi-lateral ocean voyage optimization for cargo vessels as a decarbonization method”. In: *Transportation Research Part D: Transport and Environment* 110 (2022). DOI: 10.1016/j.trd.2022.103407.

[29] United Nations Conference on Trade and Development. *Review of Maritime Transport 2024*. Technical report. United Nations conference on Trade and Development (UNCTAD), 2024. URL: https://unctad.org/publication/review-maritime-transport-2024.

[30] Pim Willem Antoon van Leeuwen and Rommert Dekker. “Assessing the value of vessel information sharing”. In: *Maritime Transport Research* 9 (2025). DOI: 10.1016/j.martra.2025.100140.

[31] Pim Willem Antoon van Leeuwen and Rommert Dekker. “Vessel arrival planning with full information”. In: *International Transactions in Operational Research* (Sept. 2025). DOI: 10.1111/itor.70108.

[32] Yadong Wang and Shuaian Wang. “Deploying, scheduling, and sequencing heterogeneous vessels in a liner container shipping route”. In: *Transportation Research Part E: Logistics and Transportation Review* 151 (2021). DOI: 10.1016/j.tre.2021.102365.

[33] Jingjing Yu et al. “Berth allocation and quay crane assignment considering the adoption of different green technologies”. In: *Transportation Research Part E: Logistics and Transportation Review* 176 (2023). DOI: 10.1016/j.tre.2023.103185.

A

Research Paper

The Impact of Dynamic Vessel Speed Optimisation on Berth Occupancy at a Container Terminal

J.P.J. ter Meulen, M.B. Duinkerken, R.R. Negenborn, A. de Waal, A.J. Dijkstra

Deep-sea container vessels play a crucial role in the world economy. Despite the enormous operating costs, these vessels often have to wait at the anchorage point due to limited berth availability at container terminals. This paper investigates the influence of vessel speed optimisation on the berth occupancy at the container terminal by developing a discrete-event simulation model. In this model, two configurations are compared. A baseline configuration, in which vessels determine their own optimal speed, and are served at the terminal on a *first-come, first-served* basis, and a global speed optimisation configuration, in which the terminal adjusts the velocities for every sailing vessel at time intervals, and new vessel departures. To evaluate the impact of the global speed optimisation on the berth occupancy, three scenarios are defined. The first scenario represents the operations at Rotterdam World Gateway in 2024 and acts as a baseline. In the second and third scenarios, the number of vessels arriving is increased. The results show that in scenario 2, with global speed optimisation, berth occupancy increases by 2.2%, while maintaining lower average vessel costs and statistically similar average departure delays compared to the baseline scenario. This corresponds to 44.000 additional TEU handled annually. Further increases, however, lead to congestion and a decrease in overall vessel performance. These findings imply that improved coordination between vessels and terminals can enhance the berth occupancy, without compromising service levels.

I. INTRODUCTION

Global maritime shipping plays a crucial role in the world economy [1]. In 2024, 80% of global trade volume was carried by maritime transport. The total volume of maritime transport is expected to increase by over 2% annually [1]. International shipping accounts for approximately 3% of global greenhouse gas emissions [1, 2]. The efficiency of port operations is crucial for maintaining a sustainable form of maritime transport. Several factors contribute to efficient maritime operations. In a container terminal, high quay utilisation leads to more efficient use of available equipment. However, with high quay utilisation, the probability that incoming vessels have to wait increases. Waiting time decreases the port's service level. Through robust berth planning, quay utilisation and service level can be balanced. An effective berth schedule leads to an increased port throughput and, in turn, increased port revenue [3].

Container terminals often serve vessels based on a first-come, first-served (FCFS) basis. Each vessel departs from the previous harbour at its optimal speed, without considering other vessels or possible berthing slots. As it is possible that no berthing slot will be available upon arrival, the vessel will have to wait. This practice, also known as *Sail Fast, Then Wait* or *Rush-to-Wait*, is an energy-inefficient approach [4]. It would be more energy-efficient to sail at a lower speed and be serviced directly, rather than cruising at a higher speed to minimise sailing time and waiting for anchorage until a berthing slot becomes available. In practice, berth schedules rely on uncertain or outdated Estimated Times of Arrival (ETAs). When a vessel arrives later or earlier than expected, it can disrupt the berth schedule, leading to increased port congestion, idle time, and higher fuel consumption. As a result, vessels may spend 5 to 10% of their operational time waiting to get into port [5].

A. Problem statement

To ensure that vessels do not have to wait before being able to berth, IMO has published the Just-in-Time (JIT) Arrival Guide in 2020 [5]. The goal of the Just in Time arrival guide is to allow the vessel to sail at an optimal velocity, such that it does not have to wait at anchorage. The Just-in-Time arrival guide states that waiting time at the anchorage point can be reduced through improved communication between vessels and the terminal.

Without the JIT guide, the vessel's optimal velocity ensures that its operating costs over the sailing leg are minimal, while still allowing it to depart the next terminal before its deadline. The operating costs of a vessel include fuel, crew, and chartering/depreciation. When waiting times at the anchorage point are not accounted for, the vessel will sail at its optimal speed. However, all operating costs, besides fuel consumption, will still occur when a vessel is waiting at the anchorage point. If the waiting time is known in advance, a new optimal velocity of the vessel can be devised.

Vessel speed optimisation also leads to more accurate arrival predictions, enabling the container terminal to service more vessels. Besides more accurate arrival times, vessel speed optimisation also reduces the number of vessels waiting at the anchorage point, thereby reducing pollution from waiting vessels. Additionally, decreasing the vessel's velocity on the last sailing leg reduces pollution during the voyage.

B. Literature

The background literature for this paper is divided into two separate topics: The berth allocation problem (BAP) and vessel speed optimisation (VSO).

1. Berth allocation problem under uncertainty

The berth allocation (and scheduling) problem is an operational problem that aims to plan vessel berths at container terminal quays. Different goals can be pursued in the berth allocation problem, for instance, minimising a vessel's waiting time before reaching a quay or maximising resource utilisation. The berth allocation problem has been widely researched. For an in-depth analysis of the berth allocation problem, this paper refers to the literature survey by Bierwirth and Meisel [6]. In the survey, all relevant papers up to 2015 on the topic are discussed. For the research in this paper, the topic of uncertainty in the berth allocation problem is specifically relevant. The main sources of uncertainty are the vessel's arrival time and handling time. Rodrigues and Agra classify the different uncertainty optimisation approaches into 3 key classes: proactive, reactive and reactive/proactive [7]. In proactive planning, all planning decisions are made in advance, while in reactive planning, decisions are made when uncertainties become realised. In proactive/reactive planning, decisions are made in advance, but when a large enough uncertainty is realised, the system replans. A way to deal reactively with uncertainties is through a rolling horizon. For instance, Ji et al. propose a hybrid rolling-horizon heuristic, in which the problem is solved for the new horizon when an interval has elapsed, or a new event has been detected [8]. A new event could be the arrival of an incoming vessel.

2. Vessel speed optimisation

Vessel speed optimisation can be used to achieve just-in-time sailing. Psarafitis et al. investigated speed optimisation on a single vessel to minimise the total cost of a liner vessel's trip [9]. Andersson and Ivehammer examined the benefits of adjusting speed by studying the number of vessels lying at the anchorage point of all terminals in the Baltic Sea [10]. Both studies, however, do not look into the influence of berth availability on the problem.

Alvarez et al. do look into the combination of vessel speed optimisation with vessel berthing [11]. In total, they compare 3 strategies. In their vessel speed optimisation strategy, called GOSBEA (Global Optimisation of Speed, Berth, and Equipment Allocation), vessels set sail at a recommended velocity when they are 500 nautical miles from the focal port. Every time a new vessel enters, a cost function is solved, and the vessel chooses a velocity with minimal costs. The velocity of a vessel is only optimised once. Van Leeuwen and Dekker research the influence of vessel speed optimisation of a stream of vessels arriving on a single berth terminal [12]. In the research, they use a cost function, which minimises total costs over the remaining sailing leg. The total costs consist of waiting, fuel, and operating costs, and are minimal at $\tilde{v}$ when there are no waiting times. The model's goal is to minimise total vessel costs, and it has information on every vessel arriving within a predetermined time window. So, in contrast to the research by Alvarez et al., vessels can speed up, resulting in higher savings for vessels arriving later in the stream. As the model has information on the entire stream of vessels, the velocity only has to be selected once.

3. Literature gap

Although prior studies consider berth allocation under uncertainty with vessel speed optimisation, there is limited research that integrates both perspectives. Most of the research focuses on either heterogeneous vessels or vessels departing from a fixed point. This paper aims to directly integrate vessel speed optimisation into the berth allocation problem under realistic conditions in container terminals.

C. Contributions

This paper develops a simulation-based model of a deep-sea container terminal to analyse the impact of vessel speed optimisation on quay occupancy. In contrast to existing studies, vessel speeds are dynamically optimised upon each new arrival and time interval, accounting for uncertainty in arrival times. The analysis evaluates whether vessel speed optimisation enables higher quay occupancy levels under increasing traffic intensity without negatively affecting vessel performance.

Can vessel speed optimisation enable higher berth occupancy levels without reducing vessel performance?

II. METHODOLOGY

This section explains the simulation tools used to answer the research question. The model consists of the vessel's sailing leg and the berthing and handling of vessels. The model is built in TRAFALQUAR, a verified and validated discrete-event simulation software, which simulates vessel arrivals and berthing, container handling, and vessel undocking and departure. The model is verified using AIS data on vessels arriving at Rotterdam World Gateway. The different scenarios and KPIs used to test the model are introduced.

A. Simulation scope

The model consists of two connected sub-models. A sailing leg model and a seaside operations model. The sailing leg model captures vessel movements and speed decisions, while the seaside model represents berthing

and handling processes at the container terminal. Routing, weather effects, and vessel interaction at sea are excluded from the model. The actual yard operations at the container terminal are also excluded from the model; the terminal's seaward side is simplified to a quay with quay cranes that move a fixed number of containers per hour. The goal of the combined model is to represent the sailing time, fuel consumption, and waiting time of all deep-sea vessels arriving at a terminal under realistic handling conditions.

B. Sailing leg configurations

There are two main sailing leg configurations: the baseline configuration, in which vessels depart from the previous terminal at a fixed velocity, and the global speed optimisation configuration, in which the lowest cost for all vessels is devised. The costs for the vessels are calculated by Equation 1, which is a modified variant of the formula used by van Leeuwen and Dekker [12]. The formula is modified to accommodate the vessels sailing at different velocities at different time intervals. This cost function captures the trade-off between fuel consumption, operating time, and waiting time.

$$C(v) = t_w \cdot c_w + t_s \cdot c_o + \sum_{i=1}^{n} a.S_i \cdot v_i^{b-1} \cdot c_f \qquad (1)$$

$$t_s = \sum_{i=1}^{n} \frac{S_i}{v_i} \qquad (2)$$

In which $C(v)$ is the total cost for a vessel in dollars, t_w is the waiting time in hours, c_w is the cost of waiting in dollars, t_s is the sailing time in hours, and c_o is the cost of operating in dollars. a and b are vessel parameters, v_i is the selected velocity in nautical miles per hour on a segment, S_i is the length of the segment in nautical miles, c_f is the cost of fuel in dollars, and n is the number of segments.

1. *Baseline*

This velocity is chosen to minimise operating and fuel costs over the sailing leg, without accounting for waiting time at the terminal. The baseline strategy represents the current "Sail fast, Then Wait" practice. The process of the baseline strategy is shown in Figure 1.

2. *Global speed optimisation*

The global speed optimisation strategy requires coordinating all sailing vessels and the terminal to minimise overall cost. The process is shown in Figure 2. In this

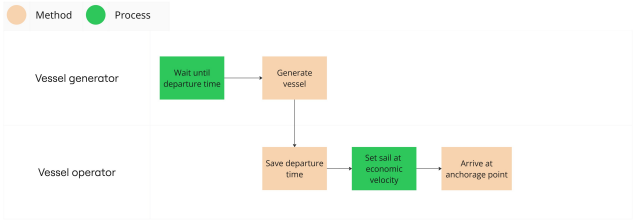


FIG. 1. Swimlane diagram of the sailing operations of a vessel without speed optimisation

strategy, the re-optimisation triggers at vessel departures from their previous ports and at fixed planning intervals. At each planning moment, all vessels' speeds (except those within 50 nautical miles of the focal port) are considered simultaneously. The goal of the function to minimise the cost of all sailing vessels.

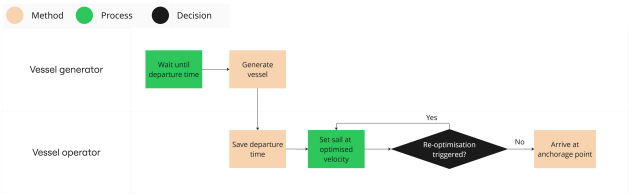


FIG. 2. Swimlane diagram of the global speed optimisation configuration

To solve this problem, the optimisation is performed in two stages. First, the berth sequence is determined by evaluating different arrival orders and selecting the one with the lowest overall cost.

Given the arrival sequence, the vessels' speeds are optimised. From this sequence, all speed-up possibilities from every vessel (except those that arrive at the end of a sequence) are evaluated. The solution with the lowest overall cost is then selected.

C. Seaside operations

The seaside operations of the container terminal are modelled in TRAFALQUAR. In the seaside operations, berthing assignments for vessels and container loading and unloading are modelled. In Figure 6, the swimlane diagram for the berth assignment is displayed and in Figure 7, the swimlane diagram for the loading and unloading of vessels is shown. These swimlanes are connected. The final process in Figure 1 and Figure 2 are the first process of Figure 6, and the line from the final process of Figure 6 goes to the first method of Figure 7.

The model of the seaside operations should represent the deep-sea quay at Rotterdam World Gateway. The quay is 1150 meters long. On the quay, 11 deep-sea cranes are stationed, each capable of moving 26 containers per hour. These cranes work continuously in shifts for 7:30

hours, then take a 0:30-hour break. Each crane can only work on 1 bay of the vessel. Each vessel can berth anywhere on the quay; however, once berthed, it cannot be moved.

D. Verification and validation

The model is built upon the verified and validated TRAFALQUAR model. The baseline model's input is validated using AIS data on vessel arrivals at Rotterdam World Gateway in 2024. The output values for waiting time and service time are compared with the data values. The logic of the speed optimisation model is verified through code consistency checks and inspection of simulation outputs.

E. Scenarios

The model is tested on three different scenarios. Each scenario consists of a number of vessels arriving every 2 weeks. The set of vessels is repeated to simulate a year of vessel arrivals. Scenario 1 consists of 11 vessels, scenario 2 consists of 12 vessels, and scenario 3 consists of 13 vessels. In Table I, the vessel class and previous port are defined. The vessels without a star (*) on the day of arrival are all in scenario 1; the vessel with 1 star is added to this set for scenario 2; and, for scenario 3, all vessels in Table I are in the set. In scenario 2, an additional 44.000 additional TEU are handled annually, and in scenario 3, an additional 130.000 TEU are handled annually compared to scenario 1.

