

Stochastic In-Port Routing in Chemical Shipping

Aurora Smith Elgesem Eline Sophie Skogen

Industrial Economics and Technology Management Submission date: June 2016 Supervisor: Kjetil Fagerholt, IØT Co-supervisor: Xin Wang, IØT

Norwegian University of Science and Technology Department of Industrial Economics and Technology Management

Description of the Thesis

The purpose of this thesis is to investigate the benefits of including uncertainty in the in-port routing problem of a chemical tanker. The problem entails finding a sequence to service a given number of cargoes located in several different terminals. Each cargo will either be picked up or delivered. The ship must service all cargoes while complying with capacity and draft limit constraints. The case port is particularly long and narrow, and the movement of the ship can be considered as movement along a straight line. Accounting for uncertainty in the problem implies modelling the waiting times at the terminals as stochastic variables. Both a static and a dynamic version of the problem are solved. The objective is to find the route that maximizes the probability of route completion within a given threshold. The static stochastic version is solved using an exact solution algorithm presented by Nikolova et al. (2006), and the dynamic stochastic version is solved by simulation, where the exact solution algorithm is used in each stage to decide the next cargo to service. Various conditions have different effects on the uncertainty pertaining to a route, and these conditions and effects are examined. Analysis and simulations are used to evaluate the performance of the model, the stochastic solutions, and the applied approximations.

Preface

This thesis is written at the Norwegian University of Science and Technology, NTNU, in Trondheim. It is the final work of our Master's degree in Industrial Economics and Technology Management. The work is done in the field of Managerial Economics and Operation Research. The thesis is written during the spring semester of 2016 and is based on the project thesis written during the fall semester 2015.

The benefits of including uncertainty in the in-port ship routing problem of a chemical tanker is investigated in our work. The thesis is part of the ongoing project GREENSHIPRISK: Green Shipping Under Uncertainty, which is a collaboration between MARINTEK/NTNU, the Centre for Applied Research at NHH, Odfjell SE, Western Bulk, and Bergen Shipowners' Association. The work conducted by Mari Jevne Arnesen and Magnhild Gjestvang for their master thesis as part of the GREENSHIPRISK project has been of great support during our research.

We would like to thank our supervisor and co-supervisor, Kjetil Fagerholt and Xin Wang, for their sincere enthusiasm, guidance and valuable input.

We also wish to thank representatives from Odfjell SE for their contributions to the input data.

Trondheim, June 2016

Algesem

Aurora Smith Elgesem

Elin, flig.

Eline Sophie Skogen

Abstract

Chemical tankers spend a substantial amount of time in port as the development of port infrastructure has not followed the fast paced increase in the world fleet. The resulting traffic at terminals causes ships to wait for a long time before a terminal is ready to accommodate it. A large amount of uncertainty is associated with the waiting times, which complicates the planning of port operations.

The aim of the thesis is to investigate the benefits of including uncertainty in the in-port routing problem of a chemical tanker. The chemical tanker has to pick up and deliver a given number of cargoes located at different terminals while complying with capacity and draft limit constraints. The waiting times at the terminals are stochastic, which results in stochastic travel times between terminals. The problem studied in this thesis is a stochastic pickup and delivery problem. The problem is dynamic by nature, and both a static and a dynamic version of the problem are solved. A review of deterministic and stochastic routing problems resembling the pickup and delivery problem is presented. To our knowledge, few or none have studied the pickup and delivery problem where travel times are stochastic. Our thesis contributes to the literature by studying the pickup and delivery problem with uncertain travel times, subject to constraining draft limits. In addition, due to the particularly narrow port channel in the case port, we consider the ship's movement as movement along a straight line. This gives a unique relation between stochastic waiting times at terminals and travel times between terminals. As far as we know, this has not been studied before, and the unique conditions and aspects of the travel times this gives are examined and discussed.

The stochastic waiting times at terminals are assumed normally distributed. The stochastic travel times between cargoes in different terminals and the stochastic waiting time at the destination terminal are correlated. The distribution of the arc travel times are assumed normally distributed, and approximations are used to obtain the distributions.

Both a static and a dynamic version of the problem are solved. The objective

is to find the route that maximizes the probability of completing within a given threshold. As the objective function is non-linear fractional, which is not straightforward to handle, specialized solution methods are used. An exact algorithm presented by Nikolova et al. (2006) is used to solve both versions of the problem. What confidence level is required for a route to be optimal depends on the risk profile of the decision maker. When solving the static version of the problem, the optimal route for the given threshold is identified prior to route execution, while the dynamic problem is solved using the exact algorithm iteratively to decide which cargo to service next.

When solving the static version of the problem, a base set of 100 test instances is generated and tested. The instances are generated based on realistic input data from Houston Ship Channel. We find that for less than 20% of the instances, the optimal stochastic solution performs better than the optimal deterministic. However, the improvements in threshold and confidence level is less than 0.5% for all instances.

An evaluation of the applied approximation of the distribution of arc travel times shows that our model suggests less variance to be associated with the routes than what is the real case. This means that the value of the stochastic solution might be higher than what our results suggest. The results from solving the dynamic version of the problem support the findings from solving the static version of the problem.

Sammendrag

Utviklingen av infrastukturen i havner har ikke holdt følge med den raske veksten i verdensflåten. Kjemikalietankere bruker derfor en betydelig andel av tiden sin i havn. Dette gir trafikk ved terminalene, og medfører at skip ofte må vente lenge før en terminal er klar til å ta imot skipet. Det er usikkerhet assosiert med ventetidene. Dette kompliserer planleggingen av havneoperasjoner.

Målet med vår avhandling er å undersøke nytten av å inkludre usikkerhet i problemet med å finne optimal rute i havn for en kjemikalietanker. Kjemikalietankeren må hente og levere et gitt antall laster i ulike terminaler samtidig som den overholder begrensninger for skipets kapasitet og dypgang. Ventetidene ved terminalene er stokastiske. Dette gir stokastiske reisetider mellom terminalene. Problemet vi studerer i denne avhandlingen er dermed et stokastisk pickup and delivery-problem. Problemet er dynamisk av natur, og både en dynamisk og en statisk versjon av problemet er løst. Et litteraturstudie av dynamiske og stokastiske ruteproblemer som ligner på pickup and deliveryproblemer er inkludert. Så vidt vi vet har ingen, eventuelt svært få, studert pickup and delivery-problemer med stokastiske reisetider. Denne avhandlingen bidrar derfor til litteraturen ved å studere et pickup and delivery-problem med usikre reisetider begrenset av kapasitet og dypgang. På grunn av den spesielt lange og smale havnekanalen som brukes som eksempel i denne avhandlingen anser vi bevegelsene til skipet som bevegelser langs en rett linje. Dette gir et unikt forhold mellom den stokastiske ventetiden ved terminalene og reisetiden mellom terminalene. Så vidt vi vet har ikke dette blitt studert tidligere. De unike forholdene og aspektene ved resitidene dette medfører er undersøkt og diskutert.

De stokastiske ventetidene ved terminalene er antatt å være normalfordelt. De stokastiske reisetidene mellom laster i ulike terminaler og den stokastiske ventetiden er korrelert. Distribusjonene til reisetidene langs kanter er antatt å være normalfordelt, og approksimasjoner er brukt for å konstruere dem.

Både en statisk og en dynamisk versjon av problemet er løst. Formålet er

å finne den ruten som maksimerer sannsynligheten for å fullføre ruten innen en gitt tid. Det er utfordrende å håndtere den ikke-lineære fraksjonelle objektivfunksjonen, og en spesialtilpasset løsningsmetode er nødvendig. En eksakt algoritme, presentert av Nikolova et al. (2006), er brukt for å løse både den statiske og den dynamiske versjonen av problemet. Hvilken rute som anses for å være optimal avhenger av kravet til konfidensnivå. Hvilket konfidensnivå som anvendes avhenger av risikoprofilen til beslutningstakeren. For den statiske versjonen av problemet blir den optimale ruten identifisert før ruten er påbegynt. Det dynamiske problemet løses iterativt med den eksakte løsningsalgoritmen. I hver iterasjon blir den optimale neste lasten å betjene identifisert.

Vi genererer et basesett med 100 testinstanser. Disse instansene utgjør grunnlaget for testing og analyse av den statiske versjonen av problemet. Baseinstansene er generert ved å bruke realistisk data fra eksempelhavnen, Houston Ship Channel. Vi finner at for mindre enn 20% av baseinstansene er den optimale løsningen på det stokastiske problemet ulik fra den optimal determiniske løsningen. Men forbedringen i tidskravet man oppnår ved å bruke den optimale stokastiske løsning fremfor den optimale determiniske løsningen er mindre enn 0.5% for alle instanser.

Approksimasjonen brukt for å konstruere distribusjonene for reisetider langs kantene blir evaluert. Evalueringen viser at vår modell foreslår mindre varianse assosiert med hver kant enn hva som er tilfellet i virkeligheten. Dette innebærer at verdien av å bruke den stokastiske løsningen *kan* være større enn det resultatene våre tilsier. Betydningen av resultatene vi får ved å løse den dynamiske versjonen av problemet samsvarer med betydningen av resultatene vi får når vi løser den statiske versjonen av problemet.

Contents

1	Intr	roduction	1				
	1.1	Background	1				
		1.1.1 Maritime Transportation	1				
		1.1.2 The Chemical Tanker Market	3				
	1.2	The Thesis Problem	6				
	1.3	Upcoming Chapters	7				
2	Problem Description						
	2.1	The Stochastic Problem	9				
	2.2	Implications of Including Uncertainty	14				
	2.3	Summary	17				
3	Lite	erature Review	19				
	3.1	Pickup and Delivery Problems	19				
		3.1.1 Classification	20				
		3.1.2 General Notation	22				
		3.1.3 Solution Methods	22				
	3.2	Stochastic Routing Problems	24				
		3.2.1 Definition of Optimality	24				
		3.2.2 Modelling Approaches	25				
		3.2.3 Travelling Salesman Problems	26				
		3.2.4 Vehicle Routing Problems	27				
		3.2.5 Shortest Path Problems	29				
		3.2.6 Distribution of Stochastic Travel Times	29				
	3.3	Static and Dynamic Routing Problems	30				
	3.4	Contribution to Literature	32				
4	Ma	thematical Formulation	33				
	4.1	Modelling Assumptions	33				
	4.2	Modelling Notation	35				
	4.3	The Mathematical Model Formulation	36				
	4.4	Preprocessing	39				
5	Solı	ation Methods	43				

	$5.1 \\ 5.2$	Nikolova's Method	43 48
C	0.2 D		-10 F-1
0	Pro	Dapling Distributions of the Stochastic Variables	51
	0.1 6 9	Distribution of the Stochastic Sailing Times	02 E4
	0.2	6.2.1 The Dile Up Effect	55
		6.2.2 The Anchorage Effect	56
		6.2.3 Simulation of Arc Distributions	58
		6.2.4 The Pile-Up Effect and Anchorage Effect for Various	00
		Port Characteristics	60
	6.3	Summary	65
7	Inp	ut Data and Generation of Test Instances	67
	7.1^{-1}	Input Data	67
	7.2	Summary of the Test Instances	74
8	Con	nputational Study: Static Version of the Problem	77
	8.1	Model Output and Interpretation	78
	8.2	Aggregated Analysis of the Base Set	80
	8.3	Detailed Analysis	82
		8.3.1 Comparing Routes for Instance 46 from the Base Set	83
		8.3.2 Comparing Routes for Instance 1 from the Base Set	89
	8.4	Sensitivity Analysis	93
		8.4.1 Geography	94
		8.4.2 Draft Limits	98
	0 E	8.4.3 Risk Profile of the Decision Maker	99
	0.0 9.6	Summary of Reasons to Different variances	100
	$8.0 \\ 8.7$	Practical Implications	102
0	Com	way to the Demonstration of the Duchlam	100
9	0.1	Modifications	110
	9.1	9.1.1 The Stochastic Planning Strategy	111
		9.1.2 The Deterministic Planning Strategy	112
	9.2	Results	113
	0.2	9.2.1 Instance 46 from the Base Set	114
		9.2.2 Instance 26 from Variation Set 2 with Split Geography.	115
		9.2.3 Instance 5 from the Base Set	116
	9.3	Practical Implications	118
10) Con	cluding Remarks	119
	10.1	Summary and Conclusion	120
	10.2	Future Research	122
B	ibliog	graphy	122

A Model Formulation

127

List of Figures

1.1	The OECD Industrial Production Index and indices for world GDP, merchandise trade and seaborne shipments, 1975-2014 (UNCTAD, 2015).	2
1.2	World seaborne trade in cargo ton-miles by cargo type, 2000-2015 (UNCTAD, 2015).	3
2.1	Houston Ship Channel in the Port of Houston.	10
2.2 2.3	Additional sailing distance due to waiting times	12
2.0	al., 2012).	13
2.4	An example of two feasible routes with different means and variances of route completion times	15
3.1	Classification of PDPs by Berbeglia et al. (2007)	20
4.1	Components of the travel time between two customers in dif-	94
4.2	The probability of the completion time, T , being lower than the threshold, H .	34 37
4.3	Arcs between delivery nodes in the same terminal to be removed to reduce symmetry	41
4.4	Arcs between pickup nodes in the same terminal to be removed to reduce symmetry.	41
5.1	Projection of the unit hypercube and the path polytope in the (μ, σ^2) -plane (Nikolova et al., 2006).	44
5.2	Points with lower z-scores than a given z-score, γ , for a given deadline, t_1 , \dots	46
5.3	Feasible combinations of mean and variance which do not exist.	47
5.4	Finding and enumerating the candidate solution-paths (Nikolova,	
	$2009). \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots $	48
5.5	The order for discovering routes and using edges between adja- cent nodes to update the weight of the objective function and	
	search for new routes.	49

5.6	The upper and lower limits for the range of thresholds for witch the confidence level for all identified routes are calculated for a given instance
6.1	Sailing time and waiting time dependency for arcs directed away from anchorage
6.2	Sailing time and waiting time dependency for arcs directed to- wards anchorage
6.3	Possible outcomes for arc (B, A) , which is directed towards an- chorage
6.4	Possible outcomes for arc (A, B) , which is directed away from anchorage
6.5	Characteristics of the port used to simulate and explain the distribution of arc sailing times
6.6	Distributions of realised sailing times for two symmetric arcs given realised waiting times at the destination terminals 59
6.7	Characteristics of the port used to simulate the Pile-Up and Anchorage Effect
7.1	Terminals in Houston Ship Channel used in the test instances (Arnesen and Gjestvang, 2015)
7.2	Terminals in Houston Ship Channel the way they are modelled in the test instances
7.3 7.4	The three different geographies, Even, Far Away, and Split 69 Distribution of the realised waiting times for a collection of terminals at Houston Ship Channel, based on a set of data received from Odfjell SE
8.1	Example of model output for a test instance from the Base Set. 79
8.2	Details of Instance 46 from the Base Set
8.3	Model output for Instance 46 from the Base Set
8.4	Cargo and terminal visiting sequence of the optimal determin- istic route, (Route 1), and the optimal stochastic route (Route 3) of Instance 46 from the Base Set
8.5	Details of Instance 1 from the Base Set
8.6	Model output for Instance 1 from the Base Set 90
8.7	Cargo and terminal visiting sequence of the optimal determin- istic route (Route 1) and the optimal stochastic route (Route 4) of Instance 1 from the Base Set
8.8	Distributions of arcs (I, K) and (A, K) used for Routes 1 and 4, respectively, for Instance 1 from the Base Set. $\dots \dots \dots$
8.9	Sample distribution and distribution outputted by our model for route completion times for the optimal stochastic route, Route 3, for Instance 46

Sample distribution and distribution outputted by our model for route completion times for the optimal deterministic route,
Route 1, for Instance 46
Distribution of route completion times for the dynamic prob-
lem with deterministic and stochastic planning strategies for
Instance 46 from the Base Set
Distribution of route completion times for the dynamic prob-
lem with deterministic and stochastic planning strategies for
Instance 26 from Variation Set 2
Distribution of route completion times for the dynamic prob-
lem with deterministic and stochastic planning strategies for
Instance 5 from the Base Set

List of Tables

$6.1 \\ 6.2$	The Pile-Up Effect for arc (B, A) directed towards anchorage The Anchorage Effect for arc (A, B) directed away from an-	63
	chorage	64
7.1	Terminals from Figure 7.1 used to approximate the geographies Even, Far Away and Split.	70
7.2	Summary of input data used to generate the Base Set of test instances.	75
8.1	Aggregated results for the Base Set	82
8.2	Cargo and terminal visiting sequence for Routes 1 and 3 of Instance 46 from the Base Set.	85
8.3	Distribution of the sailing times of the arcs which differ between Boutes 1 and 3 for Instance 46 from the Base Set	87
8.4	Cargo and terminal visiting sequence for Routes 1 and 4 of	01
8.5	Instance 1 from the Base Set	90
86	Routes 1 and 4 for Instance 1 from the Base Set	92
0.0	Away, and Split geographies.	94
8.7	Aggregated results from the set of test instances with a shorter distance between anchorage and the terminal closest to anchorage.	97
8.8	Aggregated results from the set of test instances with more binding draft limits	99
8.9	Aggregated results from the set of test instances where a lower	
8.10	Confidence level is require for routes to be optimal Characteristics of the sample set of realised route completion	100
	times and characteristics outputted by the model for the opti- mal stochastic solution. Boute 3 for Instance 46	03
8.11	Characteristics of the sample set of realised route completion	
	times and characteristics outputted by the model for the opti- mal deterministic solution, Route 1, for Instance 46 1	104

Chapter 1

Introduction

The work presented in this thesis is motivated by a real-life problem faced by the chemical shipping company Odfjell SE. Odfjell's chemical tankers spend a substantial amount of their operating time in port. Much of the time is spent waiting for terminals to become available. The unknown number of ships queueing for the different terminals and the possibility of unscheduled surveys and tank cleanings give rise to uncertain waiting times. Improving the sequence of terminal visits to handle the uncertainty in an appropriate manner can give better control of the time spent in port and thus the associated costs and environmental effects. A case from Houston Ship Channel by Galveston Bay in the Port of Houston is used to explain in more detail the problem experienced by Odfjell.

In the following sections we highlight aspects of maritime transportation in general, and chemical shipping in particular, that are relevant for this thesis. The material in this chapter is largely based on Odfjell (2003), Odfjell (2014), and the master thesis by Arnesen and Gjestvang (2015).

1.1 Background

1.1.1 Maritime Transportation

International trade heavily depends on maritime transportation. In 2014, 9.84 billion tons of goods were transported at sea, which means that around 80% of total world merchandise trade is seaborne (UNCTAD, 2015). Demands for maritime transport services and seaborne trade volumes are shaped by global economic growth. Thus, population growth, increasing standards of

living, rapid industrialization, exhaustion of local resources, road congestion, and elimination of trade barriers all contribute to the continuing growth in maritime transportation (Christiansen et al., 2007). Preliminary UNCTAD estimates indicate that global seaborne shipment has increased by 3.4% in 2014, which is similar to the increase in 2013. Figure 1.1 highlights the association between economic growth, industrial activity, merchandise trade, and seaborne shipments. Industrial activity is here measured by the Organization for Economic Cooperation and Development's (OECD) Industrial Production Index. The development of the associated port infrastructure does, however, not follow the same pace.



Figure 1.1: The OECD Industrial Production Index and indices for world GDP, merchandise trade and seaborne shipments, 1975-2014 (base year 1990 = 100) (UNC-TAD, 2015).

Strategical, Tactical, and Operational Planning

Planning problems in maritime transportation can be classified into strategical, tactical, and operational problems depending on the planning horizon (Christiansen et al., 2007). Problems at the strategic planning level have a time horizon of more than a year, and typical problems could be market and trade selection, fleet size and mix decisions, or ship design. Tactical planning problems have a time horizon of one week to one year and include decisions like fleet deployment, and ship routing and scheduling. At the operational level, decisions for the next day or week are addressed. This can be decisions on the speed for a sailing leg, the allocation of products to compartments, the next customer request to serve, or decisions regarding whether to take spot loads or not.

Shipping Modes

Shipping companies can operate its fleet in different *modes*; tramp, liner, or industrial shipping (Lawrence, 1972). Vessels in liner shipping follow a fixed route according to a public schedule. The object of a shipping company operating its vessels in a liner mode is to maximize profits. Vessels in tramp shipping, on the other hand, trade on the spot marked with no fixed schedule. They operate as taxis and the routes are decided based on the requirements to the chartered cargo. The objective of tramp shipping is also to maximize profits. In industrial shipping, the shipping company owns both the ships *and* the cargoes they transport. The objective of the company is to transport all their cargo while minimizing costs. A shipping company can operate its fleet in multiple modes and transfer ships from one mode to another depending on its strategy.

1.1.2 The Chemical Tanker Market

The shipping market can be categorized into different segments based on the type of cargo being shipped. As seen from Figure 1.2, dry bulk, containerized cargo, and oil represent the larger part of seaborne trade. Transportation of chemicals, gas, and oil are often categorized as part of the *tanker market*, as they are transported by tanker ships.



Figure 1.2: World seaborne trade in cargo ton-miles by cargo type, 2000-2015 (billions of ton-miles) (UNCTAD, 2015).

Products

The chemical tanker market is a niche of the much bigger tanker market, and includes carriage of a range of products (Odfjell, 2003). More than 600 types of products are transported by Odfjell's chemical tankers. Odfjell's main chemical tanker cargoes are organic chemicals, inorganic chemicals, vegetable oils and animal fats, clean petroleum products and a few other assorted products. Chemical tankers are ships constructed or adapted for carrying, in bulk, any liquid product listed in the International Bulk Chemical Code (International Maritime Organization, 2014). The products have different physical properties, values and handling requirements.

Requirements and Regulations

What products are allowed onto various types of ships and tanks is regulated by strict international rules. The rules are subject to requirements to pollution prevention, vapour handling, tank cleaning, and disposal of waste water etc. In addition, there are rules and regulations pertaining to stowage, adjacent cargoes and previous cargoes. What material tanks are made from also influences what sorts of products ships may carry. Tanks made from stainless steel are most common, whilst coated tanks have various compatibility restrictions (Odfjell, 2003).

Contracts

What cargoes are transported depends on the business arrangements. These can be either contracts of affreightment (CoA) or spot cargoes. Contracts are binding long-term agreements between customers (charterers) and ship owners (such as Odfjell). Cargo space for a specified time and freight is reserved and at the customer's disposal in exchange for payment to the ship owner. The CoA does typically not specify what ship is to be used. Spot cargo contracts are determined by supply and demand. Transportation of spot cargoes is optional. Ship utilization and profits can be optimized by transporting a mix of contracted cargoes and spot cargoes.

Logistics

Within a port there are several terminals. Different customers require delivery or pickup of their cargo at different terminals, but some customer requests can be located within the same terminal. The chemical tankers have between 30 and 55 segregated compartments enabling them to transport different chemicals simultaneously. The freight usually constitutes of cargoes designated to a number of different customers, implying the need for tankers to visit a large number of terminals within a port. A terminal visit may require loading, unloading or both.

When in port, scheduling conflicts and logistics challenges often cause unnecessary transits between terminals and anchorage or layberths. These transits are time consuming. Consequently, a substantial amount of time is spent in port. Odfjell's tankers spend more than 40% of time in port. Optimizing the sequence in which terminals are visited is thus crucial to minimize the time spent in port.

Economics

The running of chemical tankers is in many respects similar to running any transportation vehicle. The key to profitability lies in the highest degree of employment with the most revenue-making business. The revenues from running chemical tankers mainly come from freight (Odfjell, 2003).

An important profitability indicator for a ship owner like Odfjell is the timecharter (T/C) result. It is used to compare the economic performance between different geographical areas, ship types, and segments. The voyage earnings in terms of T/C results are calculated as

$$T/C \text{ result} = \frac{\text{freight income - voyage costs}}{\text{voyage duration}}.$$
 (1.1)

Freight income is decided by the freight rates fixed in the CoAs. The company aims at finding a combination of CoAs that results in a high utilization of their ships. If a ship is not fully utilized, the company may include spot cargoes. The fixing of spot cargoes is a marginal consideration, where the owners have to evaluate the effect on the bottom line result of the potential income as opposed to the anticipated extra costs and time spent.

There are numerous voyage related costs, pertaining to port charges including agency fees, pilot and tug boat assistance, commissions/fees, and special tank cleaning material. By using good purchasing routines and proper voyage planning, the company tries to minimize these costs (Odfjell, 2003).

It is not only increasing freight income and reducing voyage costs that is important for the company's profits. Reducing the duration of the voyage is equally important. Thus, the key to optimal T/C results is proper fleet scheduling and planning. Making the ships carry out their voyages as fast as possible is an important contribution to good results.

1.2 The Thesis Problem

The problem studied in this thesis is motivated by the problem faced by the chemical shipping company, Odfjell SE. When visiting a port, a chemical tanker has a set of contracted cargoes to pick up or deliver at various terminals. Deciding the sequence in which the cargoes will be serviced is a key activity in the tactical planning procedure. In the chemical tanker market, time is a valuable resource. Voyage duration, including time spent in port, is a key component of a shipping company's profitability. Odfjell's chemical tankers spend more than 40% of their time in port. A large amount of time in port is spent waiting for terminals to become available. There is much uncertainty related to waiting for terminal availability, as there is an uncertain amount of traffic and queuing related to each terminal. This uncertainty is an important element when planning the sequence in which port operations will be conducted. Other important considerations to take into account are requirements and regulations, contracts and business arrangements, draft limits and ship capacity. Houston Ship Channel in The port of Houston is used as a case port to explain the situation faced by Odfjell. The particularly long and narrow port channel affects the nature of the uncertain elements, advocating the importance of taking uncertainty into account during the planning process.

The aim of the thesis is to investigate the benefits of including uncertainty in the in-port routing problem of a chemical tanker. When uncertainty is at play, it is not obvious what is regarded as an optimal route. A route may have a low expected completion time, but be subject to a large amount of uncertainty. For different geographic locations of customer requests, relative sizes of terminal waiting times and distances between terminals, different route sequences and characteristics may be favoured. To investigate the effects of these variations, both a static and a dynamic version of the problem are solved. The static model identifies the optimal route prior to route execution, while the dynamic simulation utilizes information about waiting times revealed during the execution of the route to decide what cargo to service next. In reality, information about realised waiting times are revealed during route execution, and the planning problem is dynamic by nature.

The objective is to maximize the probability of route completion within a required threshold (i.e. a deadline). By including uncertainty in the planning process in this manner, it is possible to find routes which are better positioned with respect to uncertainty than routes found by using a deterministic approach. Models accounting for uncertainty also provide information about the level of risk associated with routes. The focus of the thesis is on the practical analysis and decision support rather than on technical performance and model efficiency.

1.3 Upcoming Chapters

In Chapter 2, the problem faced by Odfjell is described in more detail, and aspects important for the mathematical model formulation are highlighted. Chapter 3 reviews literature relevant to stochastic routing problems, while the mathematical model formulation is presented in Chapter 4. In Chapter 5, the method used to solve the mathematical model presented in Chapter 4 is explained. The nature of the stochastic variables is discussed in Chapter 6. A description of the input data applied to generate test instances is presented in Chapter 7. The results obtained by solving the static version of the problem are presented and analysed in Chapter 8. In Chapter 9, the results from solving the dynamic version of the problem are presented and analysed. Conclusion remarks and recommendation for future research is presented in Chapter 10.