Vessel class	**Previous port**	**Distance to previous port** (NM)	**Day of arrival**
TEU 7	Singapore	1200	1
TEU 4	Le Havre	250	1*
TEU 4	Antwerp	140	3
TEU 7	Tanger	1200	4
TEU 6	Algeciras	1200	6
TEU 7	Dunkerque	140	7
TEU 4	Bremerhaven	250	8
TEU 7	Singapore	1200	9
TEU 5	Antwerpen	140	10**
TEU 6	Gdansk	850	11
TEU 7	Algeciras	1200	12
TEU 5	Southampton	250	13
TEU 4	Hamburg	300	14

TABLE I. Vessel input parameters; the days of arrival marked with a star (*) are part of scenario 2 and scenario 3

The actual moment of arrival and the call size of a vessel are determined by the realisation of stochastic parameters at the start of a simulation. The moment of departure at the previous terminal is determined by the realised moment of arrival minus the distance to the previous port, divided by the mean velocity of a vessel. Each vessel has the same minimum, mean ($\tilde{v}$), and maximum velocities at 10, 16, and 22 NM/h. The other vessel parameters for each vessel class are listed in Table I. $\tilde{\boldsymbol{f}}$ is the fuel consumption of a vessel in metric tonnes sailing at its mean velocity for a day. The fuel consumption per vessel class is derived from the research by Brouer et al. [13].

Vessel class	**Length** (m)	**Call size** (TEU)	$\tilde{\boldsymbol{f}}$ (mt/day)
TEU 4	285	1650	75
TEU 5	345	3300	100
TEU 6	365	5000	125
TEU 7	400	8500	200

TABLE II. Vessel parameters

The cost of one metric ton of fuel is taken from Ship & Bunker. The average cost of one metric tonne of Very Low Sulphur Fuel Oil in Rotterdam during 2024 was \$555.5 dollar [14], which is used as c_f in the simulation.

F. Performance indicators

The model is tested using 3 performance indicators: berth occupancy, average vessel cost, and average vessel departure delay. The berth occupancy is directly dependent on the scenario. It is defined as the fraction of total quay length occupied over time, so an increase in the number of arrivals also increases the berth occupancy. The average vessel cost is calculated per simulation and divided into waiting, operating, and fuel costs, based on all vessels in the simulation. The departure delay is also taken as the average across the entire simulation. Vessel performance in the research question depends on average vessel costs and average departure delay.

III. RESULTS

In this section, the results of testing the model across configurations and scenarios are presented. In total, 4 sets with 10 replications were simulated. The baseline configuration under scenario 1, and the global speed optimisation under scenarios 1, 2 and 3. For every simulation, 95 % confidence intervals are constructed around the average vessel cost and departure delay.

A. Berth occupancy

In Figure 3, the berth occupancy of all four scenarios and configurations is displayed. The berth occupancy of scenario 2 increases by 2.2%, and the berth occupancy of scenario 3 increases by another 3.5%, as a result of the extra vessel arrivals.

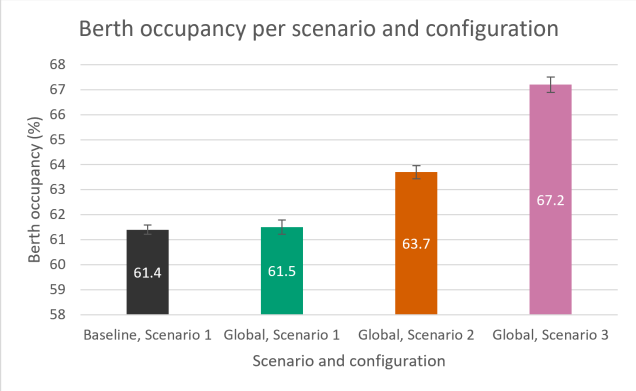


FIG. 3. Berth occupancy for the different scenarios and configurations

B. Average vessel cost

In Figure 4, the average vessel cost per scenario and configuration is displayed. The average vessel cost excludes the additional vessels (TEU 4 from Le Havre and TEU 5 from Antwerp), as they are not included in scenario 1 and would skew the average cost. The costs are thus compared to the cost of the original set of vessels, to ensure comparability with the other scenarios.

By comparing the three scenarios in Figure 4, the global speed optimisation configuration shows an increase in waiting cost and a decrease in fuel cost when the number of vessels increases. This indicates that vessels adapt by sailing more slowly to avoid congestion, thereby reducing fuel costs and increasing operating costs. When comparing the global speed optimisation configuration, scenario 2, to the baseline, scenario 1, it performs at a lower average cost, whereas the global speed optimisation configuration, scenario 3, results in a higher average cost, indicating worse performance.

The 95% confidence intervals half-widths of the average vessel costs are at \$5285 for baseline, scenario 1, \$4782 for global speed optimisation, scenario 1, \$3318 for global speed optimisation, scenario 2, and \$5449 for global speed optimisation, scenario 3.

C. Average departure delay

By comparing the three scenarios in Figure 5, the global speed optimisation configuration shows an increase in average departure delay as the number of vessels increases. However, the global speed optimisation configuration in scenario 2 results in a lower average departure delay than the baseline configuration, although the difference is not statistically significant. In contrast, the global speed optimisation configuration, under scenario 3, performs worse than the baseline in terms of departure delay.

The 95% confidence intervals half-widths for the departure delay are 0.7 hours for baseline, 1.1 hours for global speed optimisation, scenario 1, 1.6 hours for global speed optimisation, scenario 2, and 1.3 hours for global speed optimisation, scenario 3.

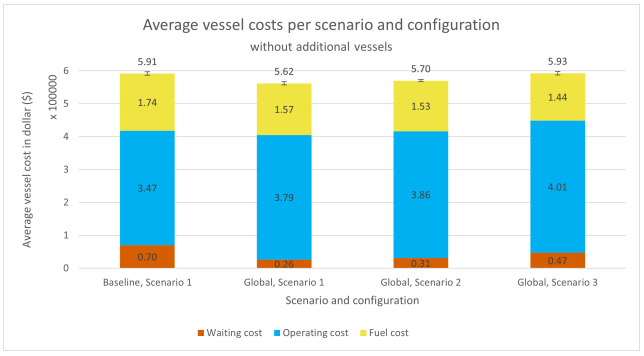


FIG. 4. Average vessel cost for the different scenarios and configurations

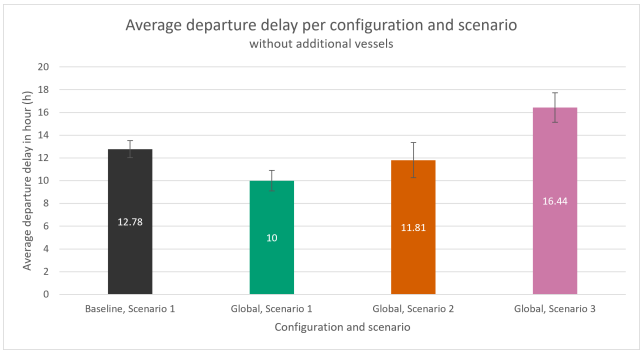


FIG. 5. Average vessel departure delay for the different scenarios and configurations

IV. DISCUSSION

This section reflects on the simulation results and discusses their implications. Further, reflects on the study as a whole and recommends interesting parts for future research.

A. Discussion on results

As Figure 3 indicates, adding additional vessels in the case of scenario 2 and scenario 3 increases the terminal's berth occupancy. However, increasing the number of vessels reduces the performance of both the terminal and the vessels, as average vessel costs and average vessel departure delays increase. In the global speed optimisation scenario 2, the average vessel cost remains lower than in the baseline, and the average departure delay remains similar. However, for individual vessels, the departure

delay can increase. Concluding that vessel speed optimisation can increase the terminal's overall berth occupancy; however, the effect is limited, as the terminal's service level decreases when too many vessels are added. Scenario 2 represents the highest tested increase in vessel arrivals while maintaining lower average vessel costs and departure delays statistically similar to the baseline. While scenario 3 leads to a deterioration in both vessel costs and departure delays beyond the baseline in scenario 1. This indicates that the terminal is approaching the capacity limit, beyond which additional vessels lead to disproportionate increases in delay and costs.

The half-width of the confidence intervals around the mean of the average vessel costs is generally small. In the baseline configuration at scenario 1, the half-width is $5285, while the mean value for average vessel costs is at $591396 dollars, indicating a deviation of around 1%. However, the difference in average vessel costs between the baseline and global speed optimisation in scenario 1 is approximately $30000. This difference is about six times the confidence interval half-width. In the Appendix, Figure 8 displays the vessel costs per previous destination. In this case, the half-widths of the confidence intervals are on average larger. For instance, the average half-width of the confidence interval at the baseline is $15000. This indicates that the performance of individual vessels can vary across simulations; however, the overall system's performance remains intact. When comparing Figure 5 to the departure delay per vessel, in the Appendix, in Figure 9, a similar trend is visible. The average departure delay over the simulation is smaller than the average departure delay per vessel per simulation, a result of the stochastic realisation of arrival times.

Figure 8 and Figure 9 also give an indication of the performance of individual vessels. Showing that while the average cost per vessel, in the global speed optimisation scenario 2, is lower for most vessels, the departure delay, especially of vessels in the TEU 6 and TEU 7, is already higher than at the baseline scenario 1. Indicating an already slight drop in performance.

These findings directly answer the research question by showing that vessel speed optimisation enables a limited increase in berth occupancy without degrading average vessel performance, but only up to a critical threshold.

B. Discussion on study

The result of this study relies on several simplifications and assumptions. Some of those simplifications or assumptions can affect the research if they are unrealistic.

The main assumptions concern vessels departing at a fixed velocity. In the model, it is assumed that, in the baseline configuration, vessels depart from their previous port at a fixed velocity of 16 NM/h. However, this assumes that there is no contact between vessel operators and terminal operators, who may already influence the velocity of incoming vessels.

Besides the fixed velocity, the willingness of incoming vessels to change their velocity can also be up for debate. In one of the results, a vessel had to change its velocity up to 23 times. These velocity changes were mainly between 15 NM/h and 16 NM/h, and back, depending on the best time of arrival at the calculation point. Although 23 velocity changes are an excess, if more vessels are added to the biweekly schedule, or if the maximum distance is above 1200 NM away from the vocal port, this number may even increase. The behaviour of changing the velocity 23 times or more during a voyage may not be feasible in practice.

The exogenic input of vessel arrivals is also simplified into a schedule that repeats every 2 weeks for an entire year. However, from the AIS data, vessel arrivals were not as consistent, especially at the vessels' previous ports. For instance, 16 out of the 19 vessels that departed from Tanger arrived before 01-06-2024, while 28 out of the 33 that departed from Dunkerque arrived after 01-06-2024. These changes can be caused by the liner taking a different route through Europe. However, the remaining distance to the new port heavily influences the arrival time of incoming vessels and thus the system's reliability.

C. Future research

The first recommendation ties into one of the discussion points. If it is possible to accurately predict when vessels depart their previous port, they can be scheduled into berth planning at the new port much earlier. With more knowledge of the vessel departing from nearby, vessels from further out no longer have to change their velocity as harshly or as often.

Another recommendation is to research the influence of speed optimisation for deep-sea vessels on all arriving vessels at the container terminal. In this case, short-sea vessels and barges must be included in the simulation. This could be done in two ways: either by letting the vessels arrive directly at the terminal without accounting for their sailing leg, or by simulating their sailing leg and optimising their velocity. However, by also optimising the velocity of the short-sea vessels, another 50 vessels must be simulated every 2 weeks, and by optimising the barges, another 154 vessels must be simulated. The model's current solution mechanism is not suited to this, so heuristics must be implemented to solve the problem.

In this research, every deep-sea vessel has the same $\tilde{v}$ at 16 NM/h. This velocity has been chosen so that the fuel consumption and costs of each vessel within the same class remain the same. However, the vessels within a class do not always set sail at the same velocity, so $\tilde{v}$ can be made a variable to range from 14 NM/h to 18 NM/h, with the maximum velocity being $\tilde{v}$ + 6 NM/h and the minimum velocity being $\tilde{v}$ - 6 NM/h. By varying $\tilde{v}$ across vessel classes and individual vessels, the model would yield more realistic results.

V. CONCLUSION

This paper investigated whether vessel speed optimisation could increase berth occupancy in a container terminal without deteriorating the performance of arriving vessels. The results from baseline scenario 1, global speed optimisation scenario 1, global speed optimisation scenario 2, and global speed optimisation scenario 3 are discussed. In scenario 2, the berth occupancy increases by 2.2%, and in scenario 3, the berth occupancy increases by 5.7% compared to the baseline. However, in scenario 3, the vessels perform worse on departure delays and average vessel costs than in the baseline of scenario 1. In the global speed optimisation configuration in scenario 2, the average vessel cost is lower, and the average departure delay remains statistically similar to the baseline configuration in scenario 1. Although for some specific vessels, the departure delay and vessel costs will increase compared to the baseline. Concluding that, with the global vessel speed optimisation strategy, berth occupancy can be increased by 2.2% without decreasing the average vessel performance. This increase would lead to an additional 44.000 TEU handled annually. Increasing berth occupancy further is not recommended, as it leads to congestion that deteriorates the overall system's performance.

VI. APPENDIX

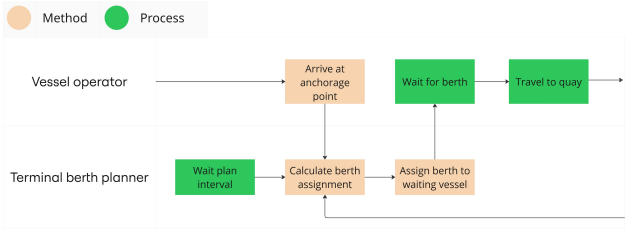


FIG. 6. Swimlane of the berth assignment

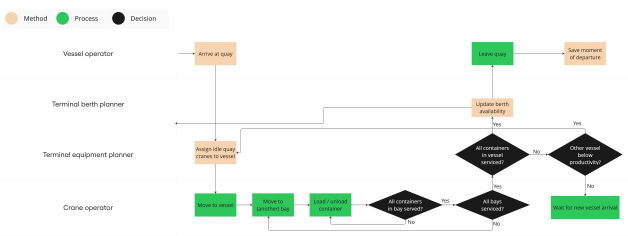


FIG. 7. Swimlane of the loading and unloading of containers

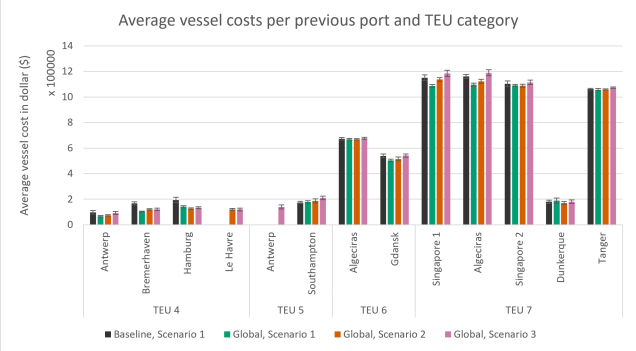


FIG. 8. Average vessel cost per vessel for the different scenarios and configurations

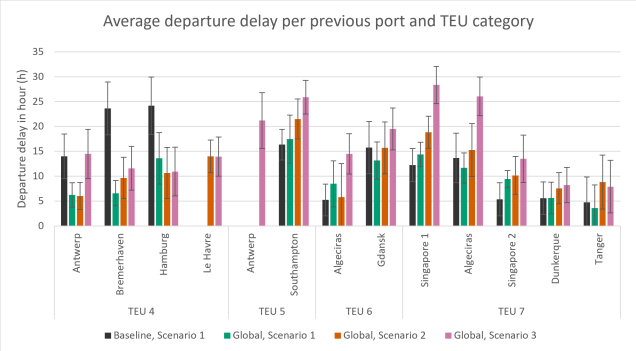


FIG. 9. Average departure delay per vessel for the different scenarios and configurations

REFERENCES

[1] U. N. C. on Trade and Development, *Review of Maritime Transport 2024*, Technical report (United Nations conference on Trade and Development (UNCTAD), 2024).
[2] International Maritime Organization, "Fourth Greenhouse Gas Study 2020," (2020).
[3] M. M. Golias, H. E. Haralambides, and M. Nienhuis, Transportation Research Part E: Logistics and Transportation Review **45**, 878 (2009).
[4] I. Sung, H. Zografakis, and P. Nielsen, Transportation Research Part D: Transport and Environment **110** (2022), 10.1016/j.trd.2022.103407.
[5] International Maritime Organization, "Just in time arrival guide," (2020).
[6] C. Bierwirth and F. Meisel, European Journal of Operational Research **244**, 675 (2015).
[7] F. Rodrigues and A. Agra, European Journal of Operational Research **303**, 501 (2022).
[8] B. Ji, M. Tang, Z. Wu, S. S. Yu, S. Zhou, and X. Fang, Advanced Engineering Informatics **54** (2022), 10.1016/j.aei.2022.101733.
[9] H. N. Psaraftis and C. A. Kontovas, Transportation Research Part C: Emerging Technologies **44**, 52 – 69 (2014).
[10] P. Andersson and P. Ivehammar, Transportation Research Part D: Transport and Environment **51**, 240 – 249 (2017).

[11] J. Alvarez, T. Longva, and E. Engebrethsen, Maritime Economics and Logistics **12**, 327 (2010).
[12] P. W. A. van Leeuwen and R. Dekker, International Transactions in Operational Research (2025), 10.1111/itor.70108.
[13] B. D. Brouer, J. F. Alvarez, C. E. M. Plum, D. Pisinger, and M. M. Sigurd, Transportation Science **48**, 281 – 312 (2014).
[14] "Rotterdam bunker prices," (2024).

B

Results

The appendix provides additional model outputs that support the interpretations of the main results presented in chapter 7. These performance indicators are not directly used to answer the research questions, but offer insight into the behaviour of the model. The results are included for completeness, and provide additional context.

B.1. T-test on sensitivity analysis

This section presents the p-values obtained from the pairwise t-test performed on the sensitivity analysis results. The p-values in Table B.3 show that no statistical differences were found between the mean departure delays for all vessels, as all p-values are greater then 0.05. Similarly Table B.2 indicates no statistical significant differences for the TEU 6 and TEU 7 vessels.

For the TEU 4 and TEU 5 vessels in Table B.1, the comparisons with the global speed optimisation scenario also result in non significant differences ($p > 0.05$), although the comparison between sensitivity 1 and sensitivity 2, as well as the comparison between sensitivity 3, and and sensitivity 2 show statical significant differences ($p = 0.001$). In contrast, sensitivity 1 and 3 show significant similarities with $p = 0.775$.

	Global speed optimisation	**Sensitivity 1**	**Sensitivity 2**
Sensitivity 1	0.085	x	x
Sensitivity 2	0.057	0.001	x
Sensitivity 3	0.108	0.775	0.001

Table B.1: Pairwise t-test p-values for the sensitivity tests on scenario 1, global speed optimisation for the average departure delay on the TEU 4 and 5 vessels

	Global speed optimisation	**Sensitivity 1**	**Sensitivity 2**
Sensitivity 1	0.576	x	x
Sensitivity 2	0.880	0.580	x
Sensitivity 3	0.436	0.774	0.454

Table B.2: Pairwise t-test p-values for the sensitivity tests on scenario 1, global speed optimisation for the average departure delay on the TEU 6 and 7 vessels

	Global speed optimisation	Sensitivity 1	Sensitivity 2
Sensitivity 1	0.352	x	x
Sensitivity 2	0.339	0.123	x
Sensitivity 3	0.733	0.697	0.292

Table B.3: Pairwise t-test p-values for the sensitivity tests on scenario 1, global speed optimisation for the average departure delay on all vessels

B.2. Velocity changes

B.2.1. All configurations, scenario 1

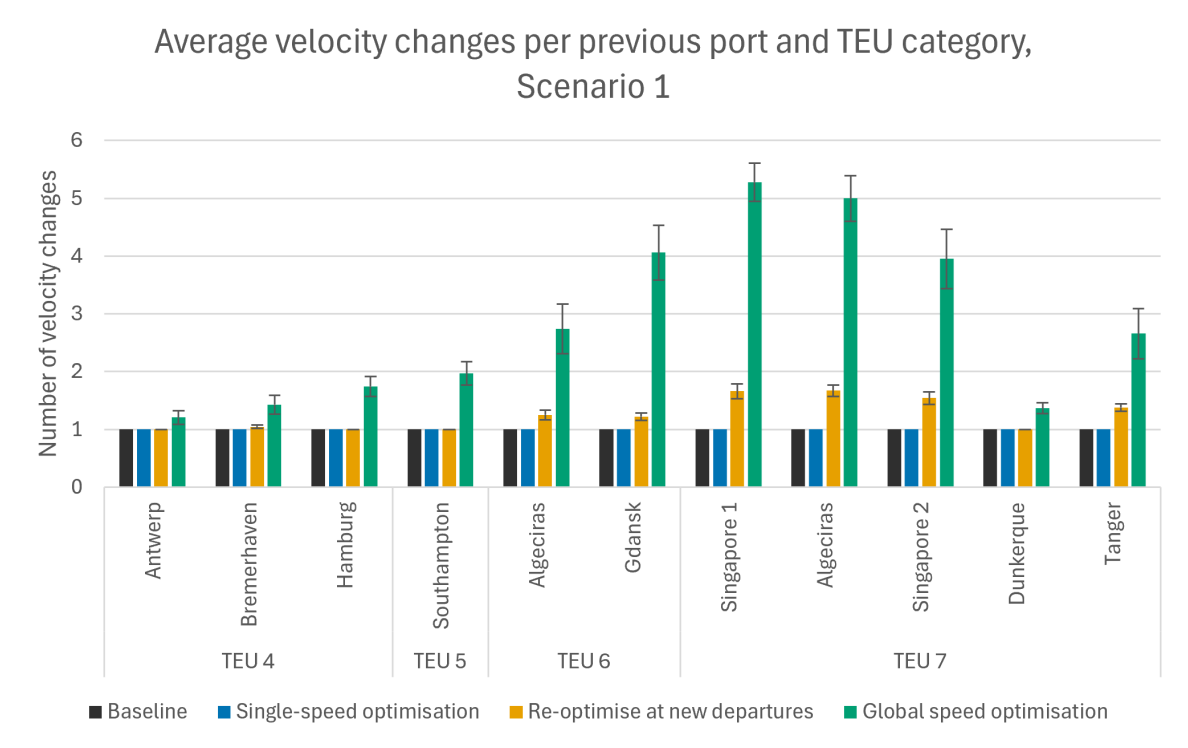


Figure B.1: The number of velocity changes in scenario 1, under the four different configurations per port and TEU category

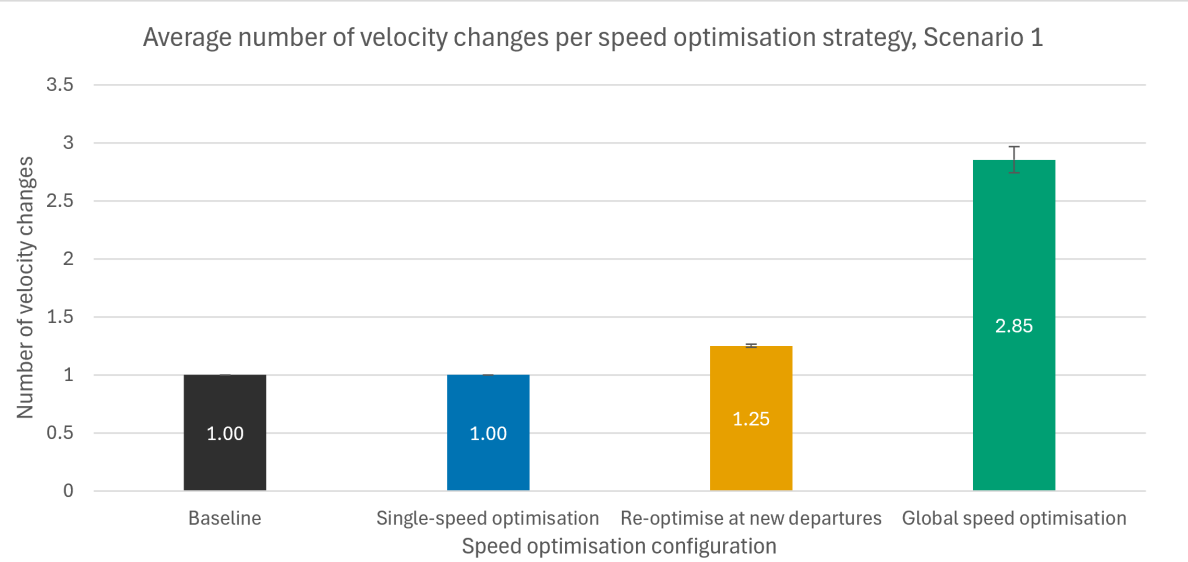


Figure B.2: The average number of velocity changes in scenario 1, under the four different configurations

The number of velocity changes for the baseline, and single-speed optimisation is 1, as both select a velocity at the moment of departure. A large increase can be observed between re-optimise at new departures and global speed-optimisation in Figure B.2. The increase is due to now optimising at new departures, and at time intervals. The peaks in Figure B.1 are generally observed for vessels with a larger starting distance, as these are often subject to re-optimisations.

B.2.2. Baseline scenario 1 and global speed optimisation scenario 1, 2 and 3

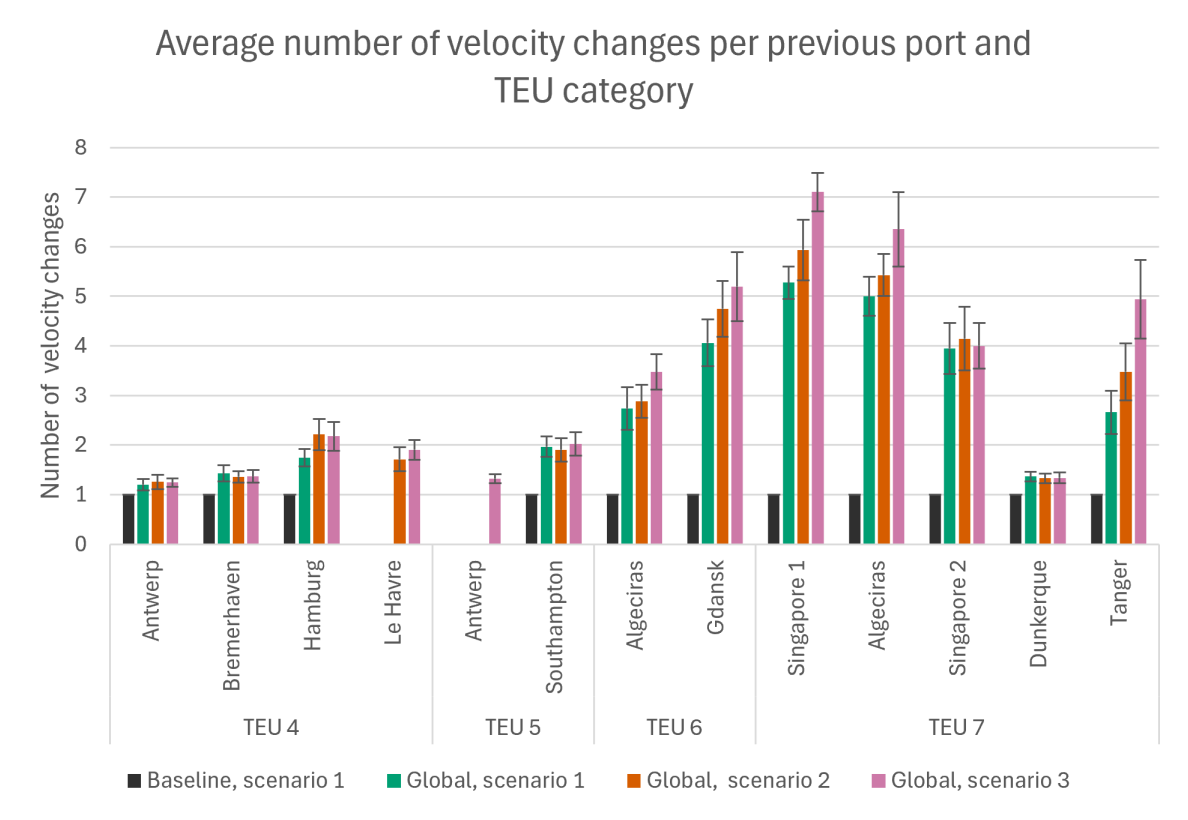


Figure B.3: The number of velocity changes in baseline scenario 1 and global speed optimisation scenarios 1, 2 and 3 per port and TEU category

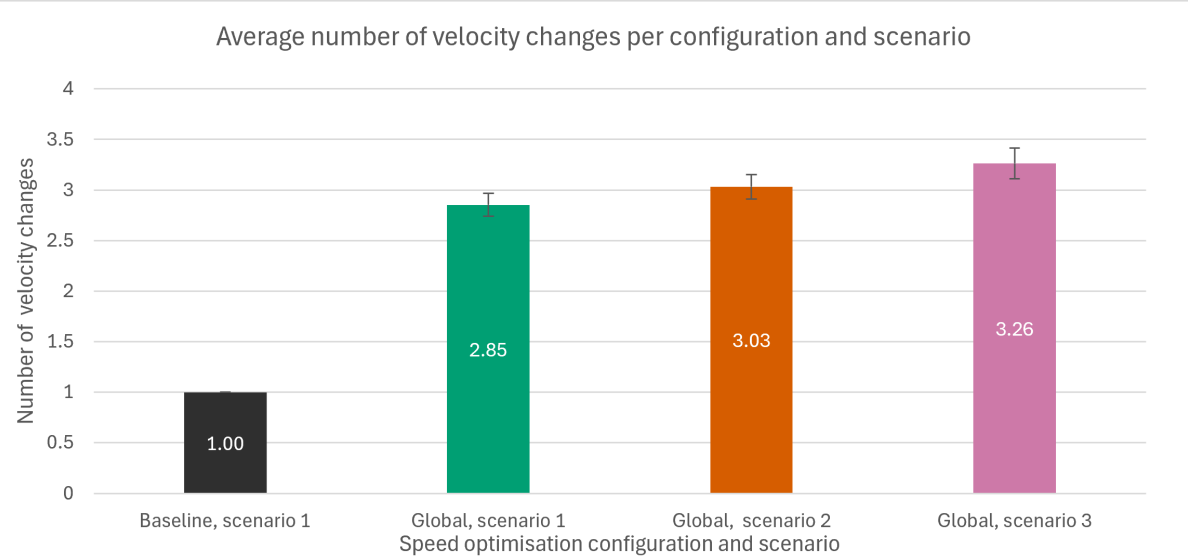


Figure B.4: The number of velocity changes in baseline scenario 1 and global speed optimisation scenarios 1, 2 and 3

The number of velocity changes, in global speed optimisation scenarios 1, 2 and 3, increases with the number of vessels.