Chapter 2

Problem Description

In this chapter, the real-life problem faced by Odfjell SE is described in more detail. Examples from Houston Ship Channel in the Port of Houston are used to highlight some of the issues. First, an overview of the problem is presented in Section 2.1. Implications of including uncertainty in the planning problem are discussed in Section 2.2. Lastly, a summary is presented in Section 2.3. Information about the problem is mostly gathered from Arnesen and Gjestvang (2015) and Kruse (2015).

2.1 The Stochastic Problem

As explained in Chapter 1, the T/C result is an important profitability indicator for a chemical shipping company like Odfjell. Voyage duration is a key factor for the profitability. Odfjell's ships spend a lot of time in port, and the majority of this time is spent waiting for an occupied terminal to become available. Uncertainty associated with the travel times is largely due to uncertainty of these waiting times. The uncertainty complicates the route planning and is a matter which must be taken into consideration. Thus, the focus in this thesis is on planning of a single ship's customer servicing sequence in a port while handling uncertainty in a proper manner. We also study how different elements affect uncertainty and how uncertainty affects the choice of optimal routes.

Geography of Houston Ship Channel

The Port of Houston is located in the state of Texas, and is one of the United State's busiest seaports. Houston Ship Channel is part of the Port of Houston. It is a long and narrow channel, 60 meters wide, 14 meters deep and 80 km long, with numerous terminals. The long and narrow geography makes the planning problem in Houston Ship Channel particularly interesting. Figure 2.1 shows a map of Houston Ship Channel, where the red line marks the channel.



Figure 2.1: Houston Ship Channel in the Port of Houston. The red line marks the channel.

Contracts and Demands

On average, 96 % of the cargo serviced during a port visit is booked by the time the ship arrives at the port. Which contracted cargoes the ship has to service in a specific port is hence considered known a priori. As the terminal location of the cargoes to be serviced are known, this implies that for each port, the ship has a known set of terminals it must visit to service the predefined customers, and quantities are known. It is also known which cargoes are to be loaded and which are to be discharged. It is not required that all cargoes to be serviced at the same terminal have to be serviced consecutively. There is no limitation on the number of times the ship can return to a terminal, but each time a terminal is entered, uncertain waiting time applies.

Tendering and Queuing

When arriving at port, the ship moors at an anchorage point outside the port while waiting for permission to visit terminals. When the ship is ready to visit terminals, notices of readiness (NOR) are tendered to relevant terminals. To be allowed send NORs, the ship must have space in suitable and satisfactory cleaned tanks to service the loading or unloading of cargoes at the terminal. When receiving NORs, the terminals reply with a *berth advice*. The berth advice includes an approximated waiting time before the ship may approach the terminal. Prior to a ship's arrival, terminals must often prepare equipment at the shore, which may take up to 12 hours (Kruse, 2015). Once a terminal is ready to receive the ship, it signals its readiness by sending a *berth readiness*. The time spent waiting for a terminal to become available depends on the number of preceding ships in the queue occupying the berth, as the ships are serviced at a first-come-first-served basis. Upon receiving the berth readiness, the ship arranges for a harbour pilot to sail the ship to its designated berth at the terminal. A ship must cancel all other NORs and loses its place in line when it starts sailing towards a terminal. Once the ship has finished its operations at the terminal, it sends out NORs to the remaining terminals on its schedule. If a terminal is ready to receive the ship, the ship may sail directly to the next terminal. However, the ship often has to wait before the terminal is ready.

In most ports, ships are not allowed to wait at the terminal when it is finished servicing the cargoes in that terminal. Instead, ships must sail back to anchorage or to dedicated layberths before the next terminal visit. If the next terminal sends its berth readiness before the ship has returned to anchorage or a layberth, the ship may turn around on the spot. If not, the ship sails all the way to anchorage or a layberth where it waits until the terminal is ready. Two cases where this may happen are illustrated in Figure 2.2. For both cases, the ship first services a cargo in Terminal 1, before the next cargo located in Terminal 2 is serviced. The full lines illustrate the distance traveled when the ship sails directly from Terminal 1 to Terminal 2. The dashed lines illustrate scenarios where a long waiting time at Terminal 2 forces the ship to sail all the way back to anchorage, where it waits until berth readiness is received. Then it may sail back to Terminal 2.

In Figure 2.2a, Terminal 2 is located further away from anchorage than Terminal 1. If the ship has to wait before Terminal 2 is ready to accommodate it, the ship must sail towards anchorage and possibly wait there until the terminal is ready. In this case, sailing towards anchorage means sailing *away* from the destination terminal. This is illustrated by the upper dashed line in Figure 2.2a. Consequently, if the waiting time at Terminal 2 is larger than zero, the distance the ship has to sail to get from Terminal 1 to Terminal 2 becomes larger than the direct geographical distance between the two terminals. If Terminal 2 is located closer to anchorage than Terminal 1, as illustrated in Figure 2.2b, sailing towards anchorage means sailing *towards* the next terminal initially. For scenarios with zero waiting time at Terminal 2, the ship may sail directly to the terminal. This is illustrated by the full line. If the waiting time at Terminal 2 is short enough, i.e. berth readiness is received before the ship has passed the Terminal 2, the ship avoids unnecessary additional sailing time and may also then sail directly to the terminal. If the waiting time exceeds the time it takes to travel from Terminal 1 to Terminal 2, the ship passes Terminal 2 and sails away from it on its way to anchorage. Consequently, the distance between the ship and the next terminal increases. The ship may have to wait at anchorage for some amount of time before it can sail back to Terminal 2. This case is illustrated by the dashed line in Figure 2.2b. The dashed circle illustrates waiting at anchorage.



(a) Possible additional sailing distance due to waiting times when sailing from Terminal 1 to Terminal 2, when Terminal 2 is located further away from the anchorage than Terminal 1.



(b) Possible additional sailing distance due to waiting times when sailing from Terminal 1 to Terminal 2, when Terminal 2 is located closer to the anchorage than Terminal 1.

Figure 2.2: Additional sailing distance due to waiting times. The full line represent the direct sailing distance if no waiting times incur, while the dashed lines illustrate the possible additional sailing distance due to waiting times.

Thus, in both cases of Figure 2.2, the actual distance traveled between terminals may exceed the *direct* geographical distance between the terminals. Hence, the uncertain waiting time causes the total distance to be sailed to be uncertain which results in uncertain sailing time.

Draft Limits

Draft is the height of the submerged ship, from the lowest part of the ship to the waterline. $Draft \ limit$ is the depth of the water, from the seabed to the waterline. The terminals have different draft limits. The draft of the ship depends on the weight of the ship, including the cargo it is carrying, as can be seen from Figure 2.3.

Based on the ship's weight and other characteristics, the draft limits can be expressed in terms of the weight of the cargoes on board the ship. Thus, in this thesis, the draft limits are expressed as tonnes of load on board the ship. The route planning problem we consider includes both pickup and delivery of cargoes, making it uncertain whether the ship's draft is largest when arriving or departing a terminal. This implies that the ship's draft must be evaluated both when arriving and departing the terminals, and may constrain the problem (Arnesen and Gjestvang, 2015).



Figure 2.3: Draft and draft limit of a) an unladen ship, and b) a laden ship (Rakke et al., 2012).

Ship Capacity

For each customer request, the quantity to be serviced is known, and a request cannot be split. The ship has limited capacity, and this capacity must be complied to at all times. When arriving at a port, the ship is loaded with cargoes to be discharged, and the ship capacity may constrain the order in which cargoes can be loaded and discharged throughout the port visit. If all of the load on board when arriving at the port is discharged during the port visit, the ship is loaded only with the total load picked up during the port visit when leaving the port. If some of the load on board at arrival is not discharged, the load when leaving the port exceeds the total load picked up during the port visit. Regardless, the ship capacity can never be exceeded.

Inspections

When arriving at port, ships go through inspections. The ship may only enter the port after the tanks and cargo on board has been inspected and approved. The coast guard must also clear the ship before it may proceed to its business. In addition, customs and immigration inspect the ship, its documents and cargo at the terminal. Inspections may be visual, or there may be a more thorough chemical analysis. If a ship does not satisfy the requirements, it must sail to anchorage and either wash or change the tanks.

Terminal Operations

There is a number of operations to be performed once the ship arrives at the terminal. These include customer-required inspections, loading and unloading of cargo and cleaning of tanks. If there are enough resources available, several cargoes may be serviced at the same time. The time it takes to load or unload is referred to as *loading time*. When cargo has been unloaded, the tanks must be cleaned. The time it takes to clean a tank is referred to as *cleaning time*, and depends on the type and size of cargo unloaded. Total servicing times at terminals in the Port of Houston are on average 24 to 26 hours (Kruse, 2015).

2.2 Implications of Including Uncertainty

In the above section we identify five major time consuming activities a ship goes through several times during a port visit: sailing between terminals, loading/unloading of cargo, cleaning of tanks, waiting for terminals to become available, and the possible *additional* sailing time incurred because of waiting times. In reality, all these activities include uncertainty, and the waiting time is particularly uncertain. The waiting time depends on the amount of traffic, that is the number of ships in line to be serviced, and can affect the choice of optimal route. Note that we distinguish between *sailing time* and *travel time*. The sailing time is the time spent sailing or waiting at anchorage, while the travel time includes loading/unloading and cleaning in addition to sailing time. Hence, the travel time always exceeds the sailing time. The problem we study can be regarded as finding a route through a network of nodes, where the arcs between nodes have given weights, or costs, associated with them. For our problem, we let the weight of an arc be the travel time.

In additional to uncertain waiting times, travel times are subject to the possibility of unscheduled events or failure. The ship could meet another vessel in a narrow part of the port and be forced to turn around or do another manoeuvre that leads the ship away from its planned path. The engine could break down or equipment onshore needed for the loading/unloading could break down. But there is a very low probability for these events to occur, and their effect on what solution is optimal is expected to be significantly less than the uncertain waiting times. Due to the uncertain waiting times that apply when entering a new terminal, other routes may be optimal for the stochastic problem than for the corresponding deterministic problem. What route is optimal is dependent on how optimality is defined. From a business point of view, it is not necessarily only the expected completion time of the route that should be optimized. The reality companies face is complex, and the ability to plan ahead can be of equally great importance. Reducing uncertainty as much as possible at the cost of a somewhat higher expected completion time may thus yield better results. Probabilities are affected by both the expected value and the variance of the completion time, and each route has a unique combination of these two characteristics. What is regarded as an optimal route depends on the company's risk profile.

As illustrated in Figure 2.2 and explained above, the time it takes to sail from one terminal to another is closely related to distances, directions and traffic and the waiting times at terminals. As such, the uncertain travel times are dependent on a combination of both uncertain elements and deterministic elements.

Figure 2.4 illustrates two possible routes for a simplified and limited route planning problem.



Figure 2.4: An example of two feasible routes with different means and variances of route completion times.

The figure shows three terminals, A, B, and C, and the anchorage. Above each grey square representing a terminal are the characteristics of the normally distributed waiting times, $\mathcal{N}(\mu, \sigma^2)$. For this simplified example waiting time only applies to Terminal A. The draft limits of each terminal is labeled by DL. The draft limits are expressed as the maximum load allowed on board the ship for it to enter (or leave) the terminal. There are three loads included in this limited problem, one in each terminal. The load of the cargo located in a terminal is shown inside the terminal icon. The sign "+" indicates a pickup cargo, while the sign "-" indicates a delivery cargo. The horizontal lines between the terminals represent the direct sailing times between terminals in hours, i.e. the time it takes to sail between terminals when waiting times are ignored or are equal to 0. The vertical lines do not impose any sailing time. The terminal visiting sequence of each of the two routes are illustrated in the figure. The sailing times along arcs are also assumed to be normally distributed, and the characteristics of the distribution of each arc sailing time is indicated above the arc.

The figure shows that the terminal visiting sequence of Route 1 is Anchorage-A - C - B - Anchorage. This amounts to an expected completion for the route of $\mu = 32$ hours, and $\sigma = 5$ hours. The terminal sequence of Route 2 is Anchorage -C - A - B - Anchorage. Relative to Route 1, Route 2 includes a detour, and gives an expected route completion time for Route 2 of $\mu = 32.5$ hours with $\sigma = 3.1$ hours. For both routes, the tight draft limit at Terminal B forces the ship to visit Terminal A before Terminal B. The variance of the two routes are not the same. Terminal A is the only terminal with uncertain waiting times, and arcs with terminal A as the destination node are the only uncertain arcs. Route 1 visits Terminal A as the first node, while Route 2 visits Terminal A after terminal C, which is further away, and then takes a detour to Terminal B before finishing. The arc from C to A is associated with less uncertainty than the arc from Anchorage to A. This is due to different sailing direction relative to anchorage, as discussed for Figure 2.2, and that some of the time spent waiting is spent sailing in the correct direction. This gives different expected completion times and variances of the routes. For one route, the expected completion time is lower at a cost of a higher uncertainty, while for the other route, the uncertainty is lower at the cost of a slightly higher expected completion time. If the objective is to merely minimize the expected route completion time, Route 1 is optimal, but if uncertainty matters, what route is optimal is not obvious.

The example above shows that routes can have different expected completion times and variances depending on what arcs are chosen. The uncertainty related to terminals, directions of arcs relative to anchorage and relative distances can cause some routes to have a lower mean, but higher variance than others. This indicates that for stochastic versions of the problem, there is a trade off between time saved and uncertainty reduced, which must be taken into account when optimality is defined.
2.3 Summary

In this thesis, we consider a problem where a single ship is going to service a set of customers within a port. Customers requires either pickup or delivery of cargoes. Both the customers and their requests (i.e. locations and quantities) are known. Each cargo is located in one out several possible terminals within the port. The ship starts from anchorage and must service all cargoes at their respective terminal locations before finalizing the port visit and returning to anchorage. The aim is to determine a sequence to service the given cargoes which is optimal with respect to both expected route completion time and uncertainty.