B.2.3. Maximum number of velocity changes

In Figure B.5, the changes in velocity from a vessel are displayed. The example is from the TEU 7 vessel, departing from Singapore. Most of the fluctuations in the velocity of +/- 1 NM/h can be explained by new information from departures on the quay.

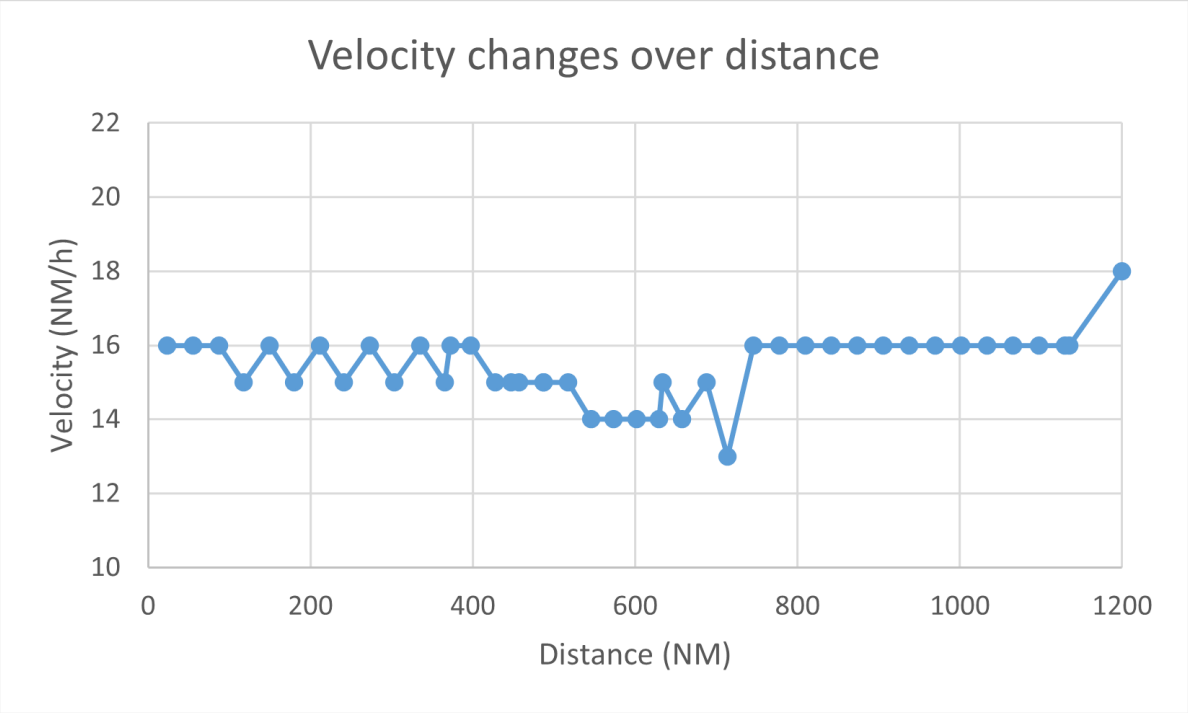


Figure B.5: Velocity changes in a TEU 7 vessel

B.3. Average velocity

B.3.1. All configurations, scenario 1

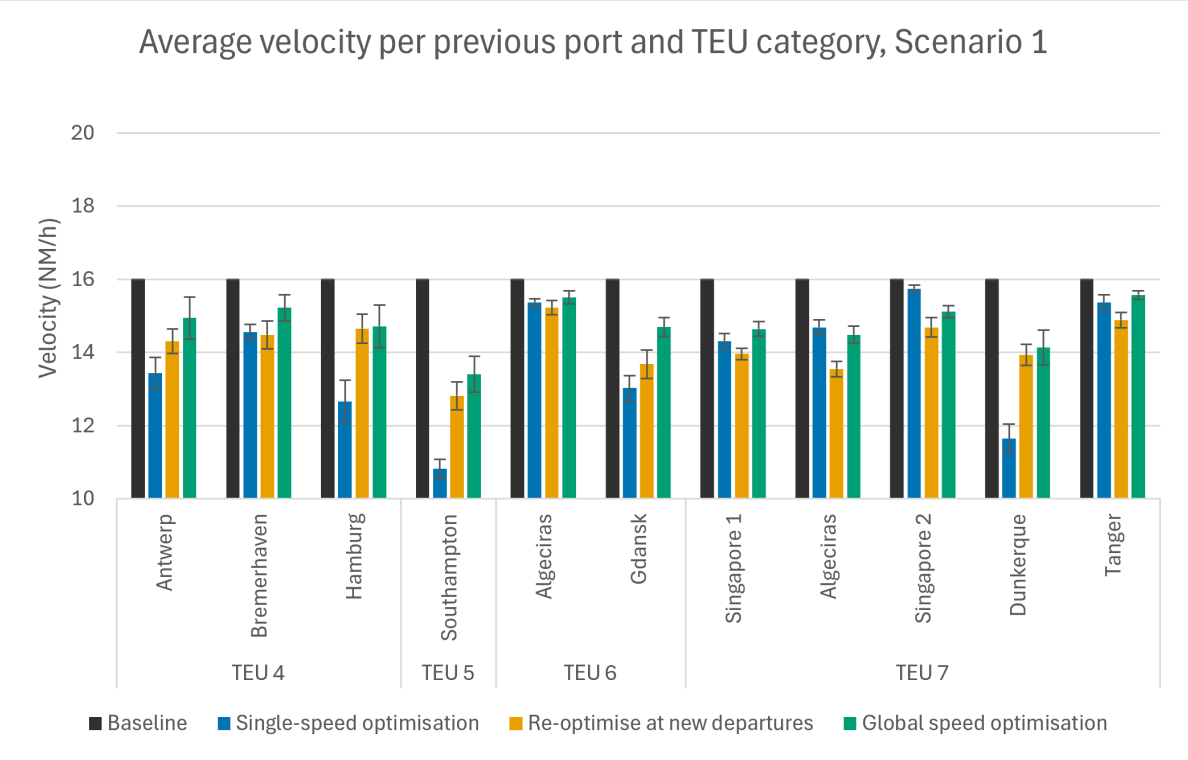


Figure B.6: The average velocity in scenario 1, under the four different configurations per port and TEU category

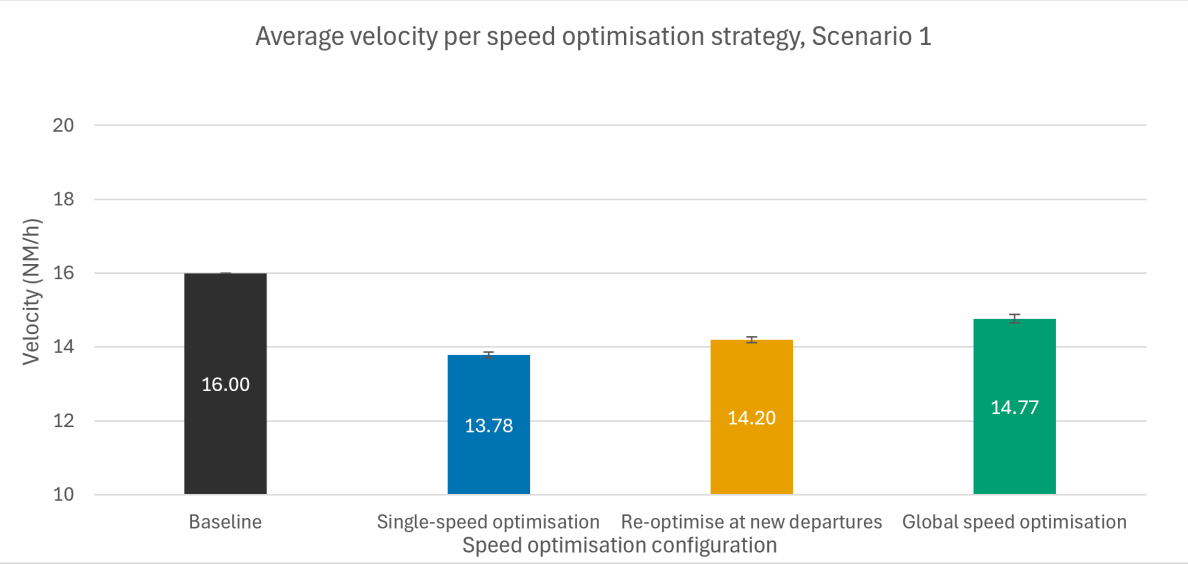


Figure B.7: The average velocity in scenario 1, under the four different configurations

The average velocity for the baseline configuration in scenario 1 is fixed at 16 ($\tilde{v}$); however, an increase in velocity is observed in re-optimisation at new departures and global speed optimisation, compared to single-speed optimisation.

B.3.2. Baseline scenario 1 and global speed optimisation scenario 1, 2 and 3

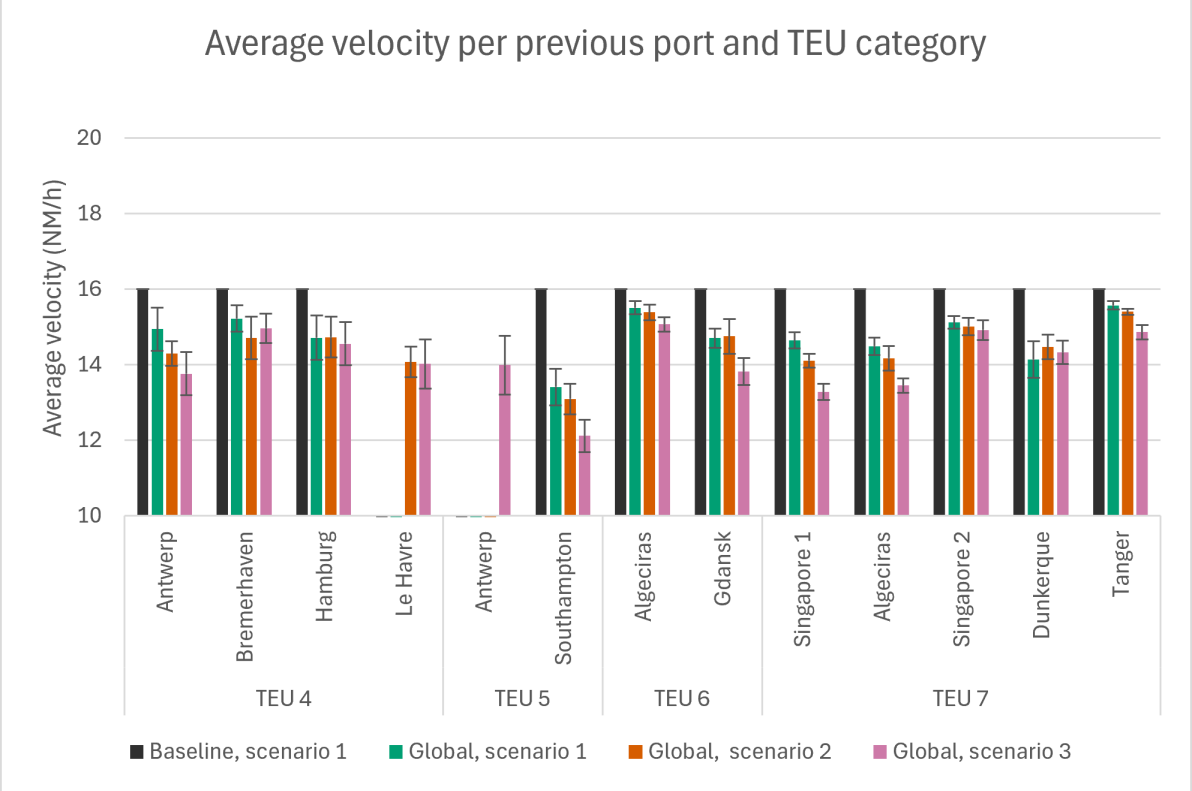


Figure B.8: The average velocity in baseline scenario 1 and global speed optimisation scenario 1, 2 and 3 per port and TEU category

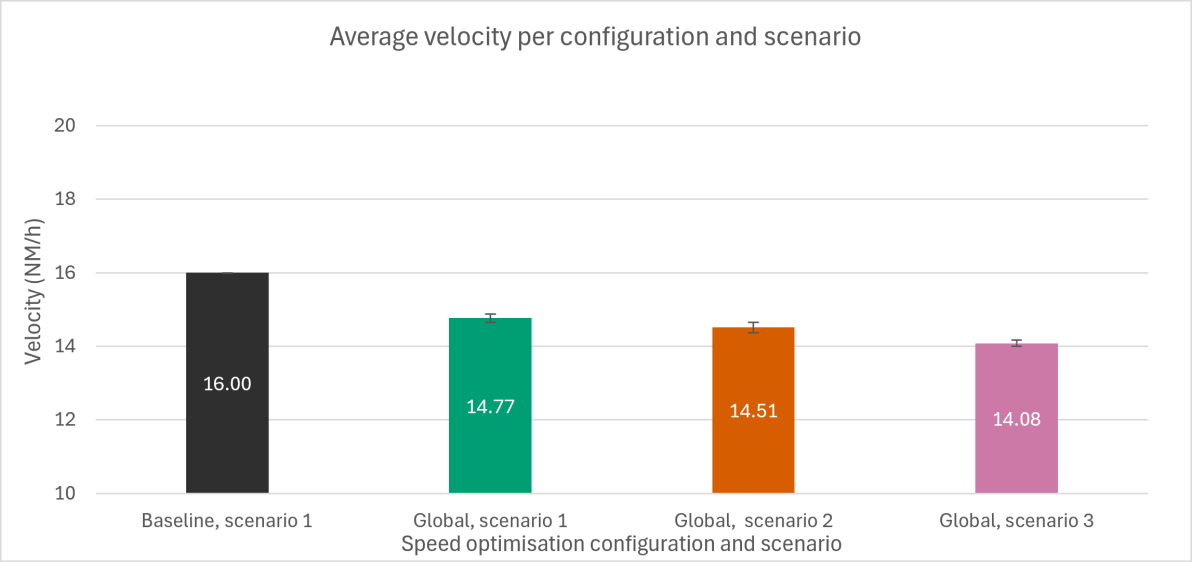


Figure B.9: The average velocity in baseline scenario 1 and global speed optimisation scenarios 1, 2 and 3

When comparing the global speed optimisation under scenarios 1, 2 and 3, the average velocity decreases. The velocity decreases due to increased congestion in the port.