Sailing between cargo locations includes a set of time consuming activities, such as waiting for the next terminal to become available when consecutive cargoes are placed in different terminals. Terminals are subject to traffic, and the time it takes before terminals are ready to receive the ship is uncertain. The ship is not allowed to wait at a terminal until the next terminal to visit becomes available, and as the length of the waiting time is uncertain, the time it takes to sail between terminals also becomes uncertain. The amount of uncertainty and the expected time associated with sailing between terminals is dependent on the distance between terminals, direction relative to anchorage and the waiting times. The resulting total travel time, which also includes deterministic terminal operations, is stochastic. This complicates the planning problem.

Chapter 3

Literature Review

In this chapter we present and discuss literature relevant for the stochastic in-port ship routing problem studied in this thesis. The problem can be modelled as a Single Vehicle Pickup and Delivery Problem with Draft Limits and Stochastic Travel Times (PDP-DLST). In this chapter, the problem studied in this thesis is referred to as *PDP-DLST*. Classification of and literature on the *Pickup and Delivery Problem* (PDP) is reviewed in Section 3.1, before stochastic routing problems are addressed in Section 3.2. As the literature on PDPs with stochastic travel times is scarce, other relevant routing problems with stochastic travel times are studied, among them the Travelling Salesman Problem (TSP), the Vehicle Routing Problem (VRP), and the Shortest Path Problem (SPP). PDP is closely related to the TSP, and the VRP is a generalization of the TSP (Dantzig and Ramser, 1959). Both theory and solution methods of stochastic routing problems is addressed. The thesis problem is a dynamic problem by nature, but it is also interesting to consider a static version of the problem. Thus, theory and literature on static and dynamic routing problems is reviewed in section 3.3. Lastly, Section 3.4 sums up the contribution of this thesis to the literature.

3.1 Pickup and Delivery Problems

The *Pickup and Delivery Problem* (PDP) is an important class of the *Vehicle Routing Problem* (VRP), in which objects or people are transported between origins and destinations (Berbeglia et al., 2007). As for the VRP, a number of variants of the PDP has been studied. In this section we first present a framework for classification of PDPs. The framework is suggested in Berbeglia et al. (2007) and presented here because of its comprehensive and usable

structure. Secondly, we present general notation for the PDP, suggested by Berbeglia et al., before literature on solution methods for PDPs is reviewed.

3.1.1 Classification

There is a vast amount of versions of the PDP. To get an overview of the different versions of the PDP, and to be able to classify the problem studied in this thesis, we present a modelling framework and classification scheme for PDPs. The framework is introduced by Berbeglia et al., and is shown in Figure 3.1. A three-field scheme to classify the problem versions is used and it is denoted by [Structure|Visits|Vehicles]. Structure indicates the number of origins and destinations. Visits specifies the way in which pickup and delivery operations are performed at customer vertices. Vehicles gives the number of vehicles used. According to the survey, the thesis problem can be classified as a version of PDP referred to as a dynamic single-vehicle one-to-many-to-one Problem with Single Demands and Mixed Solutions ([1-M-1|P/D|1], or 1-M-1-PDPSDMS). The notation is described in detail in the remainder of this section.



Figure 3.1: Classification of PDPs by Berbeglia et al. (2007).

The authors define many-to-many problems (M-M) as having several origins and several destinations for each commodity. For one-to-one problems (1-1), each commodity has exactly one pickup vertex and one delivery vertex. In oneto-many-to-one problems (1-M-1), some commodities are initially located at the depot and destined to customer vertices while other commodities supplied by the customers are brought back to the depot. Thus, there is only one origin and one destination. The problem studied in this thesis has only one anchorage which for modelling purposes represents both the origin for all deliveries and the destination for all pickups. Transhipment is not allowed, and the problem is a 1-M-1 PDP. Hence, the focus is on 1-M-1-PDPs for the remainder of this section.

Berbeglia et al. (2007) further distinguish between the demand structures single demands and combined demands. P/D indicates that each customer node is either a pickup or a delivery node, but not both. This is referred to as single demands. P-D indicates that at least one customer node is both a pickup request and a delivery request simultaneously, which is referred to as combined demands. The problem studied in this thesis is a P/D problem, and the remainder of this section focuses on this type of problem.

When delivery and pickup customers may be serviced in any order, the problem is classified by Berbeglia et al. (2007) as *mixed*. When all delivery customers must be serviced before any pickups, the problem is said to have *backhauls*. The problem we study is a mixed problem as cargoes may be serviced in any order. Note that *backhauls* may also be used to refer to pickup requests destined to a warehouse (which corresponds to an anchorage). Anily and Mosheiov (1994) use *backhauls* to refer to pickup customers destined directly to the warehouse, but does not constrain the order in which the customers may be serviced. When all pickup customers are backhauls, it implies that they do not allow delivery requirement to be satisfied using backhaul stock. This corresponds to a 1-M-1 problem, but Anily and Mosheiov (1994) do not use this notation. This somewhat different usage of the term allows Anily and Mosheiov (1994) to use the term *backhaul* when defining a problem that corresponds to what Berbeglia et al. (2007) refers to as *mixed*.

Berbeglia et al. (2007) also distinguish between *static* and *dynamic* problems. In the former, all information is assumed to be known a priori, while the latter allows information to gradually be revealed over time. The thesis problem is dynamic in reality. Hence, according to Berbeglia et al. (2007), the problem studied in this thesis is a *dynamic one-to-many-to-one Problem with Single Demands and Mixed Solutions*. A model handling the problem statically will be used as the key building block to develop a dynamic approach.

It is apparent that the deterministic version of the problem described in Chapter 2 is classified by Berbeglia et al. (2007) as one-to-many-to-one Problem with Single Demands and Mixed Solutions ([1-M-1|P/D|1], or 1-M-1-PDPSDMS). The deterministic version of this problem has been much studied, but different authors are not consistent in their denotation of the problem. Mosheiov (1994) refers to it as Travelling Salesman Problem with Pickup and Delivery (TSPD), while Anily and Mosheiov (1994) call the same problem TSP with Delivery and Backhauls (TSPDB), and Baldacci et al. (2003) as TSP with Deliveries and Collections (TSPDC). For the remaining part of the chapter we choose to use 1-M-1-PDP when referring to this type of PDP. Note, however, that the problem studied in this thesis extends the 1-M-1-PDP by including draft limits and stochastic travel times. The thesis problem is thus referred to as a PDP-DLST in this chapter.

3.1.2 General Notation

Berbeglia et al. (2007) describe PDPs by the complete and directed graph G = (V, A) in which V = 0, ..., n is the set of vertices. For this thesis, vertices and *nodes* is used interchangeably. Vertex 0 represents the depot, while the remaining vertices represent the customers. The set of arcs is defined as $A = (i, j) : i, j \in V, i \neq j$. c_{ij} denotes the non-negative cost or travel times. The set of commodities to be transported is defined as H = 1, ..., p, and the supply or request of a commodity h at vertex i is denoted as $d_{ih} \in D = (d_{ih})$. If $d_{ih} < 0, -d_{ih}$ is the amount of commodity h requested by node i. The fleet to be used to carry the commodities is the set of carriers K = 1, ..., m, each with a capacity Q. The problem described in Chapter 2 conforms with the above problem description, and is thus a PDP problem. Each $i, j \in V$ describes the set of cargoes, the anchorage point is represented by the depot, m = 1 corresponds to the single ship. Regarding the set of commodities, H, we ignore what type of commodities are handled, which corresponds to assuming p = 1, and d_{ih} is known. c_{ij} represents the travel time from cargo i to cargo j. However, there is uncertainty associated with c_{ij} for the problem described in Chapter 2. Note that c_{ij} is not used as notation for the travel times in this thesis, but they do correspond with a stochastic version of c_{ij} as described by Berbeglia et al. (2007).

3.1.3 Solution Methods

When designing a combined pickup and delivery route, one needs to consider an additional constraint: the vehicle load must remain feasible throughout the tour (Mosheiov, 1994). Hence, the 1-M-1-PDP can be considered as a generalization of the TSP. 1-M-1-PDP reduces to the classical TSP by defining the total pickup load or the total delivery load to be zero. TSP is NP-hard, and so is 1-M-1-PDP (Anily and Mosheiov, 1994). As such, exact solution methods become impractical when the problem size increases, and several authors have developed heuristics for solving the 1-M-1-PDP in polynomial time. Mosheiov (1994) presents two such heuristics. The first is called *Pickup and Delivery along Optimal Tour*, and starts with an optimal TSP tour through all customer points. Next, a feasible starting point with respect to capacity restrictions is identified by search. By inserting the depot on the arc between the start node and its predecessor, a feasible solution to the 1-M-1-PDP is obtained. The second heuristic is an extension of the well known TSP heuristic *Cheapest Insertion*, referred to as *Cheapest Feasible Insertion*. Starting from an optimal delivery tour, pickup-points are inserted one at a time, always choosing the cheapest insertion while maintaining feasibility. The author argues that if an optimal TSP tour can be found, pickup and delivery along the tour is a better technique than insertion of pickup points into an optimal delivery tour. However, the author shows that both heuristics are able to solve instances with up to 200 customer within a reasonable time.

Anily and Mosheiov (1994) suggest a heuristic based on doubling the minimum spanning tree. They prove that their heuristic has a better bound than the heuristic developed by Mosheiov (1994). However, computational experience shows that Mosheiov's heuristic performs somewhat better, while the heuristic proposed by Anily and Mosheiov has a substantial lower running time. All these heuristics are based on construction heuristics using the arc costs between the customer locations. In the problem studied in this thesis the arc costs (i.e. travel times) are uncertain. As such, it may be challenging to adapt the proposed heuristics to the thesis problem.

Two heuristics for the 1-M-1-PDP are developed by Gendreau et al. (1999). The first is based on the optimal solution of the special case arising when the graph, which represents the pickup and delivery network, is a cycle. A linear exact algorithm is developed and then used as a basis for solving general cases of the 1-M-1-PDP. The second heuristic is based on tabu search tailored for the 1-M-1-PDP. The author did several computational experiments to compare the new heuristics to the ones proposed by Mosheiov (1994) and Anily and Mosheiov (1994). The results showed that a combination of the two heuristics by Gendreau et al. (1999) outperformed the others.

Baldacci et al. (2003) suggest a branch-and-cut algorithm for 1-M-1-PDP using a two-commodity flow formulation, which is able to solve problems with up to 200 customers. But the authors define the pickup goods to be of one kind, and the delivery goods to be of another. For the problem we study, however, there is no restriction saying that pickup and delivery cargoes can not be of the same type.

Süral and Bookbinder (2003) propose a model for the single VRP with unrestricted backhauls. Serving of customers' pickup requests is optional, and the problem deals with choosing the best optional backhaul opportunities to optimize revenues. The model is based on Miller-Tucker-Zemlin subtour eliminating constraints, and several tight LP-relaxations are considered. They propose an exact solution method which is able to solve medium-sized practical problems. The problem they study applies to the real-life problem faced by Odfjell. But for this thesis, the decision regarding which customers to serve is beyond the scope of the planning problem, i.e. spot cargoes are ignored.

An important distinction between 1-M-1-PDP and the problem described in

Chapter 2, is the stochastic travel times. To our knowledge, few or none have studied PDP with stochastic travel times. However, PDP with stochastic demands has been studied by Swihart and Papastavrou (1999), among others. The problem they study handles stochastic demands which arrive according to a Poisson process, while the problem studied in this thesis regards stochastic travel times.

The models and solution methods described in this section handle *deterministic* problems in which all information is known deterministically. As described in Chapter 2, the problem studied in this thesis is subject to uncertain waiting and travel times and is thus a stochastic problem. Literature on stochastic routing problems is presented in the next section.

3.2 Stochastic Routing Problems

In this section, we take a closer look at *stochastic* routing problems. Realworld routing problems and applications often include uncertain data. Several variables of a routing problem may be subject to uncertainty, but most commonly it is the number and/or location of the customers, the size of their requests, or, as in our problem, the travel time along an arc which is subject to uncertainty. Hence, routing problems in which travel times or travel distances are stochastic are of interest in our work. Literature on stochastic PDPs is scarce. The PDP is similar to the much studied TSP, and VRP is a generalization of the TSP (Dantzig and Ramser, 1959). Studies of VRP and TSP with stochastic travel times or distances may thus be relevant for the problem studied in this thesis. The problem of finding the shortest path (SPP) in a stochastic network is a problem that has received a lot of attention in recent years. Although the SPP is a different problem than the PDP, a lot can be learned from the literature on this problem. As such, stochastic VRPs, TSPs and SPPs are reviewed in this section. When uncertainty is at play, it is not obvious what is an optimal route. We first discuss what may be regarded as an optimal route. Secondly, we present various approaches for modelling stochastic routing problems in general, before relevant research on the above mentioned problems is reviewed. A short review of various probability distributions used in routing problems with stochastic travel times is included at the end of this section.

3.2.1 Definition of Optimality

When dealing with routing in stochastic networks the definition of what is an optimal path is not straight forward. Loui (1983) reports that the standard procedure for addressing the time independent stochastic SPP is to identify

the route with least *expected* travel time (LET). Others who support this formulation of optimality for stochastic SPPs are Fu and Rilett (1998), Miller-Hooks and Mahmassani (2000) and Waller and Ziliaskopoulos (2002). The approach of minimizing expected completion time is computationally equal to the deterministic problem (Fan et al., 2005a), in which the deterministic variables take the expected value of their stochastic equivalent. However, the route with the least expected travel time could be subject to a large variance in travel time and thus have a substantial amount of risk associated with it. This would not necessarily be an optimal route for a risk-averse decision maker. If there was a path with slightly higher expected completion time with little probability of realising very large completion times, it might conceivably be preferable (Loui, 1983). As is apparent, the definition of an optimal path is not always obvious in the stochastic context. As it incorporates notions of both expected time and reliability, we consider a multi-criterion optimization problem (Samaranayake et al., 2011). Frank (1969) suggest defining the optimal path to be the path that maximizes the probability of realising a travel time less than a constant k. This definition of optimality in stochastic routing is supported by Mirchandani, Fan et al. (2005a), Nie and Wu (2009), Nikolova et al. (2006) and Samaranayake et al. (2011), among others. While a problem modelled as an LET might be easier to solve, the formulation where the probability of on-time arrival is maximized might be easier to interpret.

3.2.2 Modelling Approaches

Gendreau et al. (1996) state that a stochastic problem is usually modelled as a stochastic program with recourse (SPR) or as a chance constrained program (CCP). In SPRs, one aims to make first stage decisions that take the costs of consecutive stages into account. The total costs to be optimized are thus the cost of the first stage decisions and the expected cost of recourse actions taken in consecutive stages. Charnes and Cooper were the first to introduce CCP as a means of handling uncertainty in which known or approximated distributions for the stochastic variables are used to ensure prescribed levels of probability (Charnes and Cooper, 1959). This is done by constraining the CCP by a prescribed confidence level, i.e. by including a chance constraint. Constraint (3.1) shows a chance constraint where one requires to complete a route by a deadline, b, with a prescribed probability (i.e. confidence level), α .

$$P(Ax \le b) \ge \alpha \tag{3.1}$$

In a CCP one does typically not perform corrective actions in consecutive stages when uncertain variables are realised, and the focus is on a priori planning and finding robust *enough* routes. This differs from SPRs, where recourse actions are used to respond to the realisation of uncertain variables. Another means of handling uncertainty in stochastic routing is the problem of maximizing the probability of realising a route completion time less than a prescribed deadline. It becomes implicit that reliability is of great importance when the confidence level for on-time arrival is maximized. Fan et al. (2005a) introduced an algorithm for solving the problem of maximizing the on-time arrival in a stochastic network, referred to as the SOTA algorithm. In a SOTA model, as for a CCP, known or approximated distributions are exploited to identify optimal routes. The SOTA model is further presented in Section 3.3. The remainder of this section is a review of literature on stochastic routing where the above mentioned approaches are applied.

3.2.3 Travelling Salesman Problems

Kao (1978) is one of the earlier studies of the TSP with stochastic travel times (TSPST). The problem studied involves n cities to be visited once and only once. The aim is to find a route that has the greatest probability of completion by time C, i.e. $P(T(t) \leq C)$ is maximized. This resembles the SOTA formulation. The entire tour has to be fixed a priori. When disregarding the draft limits and the limited ship capacity, the PDP-DLST corresponds well with the problem Kao studies. The author proposes two heuristics for the TSPST. The main solution method presented is a preference order dynamic program, which reduces storage requirement at the expense of computational executions. An implicit enumeration algorithm is proposed as an alternative approach. One achieves a reduction in nodes at the cost of an increase in the number of arcs emanating from each node (Kao, 1978). The author shows that the computational effort of the two approaches is the same, except that they handle nodes and arcs differently.

In some cases, Kao's preference order dynamic program yields suboptimal solutions. This was proved by Sniedovich (1981). For nontrivial TSPSTs it may be difficult to verify the completeness of the proposed preference ordering operator due to potential violation of the monotonicity condition of dynamic programming (DP). The author states that an attempt to correct the procedure transforms it into brute force enumeration. The procedure should therefore be used with caution, unless it is possible to verify the monotonicity condition. In the case of normally distributed travel times, however, it may be appropriate to use Kao's approach as a heuristic.

Carraway et al. (1989) suggest Generalized Dynamic Programming (GDP), a modification of conventional DP, that guarantees optimal solution to the counterexample presented by Sniedovich (1981). With Generalized DP it is possible to find an optimal solution to stochastic combinatorial optimization problems when the monotonicity assumption is violated. Hence, the issues identified by Sniedovich are resolved. The authors assume normally distributed inter city distances, $(\mu, \sigma^2) \in Z$, and uses both μ and σ^2 to define *local* preference conditions, as opposed to a global preference condition. For the very last stage, however, a global preference condition is used.

3.2.4 Vehicle Routing Problems

Kenyon and Morton (2003) study a version of the stochastic VRP (SVRP) that, according to Pillac et al. (2013), can be classified as static and stochastic. A fleet of one or more vehicles is to be routed through a network where travel and service times are uncertain. The routes are selected before the stochastic variables are known and the vehicles must follow the route set a priori. Two alternative models for solving the problem are presented in the article. The first model maximizes the probability that the operation is completed within a specified threshold time, which is similar to the SOTA model proposed in Fan et al. (2005a). The other model minimizes the expected completion time of the tour. This corresponds to the LET model. Kenyon and Morton's problem and the PDP-DLST have several characteristics in common, and both the proposed models may be regarded as interesting for this thesis. The authors argue that the second version is typically easier to solve, but the objective function of the maximization problem is typically easier to interpret. They also show that for problems with only one vehicle, the model where expected completion time is to be minimized is obtained by solving the deterministic equivalent where stochastic variables take their mean value. As mentioned, this is also pointed out by Fan et al. (2005a). Thus, the value of the stochastic solution (VSS) may be negligible and the deterministic solution becomes adequate. A problem with a low expected completion time may still have a high variance in travel time and yield suboptimal results for risk averse planners (Fan et al., 2005a). As such, the minimization problem ignores risk, and it is particularly the maximization model (i.e. the SOTA model) that is interesting in relation to the PDP-DLST. If the minimization model was constrained by a required level of probability (i.e. modelled as a CCP), however, this model becomes more interesting for the PDP-DLST. This is done by Laporte et al. (1992), which is presented in the following paragraph.

Laporte et al. (1992) introduce stochastic travel and service times into VRPs. Three formulations are presented: a CCP, and two distinct recourse models. All three models set routes a priori, and the problem is classified as static and stochastic according to Pillac et al. (2013). Laporte et al. (1992) regard a problem with incapacitated vehicles. Each node must be serviced, and the vehicles have a threshold for target completion time. The CCP uses this threshold to constrain the confidence level for the route. The SPRs penalize the expected value by which route travel times exceed the respective thresholds. The developed branch-and-cut algorithm can solve the CCP with up to 20 vehicles. It can handle problems where the travel times can take on a value from five discrete states. As for the problem presented by Kenyon and Morton (2003), the problem described by these authors differs from the PDP-DLST e.g. by allowing multiple and incapacitated vehicles and regarding service times as stochastic.

Li et al. (2010) study VRP with stochastic travel and service times, and include time windows. The problem they study is similar to the one studied by Laporte et al. (1992), but differs by including time windows. The problem is formulated both as a CCP and a two-stage SPR. The former includes chance constraints for both the time windows and for the route as a whole. The objective of the CCP is to minimize stochastic travel distance. In the SPR, the determination of an a priori route is considered as the first-stage decision. In the second stage, the uncertain travel and service times are realised, and it is possible to calculate the total cost of the route. The overall objective is to minimize the sum of the first-stage routing cost and the expected recourse cost. The authors argue that, since SVRPs combine the characteristics of stochastic and integer programming, they are often regarded as computationally intractable. Thus, it is reasonable to develop a heuristics algorithm to solve SVRPs. The problem is solved by a Tabu Search-based heuristic.

Another two-stage stochastic routing problem is studied by Verweij et al. (2003). To solve the problem, the authors propose a heuristic that penalizes routes that exceed the deadline. The penalty is proportional to the size of the violation. The method uses a sample average approximation technique in which a sample of instance realisations is drawn using a Monte Carlo simulation, and each realisation is solved optimally by means of a deterministic technique. By repeating the method with different samples a statistical estimate of the optimality gap can be computed.

A solution approach for the time-dependent VRPST is suggested by Nahum and Hadas (2009). The problem is modelled as a CCP, where the total average time is minimized. The authors take into account the variation in travel times throughout the day as well as the uncertainty of the travel times in their model. An efficient heuristic, which is a version of the saving algorithm, is introduced. Simulations are used to label each route's probability of being the quickest one. The stochastic and time-dependent data is transformed to deterministic cases using three filters: average values for each time period and probability intervals, the minimal time for all periods regardless of the probability. Results are then compared to the deterministic results, and the approach yields optimal solutions for rather small instances (up to seven customers). However, due to the time dependency and usage of time periods, the problem differs from the PDP-DLST even for a single vehicle case.

3.2.5 Shortest Path Problems

Frank (1969) provides an exact method to compute the continuous probability distribution for the SPP. The shortest paths are identified through a pairwise comparison in a set of enumerated paths. Frank develops efficient methods of hypothesis testing using Monte Carlo results, and was, to our knowledge, the first to present work on shortest paths in random networks.

Loui (1983) considers a general utility function of path length which is monotone and non-decreasing, and proves that the expected utility becomes separable into the edge lengths only when the utility function is linear or exponential. In that case the path which maximizes expected utility can be found using traditional shortest path algorithms. For general utility functions, Loui gives an algorithm based on enumeration of paths with a very large running time, $O(n^n)$.

An exact algorithm for the stochastic shortest path problem is proposed by Nikolova et al. (2006). The objective is to maximize the probability that the path length does not exceed some threshold value. Path lengths are drawn from normal distributions, and routes can be characterized by their mean and variance. By maximizing a quasi-convex combination of the path mean and variance, and by varying the weights used to construct the convex combination, the author is able to identify all the extreme points of the dominant of the projection of the path polytope onto the mean-variance plane. The *path polytope* is the *convex hull* of the feasible {01}-vectors $\mathbf{x} \in \mathbf{R}^m$, where the *x*-vector represents the set of all feasible routes in a network of *n* nodes and *m* edges (Nikolova, 2009). The optimal path is found within the set of these extreme points. The algorithm is applied to solve the PDP-DLST studied in this thesis, and is described in more detail in Chapter 5.

3.2.6 Distribution of Stochastic Travel Times

In stochastic programming, three common ways to characterize parameter uncertainty are distribution based, fuzzy based, and scenario based (Meng and Wang, 2010). The distribution-based approach is commonly used to describe problems with exact concepts which depend on random factors (Meng and Wang, 2010).

Mazmanyan and Trietsch (2013) argue that, according to the Central Limit Theorem (CLT), the uncertain travel times for TSP may be modelled using normal distribution as long as the number of visits is large enough. The CLT gives conditions under which the distribution function of a suitably standardized sum of independent random variables is approximately normal (Adams, 2009). If a sufficiently large number of randomly selected, independent samples are drawn from an identical distribution, the sum of those samples tends to follow a normal distribution - even if the original population does not (Dudley, 1999). In the PDP-DLST the uncertain travel times are assumed to be independent, but not identically distributed. Aleksandr Lyapunov suggests a variant of the CLT where the random variables X_i have to be independent, but not necessarily identically distributed (Adams, 2009). Lyapunov's theorem states that under certain conditions, the CLT holds if the samples are independent, but not identically distributed.

As stated by Taş et al. (2013), the distributions for the travel times most commonly applied in routing problems are normal, log-normal, shifted gamma and gamma distributions. Fan et al. (2005b), Russell and Urban (2008), Fan et al. (2005b) and Taş et al. (2013) assume gamma distribution. Russell and Urban (2008) model the stochastic travel times with shifted gamma distribution. Kaparias et al. (2008) assume log-normal distribution. Li et al. (2010), Carraway et al. (1989) and Ehmke et al. (2015) are some of the many authors who assume normal distributions. However, Ehmke et al. (2015) argue that travel times often do not follow normal distributions in practice. Kenyon and Morton (2003) argue that assuming travel times are normally distributed is inconsistent with non-negative travel times. They also argue that in the case of negligible likelihood of negative travel times, assuming normal distribution may be appropriate nonetheless. Allowing skewed normal distributions allows them to create a more realistic approximation of the distribution of the arc travel times. Kenyon and Morton find that the value of the stochastic solution tends to increase with increased skewness.

3.3 Static and Dynamic Routing Problems

Real-world applications often include two important dimensions: evolution and quality of information (Pillac et al., 2013). Quality of information reflects the possible uncertainty on the available data, and based on this we separate problems into deterministic and stochastic problems. Evolution of information relates to the fact that in some problems the information available to the planner may change during the execution of the routes. Based on the evolution of information we categorize problems as static or dynamic. The problems addressed in Section 3.1 are deterministic in nature, and stochastic problems are reviewed in Section and 3.2. In this section we take a closer look at the distinction between static and dynamic problems.

Routing in a stochastic network is typically either used to provide an *a priori* shortest paths or *adaptive* en-route guidance (Nie and Wu, 2009), depending on whether the problem is treated as static or dynamic. In a static problem, the route is decided merely based on information available initially, and no new information is accounted for during the execution of the planned route.

We say that the route is decided *a priori*, and no changes are made to the initial plan. For a stochastic static problem, *known* information implies that the probability distributions of the uncertain data are known. For a stochastic static problem, the challenge is to design a set of robust routes a priori, that will undergo minor changes during the execution (Pillac et al., 2013). In a dynamic problem, information is revealed during the execution of the route. As new information becomes available, the problem is re-optimized and improved solutions can be found. In a stochastic dynamic problem new information is stochastic data that is realised, i.e. exploitable stochastic knowledge is available on the dynamically revealed information (Pillac et al., 2013). As an example, after an arc is traversed in a stochastic network, its true length becomes known, and a new route can be calculated based on the new information. The problem studied in this this thesis is a dynamic and stochastic problem. As such, literature on dynamic and stochastic routing problems are reviewed in the remainder of this section.

Traditionallay, approaches for solving dynamic vehicle routing problems are straightforward adaptions of static procedures, where a static problem is solved each time new information becomes available (Psaraftis, 1995). This is referred to as an *adaptive* approach. As mentioned, Nikolova et al. (2006) presents an algorithm for solving the stochastic SPP, but considers a nonadaptive, i.e. static, scenario. Nikolova et al. point out that this can easily be converted to an adaptive scenario by rerunning the algorithm when new information becomes available, and the approach thereby becomes adaptive (i.e. dynamic).