B.4. Vessel waiting time

B.4.1. All configurations, scenario 1

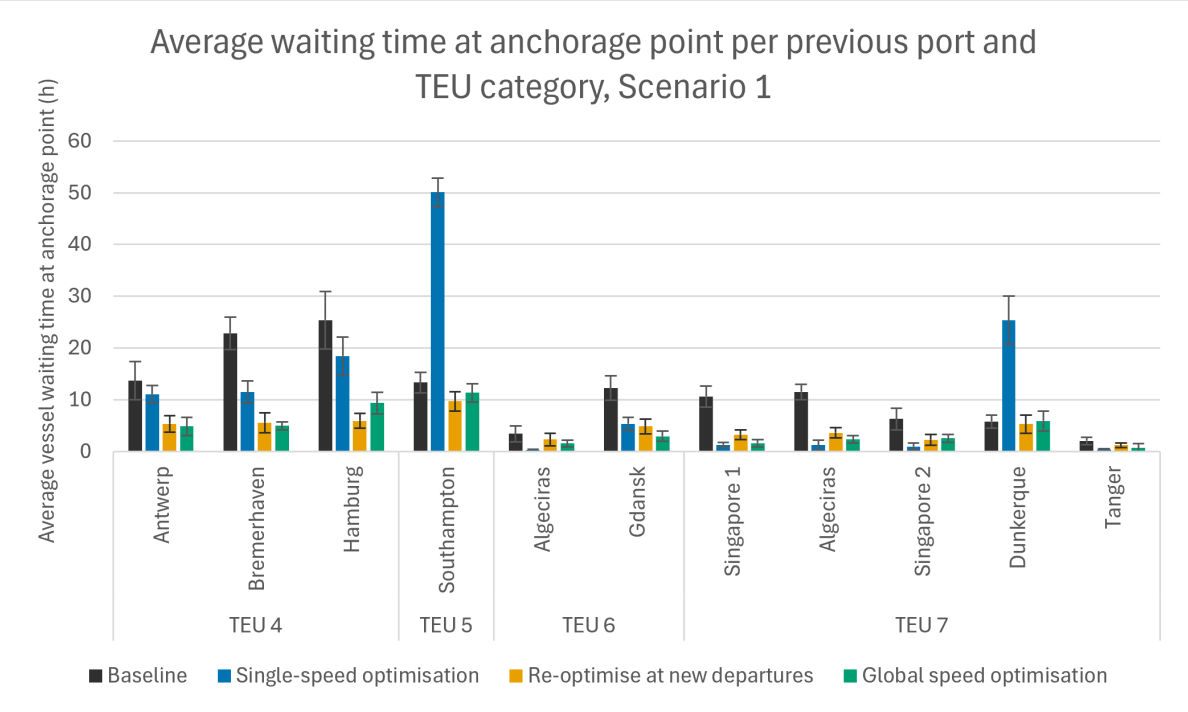


Figure B.10: The average waiting time in scenario 1, under the four different configurations per port and TEU category

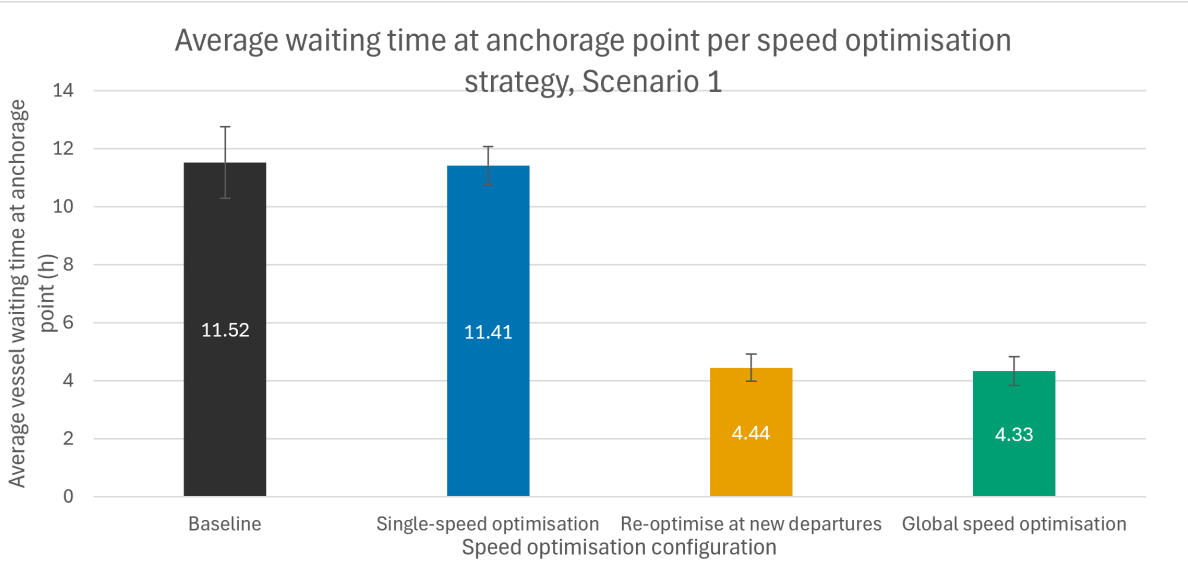


Figure B.11: The average waiting time in scenario 1, under the four different configurations

Although the average departure delay increased, the actual waiting time at the anchorage point decreased from the baseline to the single-speed optimisation. However, as can be seen in Figure B.10, the waiting times for the TEU 5 vessel from Southampton and the TEU 7 vessel from Dunkerque increase significantly. This supports the peaks in costs and departure for the same vessel observed in Chapter 7. The waiting time for both re-optimise at new departures and global speed optimisation is almost the same, and significantly decreases the waiting time at the port compared to

the baseline.

B.4.2. Baseline scenario 1 and global speed optimisation scenario 1, 2 and 3

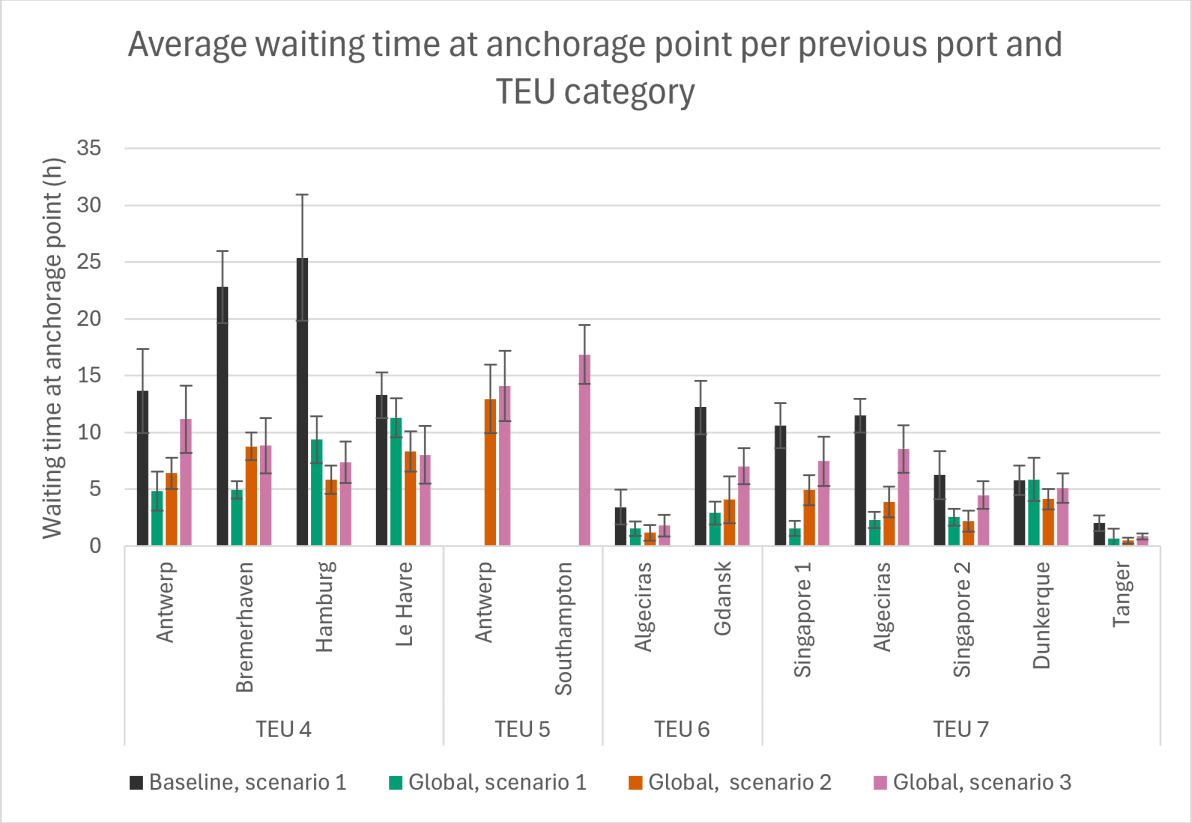


Figure B.12: The average waiting time in baseline scenario 1 and global speed optimisation scenario 1, 2 and 3 per port and TEU category

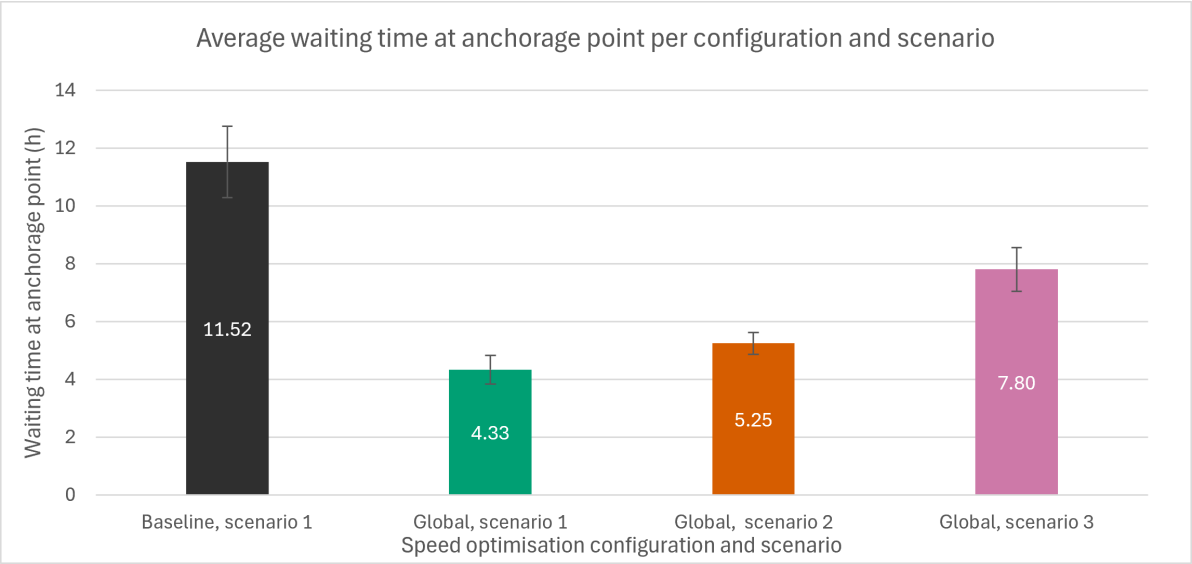


Figure B.13: The average waiting time in baseline scenario 1 and global speed optimisation scenarios 1, 2 and 3

When comparing the global speed optimisation across scenarios 1, 2, and 3, the waiting time increases with increasing vessel numbers. This increase can be explained by the increased

congestion in the port. However, the waiting time at the anchorage point is still lower for all vessels than in the baseline.

C

Verification checks

In this chapter, additional verification steps are presented to demonstrate that the implemented decision logic operates as intended. First, the logic of the speed optimisation strategies is verified by comparing their outputs within the same parameter realisation based on the realised berth schedule of the first four weeks. The realised berth schedules are split into two-week segments. For each of these weeks, the vessels' locations in time and at the quay are compared. Besides the berth schedule, examples are given for the speed selection in single-speed optimisation and for selecting the best order and best speed-up combination in global speed optimisation. Finally, the berth occupancy results are verified by showing where the deviations stem from and by calculating the expected increase in berth occupancy by hand.

C.1. Realised berth schedule

First, the realised berth schedules of the first two weeks of simulations are shown. To highlight the influence of the strategies, the location of the same vessels within different strategies is compared.

C.1.1. Weeks 1 and 2

The first two weeks are mainly compared on two vessels. The TEU 4 vessels from Antwerp, in light brown, and the TEU 6 vessel from Gdansk, in grey. In Figure C.1, the TEU 4 vessel from Antwerp has arrived before the TEU 7 vessel from Tanger, due to the higher priorities of other services. Its service started when the TEU 7 vessel from Singapore departed. The TEU 6 vessel from Gdansk has started its service, while the vessel from Dunkerque has finished its service. In the figure, for each vessel, its TEU category, port of departure, average velocity, and final velocity are specified. In the case of the baseline strategy, the average and final velocity of all vessels is at 16 NM/h.