As mentioned in Section 3.2, Kao (1978), Sniedovich (1981) and Carraway et al. (1989) study a dynamic and stochastic TSP. Carraway et al. (1989) presents a solution method refered to as Generalized Dynamic Programming (GDP). GDP is based on conventional DP, and with modifications such as usage of local preference conditions, the stochastic problem can be properly handled. In Fan et al. (2005a) a dynamic and stochastic SPP is studied. Given a current location, the goal is to identify the next location to visit so that the probability of arriving at the destination by time t or sooner is maximized. The authors propose an adaptive optimal path algorithm based on conventional DP. The problem is treated as a multistage decision process, and the Bellman principle of optimality is applied to formulate the mathematical model. The Bellman principle of optimality states that an optimal sequence of decisions has the property that, whatever the initial state and decision are, the remaining decisions must be optimal with respect to the state resulting from the initial decision (Fan et al., 2005a).

3.4 Contribution to Literature

The problem studied in this thesis is a pickup and delivery problem with draft limits and stochastic travel times (PDP-DLST). Through this literature review it has come to our attention that there has been done relatively little research on the pickup and delivery problem where travel times are stochastic. The pickup and delivery problem with stochastic *demands* has been studied by others, but the problem becomes rather different when it is the demands, not the travel times, that are subject to uncertainty. The TSP, VRP, and SPP with stochastic travel times have been studied by others as well, but these problems also differ from the problem we study. Our thesis contributes to the literature by studying both a static and dynamic version of the pickup and delivery problem with uncertain travel times, subject to constraining draft limits.

Chapter 4

Mathematical Formulation

In this chapter, the mathematical formulation used to model the problem described in Chapter 2 is presented. Both the static and dynamic versions of the problem are solved using the mathematical formulation presented here. The model is solved iteratively for the dynamic problem, but this is described in more detail in Chapter 9. The method applied for solving the mathematical model is presented in Chapter 5. Section 4.1 describes the assumptions and simplification that are made while Section 4.2 explains the modelling notation. The mathematical formulation is presented in Section 4.3. The applied preprocessing is described in Section 4.4.

4.1 Modelling Assumptions

The problem faced by Odfjell is complex. To be able to apply solution methods to solve the problem, some simplifications are made. The simplifications and assumptions used to relax the real-life problem are addressed in this section.

At each terminal there can be one or several cargoes. It is assumed that all cargoes within a terminal can be serviced from the same berth, meaning that the ship does not need to change berths during a terminal visit. It is also assumed that only one cargo can be handled at a time, i.e. parallel handling of cargoes is not possible. This also applies to the cleaning of tanks, meaning that tanks used for different cargoes have to be cleaned sequentially. We assume tanks to be cleaned once for each delivery cargo, but never for pickup cargoes.

The time it takes for the ship to move from a cargo at one terminal to a cargo

at another terminal includes a number of operations. The operations we have chosen to include are waiting for the next terminal to become available, sailing to the next terminal, loading/unloading of the cargo at the destination terminal and cleaning of the associated tank if it is a delivery cargo. This is illustrated in Figure 4.1. These operations are included because they are the major time consuming activities. Inspections and other minor time consuming activities are ignored. The loading time and cleaning time are assumed to be deterministic. Note that the time spent unloading is also referred to as loading time.



Figure 4.1: Components of the travel time between two customers in different terminals.

For cargoes located in the same terminal, the sailing time between them is set to zero. As the ship does not have to enter a new terminal, the uncertain waiting time and additional sailing are also zero. However, loading and cleaning time still incur. The time between servicing two cargoes located at the same terminal is thus deterministic.

In accordance with the Lyapunov's version of the Central Limit Theorem, as explained in Section 3.2.6, it is assumed that the probability distribution of the completion time of the entire route follows a normal distribution. The accumulated variance of the stochastic variables represents the variance for the entire route, and the mean is the sum of the means of the stochastic variables. The sailing time is the stochastic component of the travel time, and to obtain the mean of the travel time of an arc, the sailing time of the arc is shifted to the right by the sum of the deterministic cleaning and loading times. As explained above, for pickup cargoes, the cleaning time is zero, and only loading time incurs.

A single ship is considered. It is assumed that the charterer is always ready to accommodate the ship, which implies that there are no time windows. In reality, time windows are rather wide, and this assumptions does not cause significant deviation from reality. Transhipment is assumed not to be allowed. In reality there are several layberths which the ship may sail to and wait at instead of sailing all the way to anchorage. However, we assume there is only one anchorage which the ship must sail towards and possibly wait at until terminals are ready to accommodate the ship. We also ignore that the ship in reality can tender to more than one terminal at the same time. In other words, we assume that the ship only sends a NOR when the next cargo to service is in a different terminal than the current one, and a NOR is only sent to that specific terminal.

Restrictions on hazardous chemicals, the capacity of different tanks, and the stability of the ship, makes the allocation of cargoes to tanks a challenge in chemical shipping. In this thesis the tank allocation is not considered, and the problem can be regarded as always having suitable tanks available. Arnesen and Gjestvang (2015) modelled two problems - one with and one without tank allocation. The authors found that there were routes which were optimal for the problem without tank allocations, which were also feasible for the problem with tank allocation. Including tank allocation is hence not crucial to find optimal paths, which supports our choice of not including tank allocation.

As discussed, the definition of optimality is not obvious in a stochastic setting. We define the optimal route to be the route that maximizes the probability of route completion within a given threshold. As mentioned in Chapter 3, this definition is supported by Frank (1969), Mirchandani (1976), Fan et al. (2005a), Nie and Wu (2009), Nikolova et al. (2006) and Samaranayake et al. (2011), among others.

4.2 Modelling Notation

The problem is defined on a graph G = (N, A), where N is the set of all nodes and A is the set of all arcs in the network. Cargoes are modelled as nodes indexed by i and j. Let $N = \{0, ..., n + 1\}$ be the set of all nodes, and $N^C = \{1, ..., n\}$ be the set of all cargoes. The anchorage is modelled as both the origin node, i = 0, and the destination node, i = n + 1, at the same geographical place at sea outside the port. A cargo is either a pickup node or a delivery node. Let $N^+ \subset N^C$ be the set of pickup nodes, and $N^- \subset N^C$ be the set of delivery nodes.

Let Q^+ be the total amount in tonnes that will be picked up, and Q^- the total amount that will be delivered. All weights are given in tonnes. Let Q_i be the weight of cargo *i*. The weights of the pickup cargoes are given as positive numbers, and the weights of the delivery cargoes as negative numbers. The total weight capacity of the ship is denoted by K, while D_i is the draft limit at the terminal associated with cargo *i*. The draft limits are represented in terms of weight. \tilde{t}_{ij} is the time it takes from the end of loading/unloading cargo *i* to the end of loading/unloading cargo *j*. Included in \tilde{t}_{ij} is the deterministic loading time of cargo *j*, l_j , the cleaning time for the tank associated with cargo *j*, c_j , the stochastic sailing time, \tilde{s}_{ij} , and the stochastic waiting time, \tilde{w}_j , to the end node. For pickup cargoes, c_j is zero. The reason why the loading and cleaning time of cargo *j* is included in \tilde{t}_{ij} instead of the loading and cleaning time for node *i* is that, in reality, the ship decides the next node when it is finished servicing node *i*, and by choosing the next node, it also commits to the loading and cleaning time of the next node it decides to visit. Hence, it makes sense that the washing and cleaning times of node j are included in \tilde{t}_{ij} . As is explained in more detail in Chapter 6, the sailing time depends on the waiting time, and these elements are hence combined into one variable, $\tilde{s}(\tilde{w})_{ij}$. \tilde{t}_{ij} is the weight on the arc from i to j in the network, and can be expressed as

$$\tilde{t}_{ij} = \tilde{s}(\tilde{w})_{ij} + l_j + c_j.$$

The expected value of \tilde{t}_{ij} is μ_{ij} ,

$$\mu_{ij} = E[\tilde{t}_{ij}] = E[\tilde{s}(\tilde{w})_{ij}] + l_j + c_j.$$

The matrix σ_{ij}^2 holds the variance for arc $(i, j) \in A$. σ_{ij}^2 is non-zero for arcs where cargo *i* and cargo *j* are located in different terminals, and zero otherwise. How the variance of each arc is found, is further explained in Chapter 6. All arcs with node n + 1 as destination node are not subject to waiting times and are given deterministic sailing times.

The binary variable, x_{ij} , is equal to one if the ship sails from node *i* to node *j*, and zero otherwise. The continuous variable y_{ij} denotes the total weight on board the ship when sailing from node *i* to node *j*. The variable σ^2 represents the sum of the variances of the arcs used for the complete route. The expected completion time for the entire route is denoted by μ , and the realised total travel time is denoted by *T*.

4.3 The Mathematical Model Formulation

Given a threshold, H, for the completion time of a route, the objective is to maximize the probability, α , of obtaining a total travel time, T, less than the threshold. Uncertainty is associated with each route, and the completion time for each route follows a normal distribution. Hence, a route can be characterized by its expected completion time, μ , and variance, σ^2 . The objective function becomes

$$max \qquad \alpha = P(T \le H) = \Phi(\frac{H-\mu}{\sigma}), \qquad T \sim \mathcal{N}(\mu, \sigma^2)$$
(4.1)

Where $\Phi(z)$ is the cumulative distribution function (CDF) of the standard normal distribution, and $z = \frac{H-\mu}{\sigma}$ in accordance with statistics theory. Figure 4.2 illustrates the maximization problem. $P(T \leq H)$ is the probability that the route completion time, T, takes a value less than or equal to the threshold, H, when T follows a normal distribution. This probability is represented by the grey area in Figure 4.2, and the model aims at maximizing the grey area.



Figure 4.2: The grey area is the probability of the completion time, T, being lower than the threshold, H.

Flow Constraints

$$\sum_{i \in N} x_{0j} = 1, \tag{4.2}$$

$$\sum_{j \in N \mid (i,j) \in A} x_{ij} = 1, \qquad i \in N^C$$
(4.3)

$$\sum_{i \in N \mid (i,j) \in A} x_{ij} = 1, \qquad j \in N^C$$

$$(4.4)$$

$$\sum_{i \in N} x_{i,n+1} = 1, \tag{4.5}$$

Constraints (4.2)-(4.5) describe the flow along the route, making sure every node is visited once, and only once, and that the ship starts and ends its route at the anchorage.

Cargo Constraints

$$\sum_{j \in N^C} y_{0j} = Q^- \tag{4.6}$$

37

Constraint (4.6) ensures that when leaving from anchorage at the beginning of the route, the ship contains all the cargo that is to be delivered at the terminals.

$$\sum_{(i,j)\in A} y_{ij} - \sum_{(j,i)\in A} y_{ji} = -Q_j, \qquad j \in N^C$$
(4.7)

Constraint (4.7) ensures that the difference between ingoing and outgoing shipload in each node equals the weight of the cargo that has been loaded or unloaded in the node.

Capacity Constraints

0 <

$$0 \le y_{0j} \le Q^{-} x_{0j}, \qquad (0,j) \in A \mid j \in N \setminus \{0\} \qquad (4.8)$$

$$0 \le y_{i,n+1} \le Q^+ x_{i,n+1}, \qquad (i, n+1) \in A \mid i \in N \setminus \{n+1\} \qquad (4.9)$$

$$Q_i x_{i,i} \le y_{i,i} \le (K - Q_i) x_{i,i}, \qquad (i, i) \in A \mid i, i \in N^+ \qquad (4.10)$$

$$(Q_i - Q_j)x_{ij} \le y_{ij} \le Kx_{ij}, \qquad (i, j) \in A \mid i \in N^+, j \in N^-$$
(4.11)

$$-Q_j x_{ij} \le y_{ij} \le (K+Q_i) x_{ij}, \qquad (i,j) \in H + i \in \mathbb{N}, \ j \in \mathbb{N}$$

$$(4.11)$$

$$(4.12)$$

$$0 \le y_{ij} \le (K - Q_j) x_{ij}, \qquad (i, j) \in A \mid i \in N^-, j \in N^+ \qquad (4.13)$$

Constraints (4.8) and (4.9) connect the x- and y-variables in to and out from anchorage, respectively. Constraint (4.8) prevents the initial load on the ship to exceed the total delivery quantity. Constraints (4.9) prevents the load on the ship when arriving at anchorage at the end of the route to exceed the total pickup quantity. Constraints (4.10)-(4.13) ensure that the load on the ship never exceeds the total capacity of the ship.

Draft Limit Constraints

$$0 \le y_{ij} \le D_j x_{ij},$$
 $(i,j) \in A \mid j \in N^-$ (4.14)

$$0 \le y_{ij} \le D_i x_{ij}, \qquad (i,j) \in A \mid i \in N^+$$

$$(4.15)$$

(4.16)

Constraints (4.14) and (4.15) impose limits on the ship's draft when arriving at a delivery node and departing from a pickup node, respectively.

Time Constraint

$$\mu = \sum_{(i,j)\in A} \mu_{ij} x_{ij} \tag{4.17}$$

Constraint (4.17) sets the expected total time of the route, μ , equal to the sum of the expected travel time of the arcs used. μ_{ij} is the expected travel time of the arc between *i* and *j*.

Variance Constraint

$$\sum_{(i,j)\in A} \sigma_{ij}^2 x_{ij} = \sigma^2 \tag{4.18}$$

Constraint (4.18) sets the variance for the route equal to the accumulated variance of the travel times of each arc used.

Integer Constraint

$$x_{ij} \in \{0, 1\},$$
 $(i, j) \in A$ (4.19)

The decision variable x_{ij} is binary, and an arc may either be used or not. Hence the binary restriction imposed in constraint (4.19).

Subtour Eliminating Constraints

Additional constraints are included to strengthen the formulation. The intention is to reduce the computational time and more efficiently find a feasible solution. Rakke et al. (2012) show that introducing subtour elimination constraints, although it is not necessary for the completeness of the model formulation, gives a significant advantage. Arnesen and Gjestvang (2015) also proved that constraint (4.20) improved the solution time. Including the two-node subtour eliminating constraints

$$x_{ij} + x_{ji} \le 1,$$
 $(i,j) \in A$ (4.20)

is not necessary, but reduces the run times of the model and the constraints are thus included as valid inequalities.

4.4 Preprocessing

To further improve the solution speed, additional preprocessing is done before running the model. The applied preprocessing is explained in this section. An important factor affecting the solution time is the number of symmetric solutions, i.e. solutions that seem logically different to the solver, but are equal for all practical purposes (Williams, 2013). An example is when several cargoes are to be delivered in the same terminal. Then, for most practical cases, it is of no difference which sequence these are delivered in since we do not include time windows. By defining rules that describe sequences within the terminals we are able to remove a number of arcs from the network. Hence, to reduce the number of symmetric solutions, as many arcs as possible are removed from the network prior to solving the problem. As all arcs in the network are defined by the set A, removing an arc from the network is equivalent to removing it from A. The value of the removed arcs in the matrix A are set to 0, and the ship is only allowed to use an arc (i, j) if A(i, j) = 1.

When the ship is in a terminal, it is assumed that it always delivers all cargo destined to the terminal before picking up any cargo from the same terminal. Consequently, the arcs from a pickup node to a delivery node within a terminal, are removed from A. In addition, it is assumed that the delivery cargo with the largest load is to be discharged first, then the second largest and so forth. This implies that once a delivery cargo is discharged all other delivery cargoes in the same terminal are discharged as well. By defining a delivery sequence like this, several arcs are excluded and symmetry reduced. Figure 4.3 illustrates which arcs are removed in a terminal with several delivery cargoes. The figure also illustrates the arcs that are kept in the problem. The "+"-sign in a node indicates that it is a pickup cargo, and a "-"-sign indicates that it is a delivery cargo. The number inside each node icon indicates the size of the load. In reality these loads typically range from 115 to 7,225 tonnes (Arnesen and Gjestvang, 2015), but their relative sizes are more important than realistic representation to explain arc removal here.

When it comes to removing arcs, the biggest difference between delivery and pickup nodes is that once a terminal is visited the first time, *all* the delivery cargoes are discharged consecutively, while pickup cargoes can be loaded at different terminal visits and in different combinations depending on the draft limits. Hence, if there are several pickup nodes in a terminal, and since these can be picked up at different terminal visits and in various combinations, we can not remove arcs to/from nodes in other terminals like we can with delivery nodes. Even so, it is possible to remove symmetry by assuming that, for all practical reasons, if more than one pickup node is serviced during the same terminal visit, the larger cargo is always picked up before the smaller. As such, all arcs from a pickup cargo to all larger pickup cargoes in the same terminal are removed. Figure 4.4 illustrates what arcs are removed internally between pickup nodes in the same terminal to remove symmetry.

As mentioned, anchorage is modelled as two nodes, the origin node, 0, and a destination node, n + 1. It is not possible to sail to the origin node or out from the destination node. Consequently, these arcs are removed from A. It



Figure 4.3: Arcs to be removed from the matrix A (left) and those being kept in A (right) related to delivery nodes to reduce symmetry. Arcs ending or beginning outside the terminal represents all arcs to or from cargoes in other terminals. The numbers represent the size of the loads.



Figure 4.4: Arcs that are removed from the matrix A between pickup nodes in the same terminal to reduce symmetry (left), and the arcs that are kept in A (right). The numbers represent the size of the loads.

is not possible to sail from the origin directly to the final destination, n + 1, or from a node to the same node (self loops). These arcs are removed from A as well.

Chapter 5

Solution Methods

The mathematical model presented in Chapter 4 has a non-linear fractional objective function. This implies that the problem is not straightforward to solve and a specialized solution method is needed. The method presented by Nikolova, mentioned in Chapter 3, is chosen as the applied solution method, and the exact algorithm is described in detail in Section 5.1. The adaption of the algorithm to solve the problem is explained in Section 5.2, and the method involves solving the mathematical model presented in Chapter 4. Note that the solution method presented here is used iteratively to solve the dynamic version of the problem. Adaptions to solve the dynamic version of the problem are presented in Section 9.1 in Chapter 9.

5.1 Nikolova's Method

The problem described in Chapter 4 is solved using an algorithm presented by Nikolova et al. (2006). The suggested algorithm is exact and was initially used to solve the stochastic SPP. The material in this section is mainly based on the work done by Nikolova et al. (2006) and Nikolova (2009). Nikolova et al. consider the problem of finding the shortest path in a graph with independent randomly and normally distributed edge lengths. The optimal path is defined to be the one that maximizes the probability that the path length does not exceed a given threshold (or deadline), t. The optimal path to the problem they study is a solution to

$$\max \quad \frac{t - \mu x}{\sqrt{\sigma^2 x}}$$
s.t. $x \in \text{path polytope}$
 $x \in \{0, 1\}^m,$

$$(5.1)$$

where the $\{0, 1\}^m$ -vector is a set of vectors representing the edges, each with m elements which are either 0 or 1. The vectors μ and σ^2 represents the mean and variance of each of the m edges, respectively. The expression in the objective function in (5.1) corresponds to computing the z-score (also known as standard score) of the route given by the x-vector. Nikolova et al. prove that for deadlines larger than the mean of the smallest-mean path, the optimal path is an extreme point of the *dominant of the path polytope shadow*. Some definitions are given in the following paragraph.



Figure 5.1: Projection of the unit hypercube and the path polytope in the (μ, σ^2) -plane (Nikolova et al., 2006).

For a graph, G, with n nodes and m edges, a solution (i.e. a route) is represented by the $\{0, 1\}$ -vector $\{x\} \in \mathcal{R}^m$, where $x_i = 1$ if edge i is in the route and 0 otherwise. This gives 2^m possible subsets of edges, and all the subsets correspond to the vertices of what is referred to as the *unit hypercube* in \mathcal{R}^m . The convex hull of the set of *feasible x*-vectors is called the *path polytope*, and is a subset of the unit hypercube. The *shadow* of the path polytope is the convex polygon we get when the path polytope is projected onto the span of vectors $\mu = (\mu_1, ..., \mu_m)$ and $\sigma^2 = (\sigma_1^2, ..., \sigma_m^2)$ (Nikolova et al., 2006). Figure 5.1 shows the projection of the unit hypercube and the path polytope in the (μ, σ^2) -plane.

The *dominant* of a set, C, is defined by Nikolova et al. as the set of all points

that are greater than a point in C, $\{x \in \mathbb{R}^m | x \ge y \text{ for some } y \in C\}$. To explain why the solution to (5.1) is found at an extreme point of the dominant of the path polytope shadow, let us follow the example given by Nikolova et al. (2006) and consider the relaxed version of (5.1). Let $z_1 = \mu x$ and $z_2 = \sigma^2 x$. We rewrite Equation (5.1) as the continuous system

$$\max \quad \frac{t - z_1}{\sqrt{z_2}}$$
s.t. $(z_1, z_2) \in \text{path polytope shadow } S.$

$$(5.2)$$

We denote the objective function as $f = \frac{t-z_1}{\sqrt{z_2}}$. In other words, $f(z_1, z_2)$ is the z-score of the route with $z_1 = \mu$ and $z_2 = +sigma^2 x$. For a subset $\overline{S} = S \cap \{z_1 | z_1 < t\}$ which we assume to be non-empty (i.e. there exists a path with mean less than t), the set $f(z_1, z_2)$ on this feasible subset must be positive since $z_1 < t$, and it must contain the maximum value of f as the subset contains all the feasible points (z_1, z_2) for which the mean, z_1 , is lower than t. Let us consider the level set $L_{\gamma} = \{z \in R^2 | f(z) \le \gamma\}$. This is the set of points (z_1, z_2) (i.e. paths) with z-scores lower than the given level, γ (γ can be regarded as a z-score) for some given deadline, t. Lower z-scores correspond to lower probabilities of route completion within the given deadline. L_{γ} consists of points (z_1, z_2) such that

$$\frac{t-z_1}{\sqrt{z_2}} \leq \gamma \quad \Longleftrightarrow \quad z_2 \geq \left(\frac{t-z_1}{\gamma}\right)^2. \tag{5.3}$$

For positive γ and $z_1 < t$, the level set L_{γ} is convex, as the right side of (5.3) is convex. The area *above* the convex function given by $(\frac{t-z_1}{\gamma})^2$ contains points (z_1, z_2) with lower z-scores than for points on or under the line, as illustrated by Figure 5.2.

To argue why the maximum is attained at an *extreme point* of \overline{S} , we need also provide the theorem stated by Nikolova et al. (2006).

Theorem 1 Let $C \in \mathbb{R}^m$ be a compact convex set. A quasi-convex function $f: C \to \mathbb{R}$ that attains a maximum over C, attains the maximum at some extreme point of C.

Hence, $f(z_1, z_2)$ is quasi-convex on \overline{S} , which is the part of the path polytope shadow to the *left* of $z_1 = t$, and the maximum is attained at an extreme point of \overline{S} . Nikolova et al. have also shown that the optimal solution of the continuous formulation is the same as the optimal solution to the discrete



Figure 5.2: The marked area contains points with lower z-scores than a given z-score, γ , for a given deadline, t.

problem, which is necessary for the method to be valid for the problem studied in this thesis.

To find the extreme points of the dominant of the path polytope shadow, a convex combination of the paths' means and variances is minimized. The algorithm can be thought of as a search algorithm on the values of the weight to be used in the convex combination (Nikolova, 2009). The objective function to be minimized looks as follows:

$$\min \alpha \mu + (1 - \alpha)\sigma^2. \tag{5.4}$$

The algorithm first sets the weight, α , equal to 1 and then to 0 and solves the two resulting single objective problems. This way, the method is initialized by first finding the two routes with the smallest mean and with the smallest variance. These two routes are then considered as points, P_1 and P_2 , in the two dimensional (μ, σ^2) -plane. This way, the first route equals the deterministic solution. Figures 5.3 and 5.4 illustrate how P1 and P2 are typically placed in the (μ, σ^2) -plane. Obviously, there exists no path such that its (μ, σ^2) combination places it to the *left* of P1 or *below* P2 in the (μ, σ^2) -plane, as these have the lowest mean and variance, respectively. In Figure 5.3, the area with grey lines illustrates the area containing feasible (μ, σ^2) combinations that can not exist based on the two initial points.



Figure 5.3: The dashed grey area contains feasible combinations of mean and variance which do not exist.

Next, the weight, α , is updated and set such that the slope of the linear objective function is equal to the slope connecting the points P_1 and P_2 . This implies that the weight, α , is set as

$$\alpha = \frac{slope}{slope - 1}.$$
(5.5)

The model is rerun with α as in Equation 5.5, and the new resulting path, if any, is denoted by P_3 . The search for more routes continues by searching for routes using the edge between P_1 and P_3 and between P_3 and P_2 , etc., until no more new routes are found. The updating of the weight makes the algorithm a logical search on the convex hull of the path polytope shadow, and the convex hull is expanded iteratively based on the extreme points found so far. When no new routes are found and the area can not be expanded any more, we have attained the convex hull and the set of points that constitute the dominant of the shadow path polytope. The optimal path can be found in this set of extreme points.



Figure 5.4: Finding and enumerating the candidate solution-paths (Nikolova, 2009).

5.2 Adaption of Nikolova's Method to our Problem

When the algorithm is adopted to the problem studied in this thesis, the extreme points of the dominant of the shadow path polytope are identified by minimizing a convex combination of the routes' mean and variance, as seen in the expression in (5.4). This expression becomes the objective function, and is subject to equations (4.2) - (4.20) in Chapter 4. The weight used to construct the convex combination, α , is updated as described above, always using the slope between the most recent new point and its adjacent points. To choose what adjacent point to use, the slope of the edge between the most recent point and its nearest neighbor to the *left* is always used to find the new weight and thereby the new objective function. When no new point is found, the search continues by using the edge between the newest point and its nearest neighbor point to the *right*. If a new point is found, one searches to the left again, and if not, the search continues to the right. Figure 5.5 illustrates the order in which extreme points of the path polytope shadow (i.e. routes) are found and shows the resulting order in which edges are used to set α .

The extreme points of the shadow path polytope are the candidate points for the optimal route, i.e. the *possibly* optimal routes. For each threshold within some range of thresholds, the probability of completing each route within that threshold is calculated. This way, the optimal route for each given threshold can be identified.

For the model presented in Chapter 4, the threshold, H, is a constant. It is of interest to examine confidence levels for the route completion time for various



Figure 5.5: The order for discovering routes and using edges between adjacent nodes to update the weight of the objective function and search for new routes. The red letters indicate the order in which extreme points of the path polytope shadow (routes) are found, and the black numbers indicate the order of edges used for setting α .

thresholds. In reality, the threshold (or deadline) is often a requirement from customers which must be adhered to. For any given route, there is a oneto-one correspondence between a threshold and the confidence level for route completion within that threshold. Note that in a business context, there is typically a lower limit for confidence levels to be accepted. Such limits are commonly rather high. To present the model output in an understandable way, the lower limit for thresholds for which each route's confidence level is calculated for is the threshold that gives a confidence level, β , of 0.05 for the route with the lowest mean (i.e. P1 from Nikolova's method). Let us for now label this confidence level as β_l , and the corresponding threshold as H_l . The upper limit for the range of thresholds used is the threshold that yields a confidence level, β , of 0.95 for the route with the highest mean (i.e. P2 from Nikolova's method). Let us for now label this confidence level as β_h , and the corresponding thresholds as H_h . For a given test instance, each identified route's confidence level is calculated for all threshold with an integer value between these two limits, H_l and H_h . Confidence levels of 0.05 are not interesting from a business perspective, but including these low confidence levels allows us to see the full picture. For an instance with 3 possible optimal routes, the range of thresholds used is illustrated in Figure 5.6.