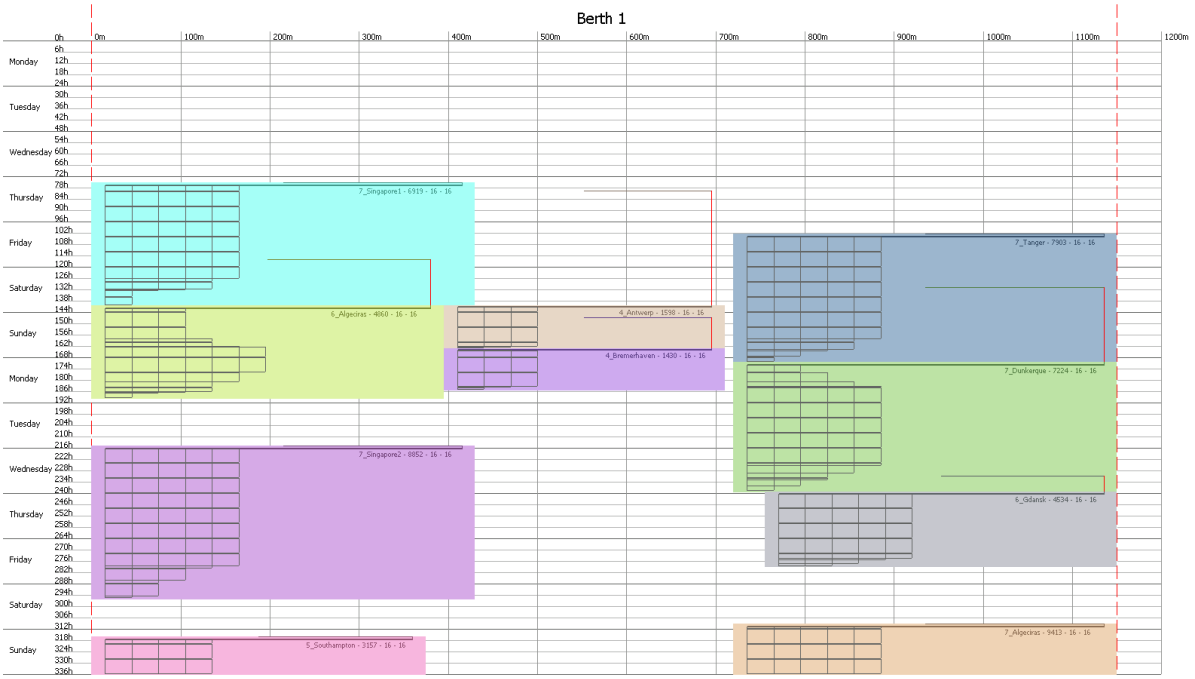


Figure C.1: Realised berth schedule of week 1 and 2 in scenario 1, configuration baseline

In the Figure C.2, the overall waiting time of vessels is decreased, due to the lower velocity of incoming vessels. Although the vessel from Antwerp has yet again arrived before the vessel from Tanger, it could not start berthing before it, as its expected end of service was after the vessel from Tanger's start of service. The reason that the vessel from Gdansk has sailed at 10 NM/h is that at the moment of its departure, the vessel of Dunkerque was being serviced by 3 cranes. Expected end of service is defined as the current productivity divided by the remaining call size. So, for the vessel from Dunkerque, the expected end of service was thus later than in reality. However, because a vessel's velocity is updated only once, the vessel from Gdansk could not speed up to avoid this idle time.

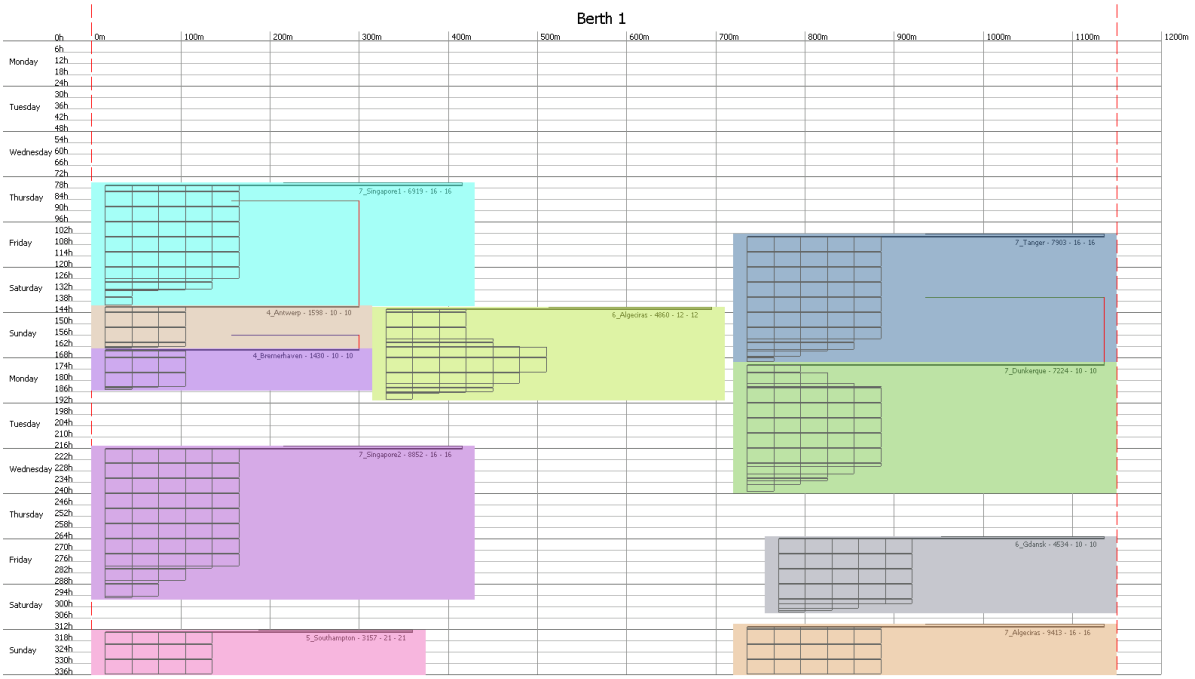


Figure C.2: Realised berth schedule of week 1 and 2 in scenario 1, configuration single speed optimisation

In the realised berth schedule from re-optimising at every new departure in Figure C.3, the vessel from Antwerp has arrived before the vessel from Tanger, as at the time of departure from Antwerp, it had fewer nautical miles remaining than the vessel from Tanger. Something similar happened to the vessel from Dunkerque, which has now started berthing ahead of the vessel from Algeciras, causing the schedule to change.

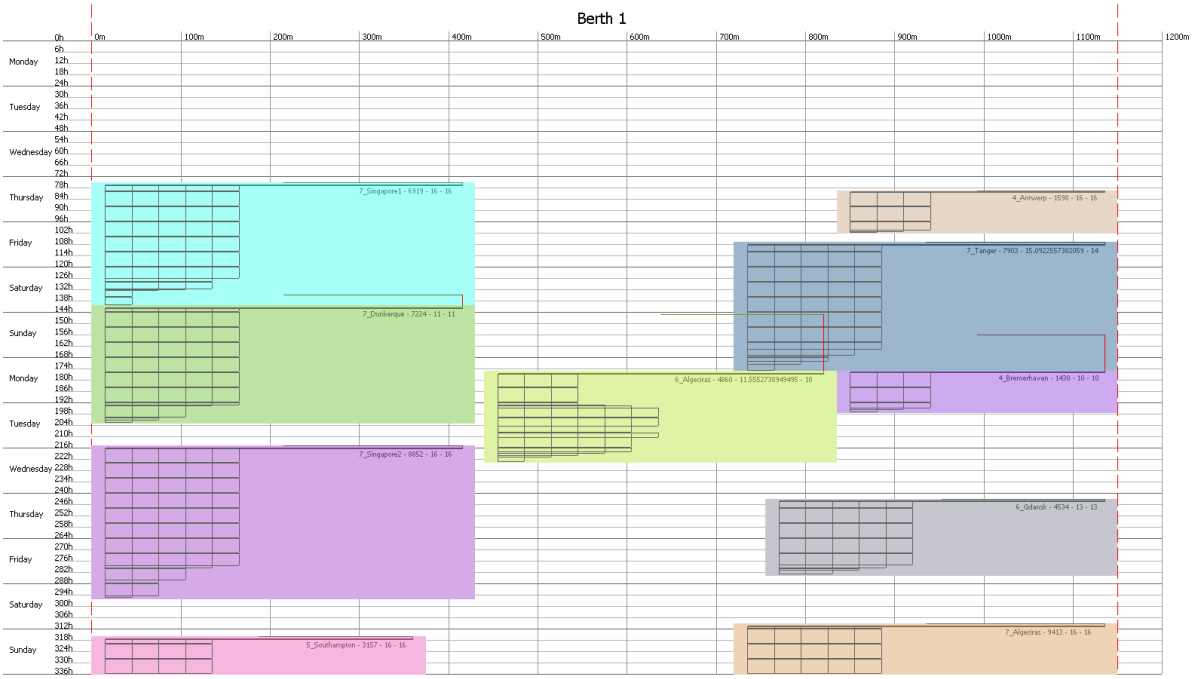


Figure C.3: Realised berth schedule of week 1 and 2 in scenario 1, configuration re-optimise at new departures

The realised berth schedule of the global speed optimisation strategy in Figure C.4, is similar to the berth schedule of re-optimise at new departures. However, vessel waiting times have decreased. This is due to individual vessels sailing above their $\tilde{v}$, which increases their own costs, but decreases the overall system costs. The vessel from Antwerp now sets sail at 17 NM/h, allowing the vessel from Tanger to sail almost at its $\tilde{v}$.

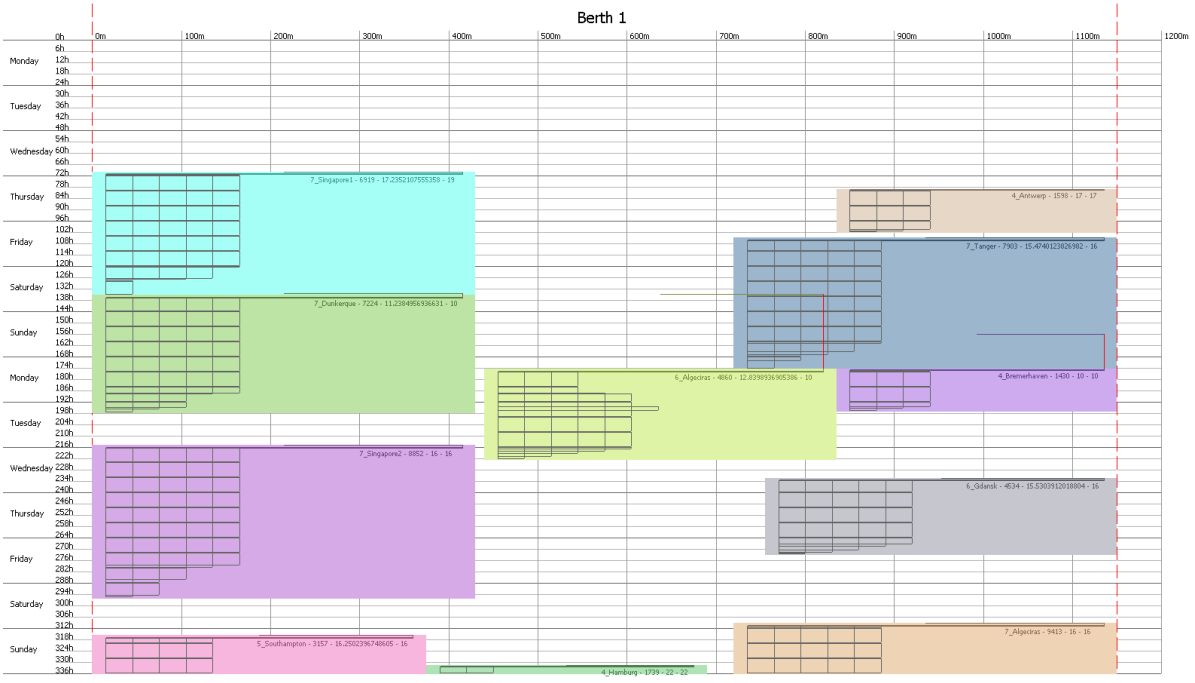


Figure C.4: Realised berth schedule of week 1 and 2 in scenario 1, configuration global speed optimisation

C.1.2. Weeks 3 and 4

For weeks 3 and 4, the three main vessels that are discussed are the green TEU 4 vessel from Hamburg, the light green TEU 7 vessel from Dunkerque, and the grey-blue TEU 7 vessel from Tanger. In the baseline configuration, in Figure C.5, the vessel from Hamburg has to wait for a long time. While the vessels from Dunkerque do not have to wait, the vessel from Tanger has to wait for a bit.

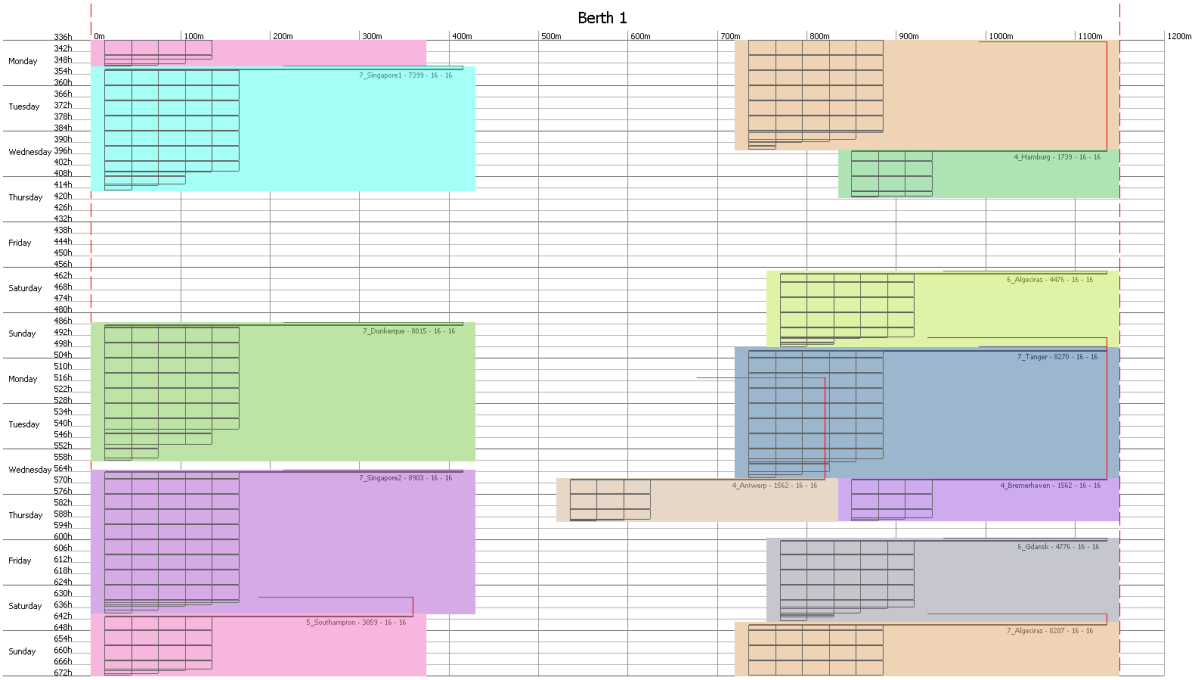


Figure C.5: Realised berth schedule of week 3 and 4 in scenario 1, configuration baseline

In the single-speed optimisation configuration, displayed in Figure C.6, the vessel from Hamburg still has to wait; however, it has set sail at a lower velocity, decreasing the waiting time. And the vessels from Tanger and Dunkerque have switched around, causing the vessel from Dunkerque to set sail at a lower velocity and wait. The switch happened as the vessel from Tanger had been assigned a berth before the vessel from Dunkerque, due to departing earlier.

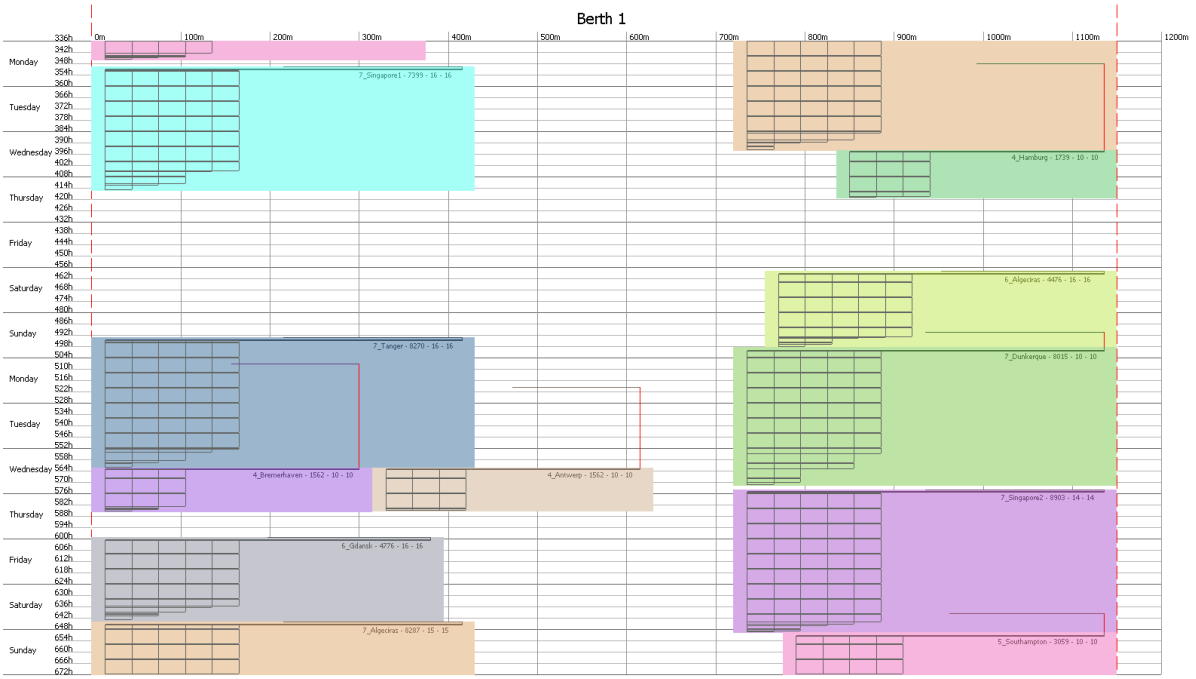


Figure C.6: Realised berth schedule of week 3 and 4 in scenario 1, configuration single speed optimisation

In Figure C.7, the vessel departing from Hamburg no longer has to wait. This does cause the TEU 7 vessel from Singapore to sail at a lower velocity. The Vessels from Dunkerque and Tanger have switched again to the same locations as in the baseline scenario. However, in this case, the vessel on the right does not have to wait. The switch happened as the vessel from Dunkerque at departure had fewer miles to travel than the vessel from Tanger.

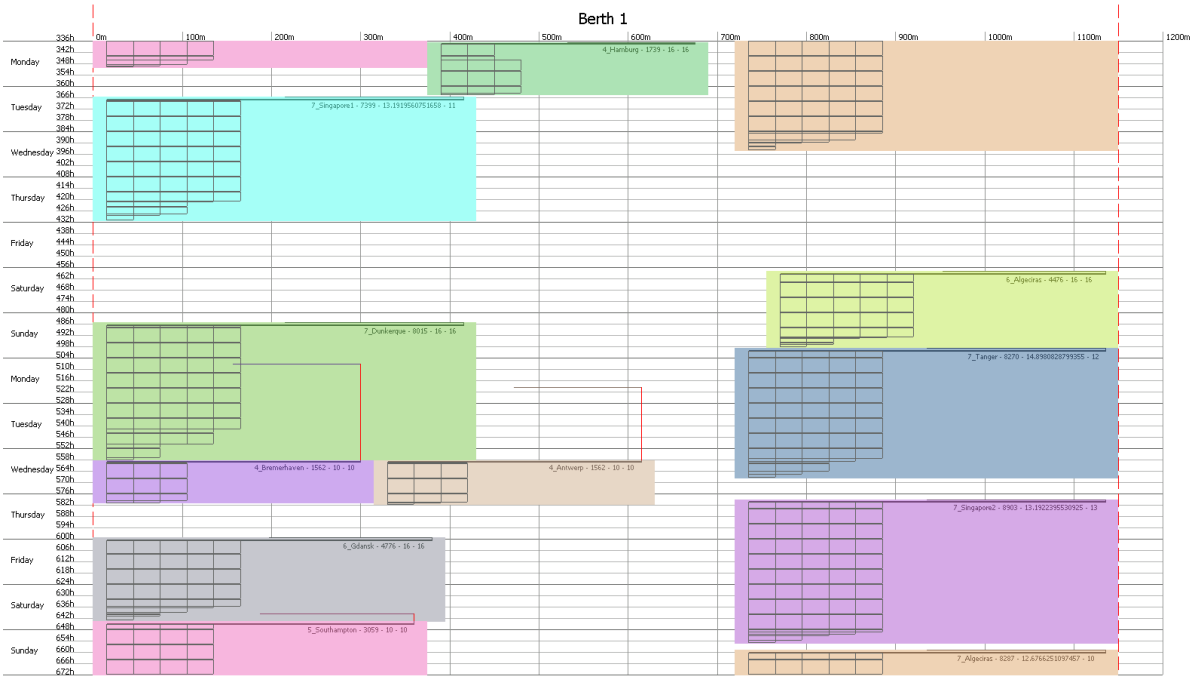


Figure C.7: Realised berth schedule of week 3 and 4 in scenario 1, configuration re-optimise at new departures

The realised berth schedule of global speed optimisation in Figure C.8 again looks similar to that of re-optimise at new departures. Main differences are that the vessel from Hamburg now travels at its maximum velocity, to ensure that the vessel from Singapore can sail at a velocity closer to its optimum. Furthermore, the gaps in time between the vessels are smaller, due to the velocity being updated every 2 hours.

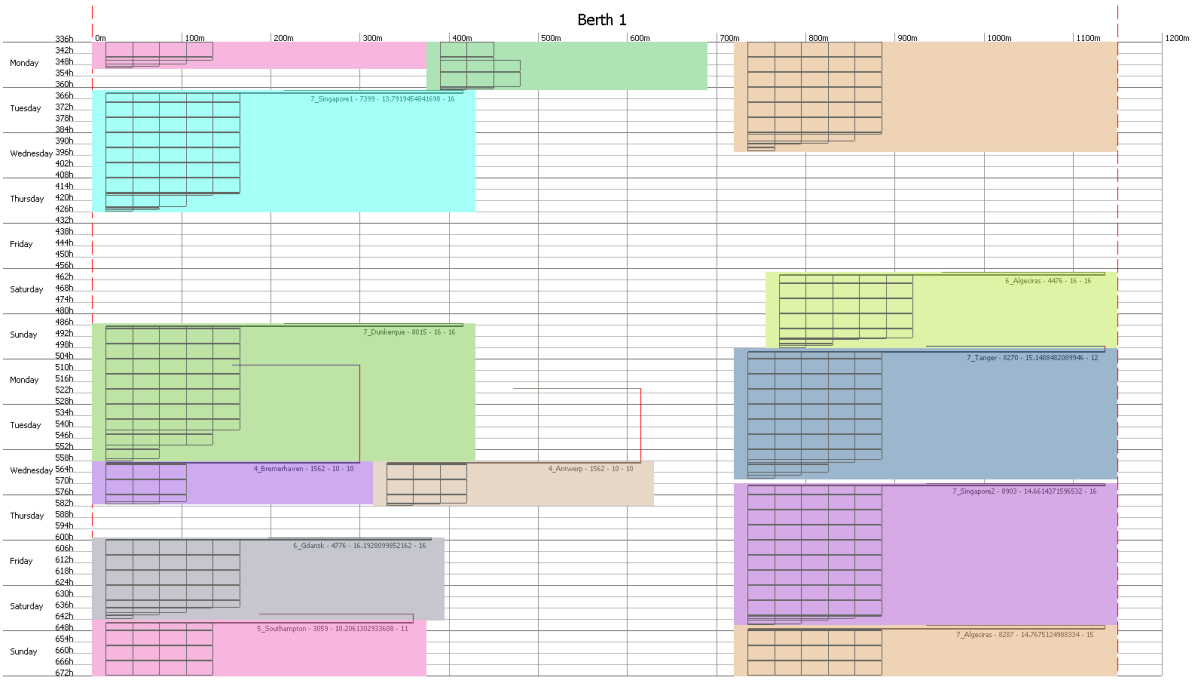


Figure C.8: Realised berth schedule of week 3 and 4 in scenario 1, configuration global speed optimisation

C.1.3. Conclusion

The realised berth schedules, displayed in Figure C.1 to Figure C.8, demonstrate that the speed optimisation configurations operate according to their design.

In the baseline configuration, all vessels set sail at their economical velocity of 16 NM/h, and are served according to the FCFS priority-based berth allocation logic. Therefore, the strategy serves as a benchmark for the other strategies.

In the single-speed optimisation configuration, the vessels' velocities are updated only at departure. As expected, this results in lower speed when the waiting time is anticipated, leading to shorter anchorage waiting time. However, as the vessels cannot react to updated berth information, either from newly departing vessels or from the estimated end of service, this leads to avoidable idle time at the terminal.

In the re-optimise at new departures configuration, the velocity of vessels is updated whenever a new vessel departs with a lower remaining distance. The variations in schedule from both the baseline and the single-speed optimisation confirm that the vessels react to the new information at the departures, and that the velocity is updated accordingly.

In the global speed optimisation configuration, the optimisation is executed periodically for all sailing vessels. The realised berth schedule shows that vessels occasionally sail above or below their reference speed to improve the overall system's performance. This is consistent with the objective function, which minimises the total system costs, rather than the costs of individual vessels.

Therefore, these realised berth schedules show that the decision logic of all four configurations has been implemented correctly and that they behave according to their intended design.

C.2. Single-speed optimisation

Another output report that is inspected is the output report at the moment of the speed optimisation. Below are three of the single-speed optimisations at the moment of departure. In Figure C.9, the expected costs of a vessel are displayed without an expected delay. This cost function is shown for a vessel (A TEU 7 vessel, departing from Singapore with 1200 NM to go, with a c_f of 555.5 dollars per

metric tonne, a $\tilde{f}$ of 200 tonne per day, and a $\tilde{v}$ of 16). When it does not experience any waiting, the terminal will select the velocity with the lowest cost, which in this case is 16 NM/h. As no waiting is involved in this case, this is equal to $\tilde{v}$. In Figure C.10, a cost function is shown for the same vessel; this vessel would experience a delay when travelling at a velocity larger than or equal to 15 NM/h. In this case, the velocity with the lowest cost, 14 NM/h, is selected.

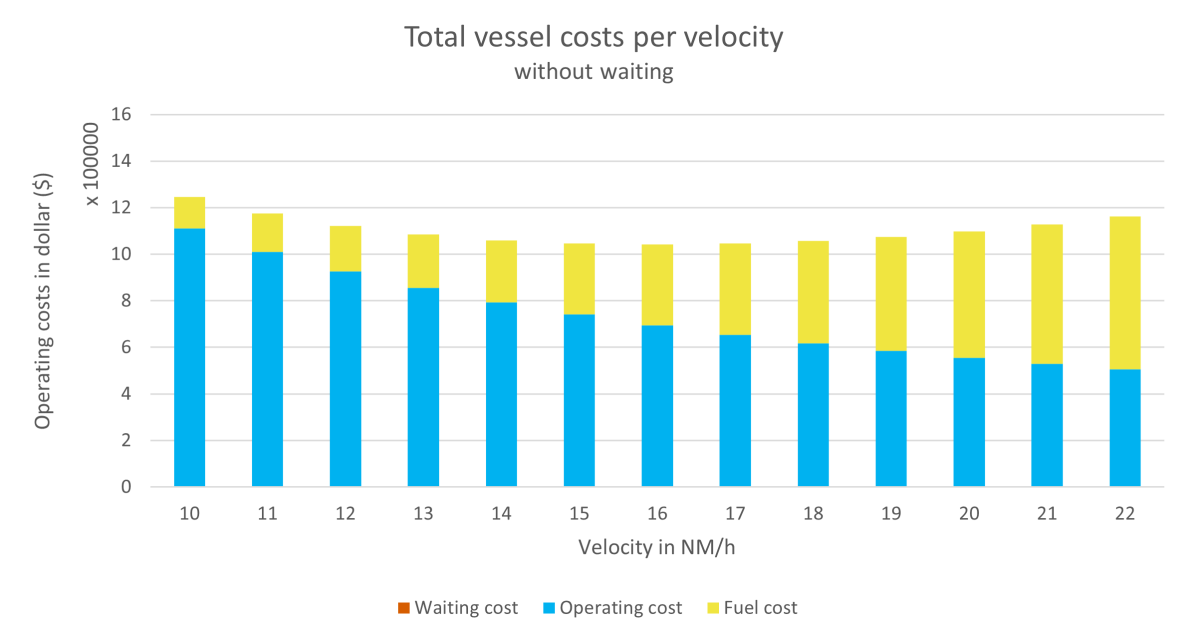


Figure C.9: Cost function without expected delay

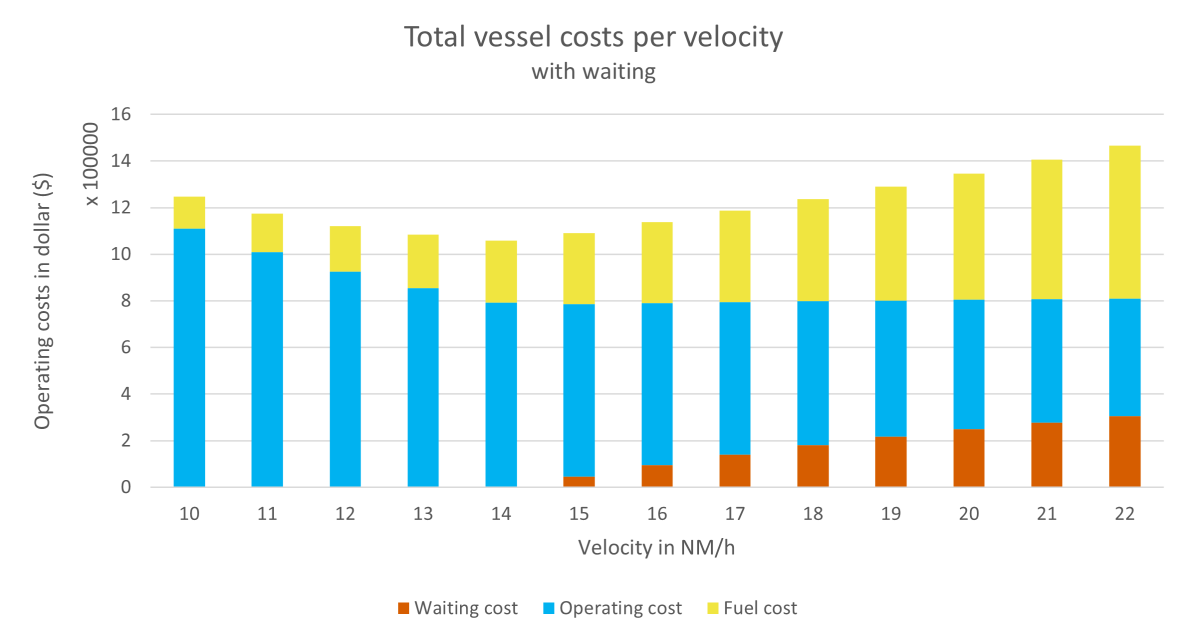


Figure C.10: Cost function with expected delay

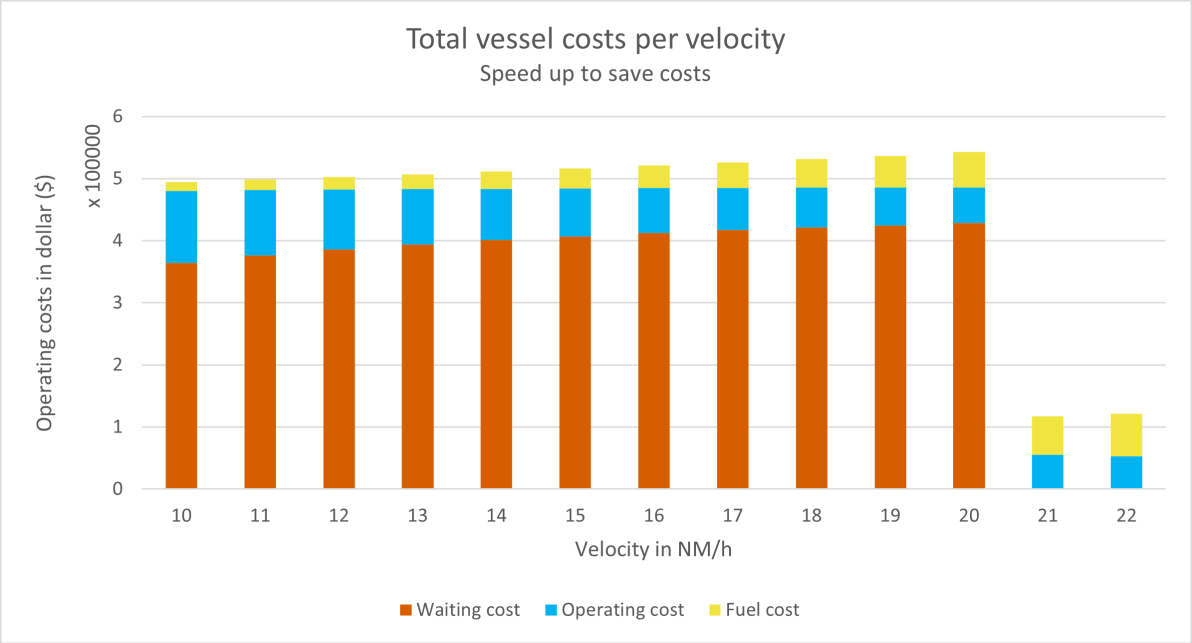


Figure C.11: Cost function in which the lowest cost would be achieved by speeding up

In Figure C.11, the expected cost function is displayed for a vessel (A TEU 5 vessel, sailing from Southampton, with 250 NM to go, with a c_f of 555.5 dollars per metric tonne, a $\tilde{f}$ of 100 tonne per day, and a $\tilde{v}$ of 16). In this cost function, a rare situation occurs: the vessel's lowest cost is when it speeds up to 21 NM/h. This is the case because the berth slots behind the vessel are already occupied by a vessel departing from farther away.

In all three cases, the vessel would select the velocity with the lowest costs, which is $\tilde{v}$ for Figure C.9, which is 14 NM/h for Figure C.10, to avoid waiting at the quay, and 21 NM/h for Figure C.11, indicating a necessity to speed up to save costs. Indicating that the speed optimisation function works correctly, and chooses the velocity with the lowest costs.

C.3. Global speed optimisation

Based on the scenario in week 1, an example shows how the speed-up function selects the best arrival order and the best index for speeding up.

C.3.1. Current sailing vessels

In Figure C.12, all the vessels that have currently departed their previous port and are sailing on route to Rotterdam World Gateway are displayed. In this table, information about the vessel is displayed. For instance, the call size, the priority (which is in global speed optimisation 0 for every vessel), the length, the average velocity, the current time, and the distance to the port. In this case, the vessel from Singapore in Figure C.4 has already berthed, but the vessel from Dunkerque has not yet departed. For these three vessels, the optimal berthing order and velocity need to be determined.

	object 0	string 1	integer 2	integer 3	real 4	integer 5	time 6	real 7	table 8	string 9
string	Vessel	Type	Callsize	Priority	Length	AverageVelocit	DepartureTime	Distance	TimeOfArrival	VesselID
1	.BAM.Vessel:4	4_Antwerp	1598	0	285	16	2:22:36:17.979	140.00	table81	Vessel_28
2	.BAM.Vessel:2	7_Tanger	7903	0	400	16	2:22:36:17.979	505.23	table82	Vessel_55
3	.BAM.Vessel:3	6_Algeciras	4860	0	365	16	2:22:36:17.979	815.19	table83	Vessel_82

Figure C.12: All sailing vessels of which the velocities have to be re-optimised

Based on the information in Figure C.12, the table shown in Figure C.13 is constructed. For each velocity, fuel and operating costs are determined based on the vessel's parameters and the distance

to the port. The cost of arrival, which is the total cost, is not yet filled in, as it is dependent on the waiting time, which is dependent on the order of berthing.

	real 0	time 1	real 2	boolean 3	real 4	real 5	real 6	real 7
string	Velocity	Time Of Arrival	Cost Of Arrival	FoundBerth	ExpectedWaitin	Waiting Cost	Fuel Cost	Operating Cost
1	10.00	3:12:36:17.979					5933.38	48606.25
2	11.00	3:11:19:56.161					7179.39	44187.50
3	12.00	3:10:16:17.979					8544.07	40505.21
4	13.00	3:09:22:27.210					10027.41	37389.42
5	14.00	3:08:36:17.979					11629.43	34718.75
6	15.00	3:07:56:17.979					13350.11	32404.17
7	16.00	3:07:21:17.979					15189.45	30378.91
8	17.00	3:06:50:25.038					17147.47	28591.91
9	18.00	3:06:22:57.979					19224.15	27003.47
10	19.00	3:05:58:24.295					21419.50	25582.24
11	20.00	3:05:36:17.979					23733.52	24303.13
12	21.00	3:05:16:17.979					26166.21	23145.83
13	22.00	3:04:58:07.070					28717.56	22093.75

Figure C.13: Arrival list with all possible times of arrival, fuel cost and operating cost for the TEU 4 vessel departing from Antwerp

C.3.2. Select order

First, all the different orders of arrival are evaluated. This is done using the tables shown in Figure C.14. First, all orders are generated, ranging from 1,2,3 to 3,2,1. These numbers correlate with the indices of vessels in Figure C.12. Based on the indices, the table in Figure C.13 is filled in first for the first vessel, and the velocity with the lowest costs is selected. Based on the new situation, the table is filled in for the second vessel until all vessels have been placed into the berth planning. For each of the orders, the costs are added together and placed into the TotalCost. The order with the lowest total cost is selected, which, in the case of Figure C.14, is the first order.

	integer 0	table 1	table 2	real 3
string	Number	Order	IndividualCo	TotalCost
1	1	table11	table21	970190.85
2	2	table12	table22	1189187.64
3	3	table13	table23	1174753.98
4	4	table14	table24	1174753.98
5	5	table15	table25	1189187.64
6	6	table16	table26	1189187.64
7				

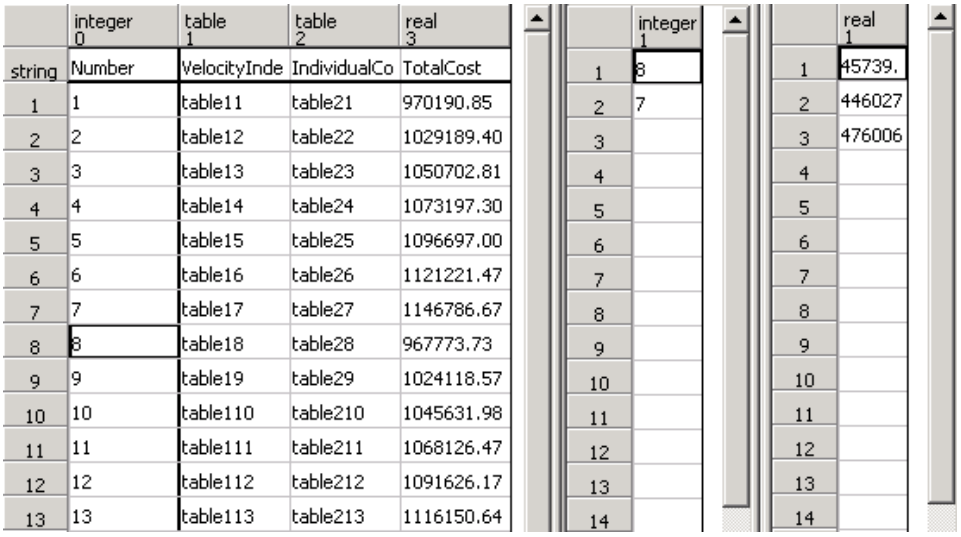
	integer 1
1	1
2	2
3	3
4	
5	
6	
7	
8	

	real 1
1	45568.36
2	448615.99
3	476006.50
4	
5	
6	
7	
8	

Figure C.14: Evaluation of all different orders of arrival of sailing vessels, from left to right: the table in which all different orders are evaluated, the order of the table with the lowest cost, the cost of individual vessels of the order with the lowest cost

C.3.3. Select velocity

Based on the order, it is investigated whether the total costs could be lower if the other vessels, except the last one, speed up. The evaluation is displayed in Figure C.15. In this case, if the first vessel were to speed up to index 8 from Figure C.13, the total cost would be the lowest. So the first vessel in the order would set sail at 17 NM/h. The second and third vessels would then select velocities for which their own costs are the lowest.



	integer 0	table 1	table 2	real 3
string	Number	VelocityInde	IndividualCo	TotalCost
1	1	table11	table21	970190.85
2	2	table12	table22	1029189.40
3	3	table13	table23	1050702.81
4	4	table14	table24	1073197.30
5	5	table15	table25	1096697.00
6	6	table16	table26	1121221.47
7	7	table17	table27	1146786.67
8	8	table18	table28	967773.73
9	9	table19	table29	1024118.57
10	10	table110	table210	1045631.98
11	11	table111	table211	1068126.47
12	12	table112	table212	1091626.17
13	13	table113	table213	1116150.64

	integer 1
1	8
2	7
3	
4	
5	
6	
7	
8	
9	
10	
11	
12	
13	
14	

	real 1
1	45739.
2	446027
3	476006
4	
5	
6	
7	
8	
9	
10	
11	
12	
13	
14	

Figure C.15: Evaluation of all different speed-up possibilities within the order with the lowest costs, from left to right: the table in which all different speed-ups are evaluated, the indices of the vessels' velocity with the lowest cost, and the cost of the velocity of individual vessels of the combination with the lowest costs

C.3.4. Conclusion

The example demonstrates that the global speed optimisation procedure evaluates all feasible arrival orders and determines the velocity combination within this order that has the lowest total system costs.

C.4. Berth occupancy

In this section, the results for the berth occupancy are verified. First, the deviation in the number of vessels arriving within a year is explained. And later, the calculations for the increase in berth occupancy are verified by calculating the expected increase in berth occupancy.

C.4.1. Number of vessels handled per simulation

From Figure 7.9, it becomes clear that the berth occupancy for different configurations within scenarios is not exactly the same. This can be explained using Table C.1. In the table, vessels from vessel class 4 arrive between 25.6 and 25.9 times. 25.6 times indicates that in 4 simulations, it has arrived 25 times, while in the other 6 simulations, it has arrived 26 times. This concerns the realisation of stochastic parameters on the arrival time of vessels. As in Table 6.1, it is stated that the TEU 4 vessel from Hamburg arrives on day 14, the latest day of the biweekly schedule. A year consists of 365 days, which consists of 26 periods of 14 days, and 1 extra day. So if the vessel from Antwerp had been delayed for more than a day, it would not have started its service before the end of the simulation. On the other hand, the TEU 7 vessel from Singapore arrives between 26.6 and 26.9 times per year. This vessel arrives on day 1 of the biweekly schedule. Thus, if the vessel is not delayed, it will arrive in time to be serviced before the end of the simulation.

The berth occupancy is calculated using Equation 5.1. If in scenario 1, configuration baseline, 26 vessels from Hamburg would have arrived, and 27 vessels from Singapore, the berth occupancy would have been increased from 61.4 % to 61.6%.

Vessel class	Scenario / Configuration / Previous port	Scenario 1 Baseline	Scenario 1 Single-speed	Scenario 1 Re-optimise	Scenario 1 Global	Scenario 2 Global	Scenario 3 Global
4	**Antwerp**	26	26	26	26	26	26
4	**Bremerhaven**	26	26	26	26	26	26
4	**Hamburg**	25.6	25.8	25.9	25.8	25.7	25.9
4	**Le Havre**	0	0	0	0	26.8	26.7
5	**Antwerp**	0	0	0	0	0	26
5	**Southampton**	26	26	26	25.7	26	25.7
6	**Algeciras**	26	26	26	26	26	26
6	**Gdansk**	26	26	26	26	26	26
7	**Singapore**	26.6	26.8	26.7	26.9	26.7	26.6
7	**Algeciras**	26	26	26	26	26	26
7	**Singapore**	26	26	26	26	26	25.9
7	**Dunkerque**	26	26	26	26	26	26
7	**Tanger**	26	26	26	26	26	26

Table C.1: Average number of arrivals per scenario and configuration

C.4.2. Increase in berth occupancy

To verify the calculations of berth occupancy, the realised increase in berth occupancy from scenarios 2 and 3 is also calculated by hand. The variables used for the vessels in scenarios 2 and 3 are shown in Table C.2. By multiplying these variables and dividing them by $1150 \cdot 24 \cdot 365$, the length of the simulation times the length of the quay, the quay occupancy would increase by 2.13 % due to the biweekly arrival of the vessel from Le Havre and by 3.60 % due to the biweekly arrival of the vessel from Antwerp, together leading to the increase of 5.7%. The 2.14 % is slightly lower than the 2.2 % difference in the results; this can be explained by the varying number of arriving vessels, as shown in subsection C.4.1, where the berth occupancy can vary at the decimal point, due to the arrival of an additional vessel.

Vessel Class	Previous Port	Total time on quay	Vessel length + mooring marging	Number of arrivals
4	Le Havre	25.64	315	26.6
5	Antwerp	37.22	375	26

Table C.2: Additional vessels in scenarios 2 and 3

The close agreement between the analytically calculated berth occupancy increase and the simulation results confirms that the berth occupancy KPI is calculated correctly within the simulation model.