Figure 5.6: The upper and lower limits for the range of thresholds for witch the confidence level for all identified routes are calculated for a given instance.

Implementation in MATLAB R2015a and Xpress-IVE

To solve the problem using Nikolova's method, the algorithm is implemented in MATLAB R2015a. MATLAB R2015a accesses Xpress-IVE Version 1.24.06 64 bit with Xpress Optimizer Version 27.01.02 to solve the problem in Chapter 4 each time a new objective function (i.e. a new weight, α), is used to search for more possible optimal routes. In other words, decisions regarding what adjacent point (i.e. route) to use when setting the weight, α , are handled in MATLAB R2015a. Every time we we search for a new point, we solve the mathematical model presented in Chapter 4 with an objective function as defined in Equation (5.4) with some weight, α , and this is done using Xpress.

Chapter 6

Probability Distributions of the Stochastic Variables

The nature of the stochastic sailing times of the arcs in the network and how they are handled is explained in more detail in this chapter. Section 6.1 explains the relation between the realised sailing times and the realised waiting times. Section 6.2 explains some aspects of the distributions associated with the stochastic sailing times of each arc. The real distributions of the sailing times are found by simulation. The (μ, σ^2) -characteristics of the real sailing time distributions are used to approximate the distributions used as input data to our model. They are also used to better understand the behaviour of the stochastic variables.

Recall that we denote the stochastic sailing time between nodes i and j as $\tilde{s}(\tilde{w})_{ij}$. This sailing time is *not* the same as the travel time, as the travel time also includes loading time and cleaning time of node j. The weight of an arc in the network equals the *travel* time. To decide the arc means and the variances, we need to separate the stochastic and the deterministic contributions to the travel times. As mentioned in Chapter 2, the stochastic sailing times (between *different* terminals) depend on the waiting time at the destination terminal, the direction from start node to end node relative to the direction from the nodes to anchorage, and the distances between terminals.

The speed of the ship is assumed to be 4.3 knots (Arnesen and Gjestvang, 2015), and the distances between terminals are known and fixed. The resulting direct sailing times (i.e. the time it takes to sail between terminals if the ship could sail directly to the next terminal) are given in hours. The waiting time at terminals is assumed to follow a normal distribution. The variance and expected value of the waiting time for each terminal is set to some value,

which is further explained in Chapter 7.

6.1 Realisations of Stochastic Sailing Times

The dependency between the sailing times, the waiting times and the direction of the arc is further explained in this section. Both the sailing times and the waiting times are stochastic. However, for a given *realised* waiting time, one can find the resulting *realised* sailing time. We show the expression for the realised sailing time with respect to the realised waiting time for arcs with a source node in one terminal and destination node in a different terminal. How the distributions of the stochastic sailing times are found and set using a set of randomly drawn realised waiting times is explained in Chapter 7. An explanation of the applied notation and concepts follows first. We then explain the relation between the realised sailing times and the realised waiting times when arcs are directed away from anchorage, before the same is explained for arcs directed towards anchorage.

In Figure 6.1 we consider sailing from Terminal A to Terminal B, and in Figure 6.2 we consider sailing from Terminal B to Terminal A. In Figure 6.1, the direction is *away* from anchorage, Terminal 0, while the direction of the arc from B to A in Figure 6.2 is *towards* anchorage. The unit we are concerned with is hours. Recall from Chapter 2 that Houston Ship Channel in the Port of Houston is long and narrow, and the movement of the ship can be regarded as movement along a straight line. Hence, the time it takes to sail along the vertical lines are considered to be zero. The grey arrow indicates the direction we would sail if there were no waiting times and we could have sailed directly to the next terminal. This is referred to as *direct sailing times*. The realised waiting time at Terminal A is denoted by w_A and for Terminal B by w_B . The realised sailing times resulting from the realised waiting times at the destination terminal are denoted as s_{AB} and s_{BA} . Because the ship is not allowed to wait at a terminal, but has to sail towards anchorage while waiting, s_{AB} and s_{BA} can take other values than the direct sailing times. As such, s_{AB} is the times it takes from the ship leaves Terminal A until it arrives at Terminal B. s_{AB} depends on direction and waiting time at the end node, w_B . Because w_B is stochastic, it makes s_{AB} stochastic as well. In Chapter 4, this stochastic variable is denoted as $\tilde{s}(\tilde{w})_{AB}$, but for simplicity, we use s_{AB} in this section.

The length of the arrow illustrating s_{AB} or s_{BA} only shows the direction and indicates that we are considering sailing from A to B, or from B to A. By adding the associated loading and cleaning times to s_{AB} and s_{BA} , we get the *travel time* along the arc, which is the weight the arc is given in the network. But for now, we are explaining the nature of the stochastic part of the travel times. In other words, cleaning and loading are ignored in this chapter and
are not included in s_{AB} or s_{BA} .

During the waiting time, the ship either sails towards anchorage, or it waits at anchorage. An arc can either be directed towards or away from anchorage, and this is illustrated in Figures 6.1 and 6.2, respectively. For these two figures, the direct sailing time between Terminals A and B is denoted by D_{AB} or D_{BA} , and D_{0A} is the direct sailing time from anchorage to Terminal A and D_{A0} from Terminal A to anchorage. Equivalently, D_{0B} is the direct sailing time from anchorage to Terminal B and D_{B0} is the direct sailing time from Terminal B to anchorage. These values $(D_{AB}, D_{BA}, D_{0A}, D_{A0}, D_{0B} \text{ and } D_{B0})$ are known and fixed for a given problem instance. Note that $D_{AB} = D_{BA}$, $D_{0A} = D_{A0}$, $D_{0B} = D_{B0}$ but the directions are opposite.

Sailing from Terminal A to Terminal B, away from anchorage

In Figure 6.1, we differ between two cases; the waiting time at Terminal B, w_B , can be less than or equal to the time it takes to sail from A to anchorage $(w_B \leq D_{A0})$, or w_B exceeds the time it takes to sail from Terminal A to anchorage $(w_B > D_{A0})$.

In the first case, if w_B is zero, the ship sails directly from A to B along the D_{AB} segment. If w_B is larger than zero but smaller than the time it takes to sail along the D_{A0} segment, the ship first sails towards anchorage from A and turns around before reaching anchorage (i.e. it turns around somewhere on the D_{A0} segment). Then it sails back along the part of the D_{0A} segment that has been travelled, passes Terminal A and sails to B along the D_{AB} segment. This gives $s_{AB} = 2w_B + D_{AB}$. Note that if there is no waiting time ($w_B = 0$) the ship still sails $2w_B + D_{AB}$, which then equals sailing directly from A to B.

In the second case, where $w_B > D_{A0}$, the ship sails to anchorage along the D_{0A} line, waits there and then sail all the way back to B. This gives $s_{AB} = w_B + D_{0A} + D_{AB}$. Equation (6.1) sums up the two possible cases for the values of the realised sailing time depending on the realised waiting time.

$$\mathbf{A} \xrightarrow{s_{AB}} \mathbf{B}$$

$$D_{0A} = D_{A0} \qquad D_{AB} = D_{BA}$$

Figure 6.1: Sailing time and waiting time dependency for arcs directed away from anchorage.

$$s_{AB} = \begin{cases} 2w_B + D_{AB}, & \text{if } w_B \le D_{A0} \\ w_B + D_{0A} + D_{AB}, & \text{otherwise.} \end{cases}$$
(6.1)

Sailing from Terminal B to Terminal A, towards anchorage

In Figure 6.2, we differ between three cases; the realised waiting time, w_A , at Terminal A can be less than or equal to the direct sailing time from B to A ($w_A \leq D_{BA}$), it can be higher than the direct sailing time from B to A but smaller or equal to the direct sailing time from B to anchorage ($D_{BA} < w_A \leq D_{BA} + D_{A0}$), or w_A can be larger than the direct sailing time from B to anchorage from B to anchorage ($w_A > D_{BA} + D_{A0}$).

In the first case, the realised sailing time equals the direct sailing time from B to A, D_{BA} . In the second case, the ship turns around somewhere between A and anchorage (on the D_{A0} segment), and realised sailing time becomes $2w_A - D_{BA}$ hours. In the last case, the realised sailing time is $w_B + D_{0A}$ hours. This connects the realised sailing time (excluding the deterministic time for loading/unloading/cleaning) to the realised waiting time. Equation (6.2) sums up the three possible cases.



Figure 6.2: Sailing time and waiting time dependency for arcs directed towards anchorage.

$$s_{BA} = \begin{cases} D_{BA}, & \text{if } w_A \le D_{BA} \\ 2w_A - D_{BA}, & \text{if } D_{BA} < w_A \le D_{BA} + D_{A0} \\ w_A + D_{0A}, & \text{otherwise.} \end{cases}$$
(6.2)

6.2 Distribution of the Stochastic Sailing Times

The previous section clarifies how different sailing directions relative to anchorage and the realised waiting times result in different realised sailing times between two terminals. In this section, the distributions of the stochastic variables are examined, and the resulting behaviour and effects this gives are explained. Keep in mind that we are still only addressing arcs between cargoes in *different* terminals, not between cargoes in the same terminal. If we let σ_{AB}^2 represent the variance of the arc from A to B and σ_{BA}^2 the variance for the arc from B to A, σ_{AB}^2 and σ_{BA}^2 are different from each other in most cases. There are especially two effects, the Pile-Up Effect and the Anchorage Effect, on which the variance of an arc depends. The unique combination of the distances between terminals, directions of the arcs, the distributions of the waiting times, and the distances to anchorage are crucial for how and to what degree the Pile-Up Effect and the Anchorage Effect contribute to the arcs' variances. These effects and conditions are explained in more detail in this section.

6.2.1 The Pile-Up Effect

Consider Figure 6.3 with Terminals A and B. The waiting time at each terminal follows independent normal distributions. Let us first consider an arc from Terminal B to Terminal A, and let us denote the arc (B, A). This is illustrated in Figure 6.3. The full line between Terminal B and A in Figure 6.3 illustrates the *direct sailing time* between Terminal B and A. The dashed line illustrates possible additional sailing time, and the dashed circle illustrates possible waiting at anchorage which happens for large enough waiting times.



Figure 6.3: Possible outcomes for arc (B, A), which is directed towards anchorage.

For this arc, (B, A), the ship initially sails in the direction of the destination, Terminal A. For all realised waiting times less than or equal to the time it takes to sail along this full line, i.e. the direct sailing time, the resulting realised sailing times are equal to each other, and their value is the direct sailing time. This gives a lower limit for possible realised sailing times, and suggests a truncated distribution of the sailing times. For this distribution, sailing times equal to the direct sailing times between B and A are piled up. A higher probability for the waiting time to be equal to or less than the direct sailing time gives a higher amount of realised sailing time equal to the direct sailing time from B to A, i.e. a higher pile. For arcs directed towards anchorage, the accumulation of sailing times with value equal to the direct sailing time depends on the size of the direct sailing time between the two terminals relative to the distribution of the waiting time at the destination terminal.

If the waiting time at A exceeds the time it takes to sail directly from B to A, the ship passes Terminal A and sails away from it and towards anchorage along the dashed line. Additional sailing time that would not have been necessary if there was no waiting time then applies. Hence, waiting times exceeding the direct sailing time contribute to a higher mean of the distribution of the sailing time. A larger difference between the direct sailing time and the mean of the distribution of the sailing times causes the pile of realised sailing times equal to the direct sailing time to contribute with more variance.

This means that the distribution of the sailing times of the arcs directed towards anchorage are truncated and piled up. This affects the variance associated with the arc. The effect this has on the variance depends on the distribution of the waiting time relative to the direct sailing time between the terminals. The distance to anchorage also matters as it affects how large the mean of the distribution of the sailing time can be. For low enough probabilities for the ship to pass the destination terminal, the pile becomes high and the variance low. For higher probabilities for the ship to pass the destination terminal, the mean of the sailing time increases. The higher the mean, the larger the variance as the pile gives more variance. This effect is henceforth referred to as the *Pile-Up Effect*.

Note that as negative waiting times is not realistic, the distributions of the *waiting* times are truncated at zero. For this reason, arcs directed *away* from anchorage will also have a pile for realised sailing times equal to the direct sailing times. This pile is smaller than for the symmetric arc, but the effect this pile has on the variance resembles the Pile-Up Effect.

6.2.2 The Anchorage Effect

In the above section, we explain how the Pile-Up Effect works for arc (B, A) in Figure 6.3 directed towards anchorage. The symmetric arc (A, B), directed towards anchorage, is illustrated in Figure 6.4. In this subsection we examine another effect that influences the variance of the stochastic sailing time of an arc, the *Anchorage Effect*. The Anchorage Effect applies both to arcs directed towards anchorage and to arcs directed away from anchorage.

For arc (A, B) the ship starts sailing away from the destination terminal from the very start and additional sailing time applies at once for waiting times larger than zero. For both arcs (B, A) in Figure 6.3 and (A, B) in Figure 6.4, the ship turns around before reaching anchorage if the waiting times are short enough. The dashed circle illustrates possible waiting at anchorage which happens for large enough waiting times. For the ship to actually reach



Figure 6.4: Possible outcomes for arc (A, B), which is directed away from anchorage.

anchorage, the waiting time must be large enough relative to the sailing time to anchorage.

Let us consider situations where the ship sails towards anchorage and away from the destination node. For each unit of distance travelled before the ship reaches anchorage, it must eventually sail this same distance in the opposite direction when it sails back to reach the destination. This implies that the *contribution* to the realised sailing time is the double of the units of waiting time spent somewhere between the anchorage and the closer of the two terminals to anchorage (Terminal A in the figures). For the amount of time spent waiting at anchorage, the contribution to realised sailing time is not doubled in the same way, as the ship stands still. This means that the realised sailing times are more spread out for lower values of realised sailing times, and less spread out (i.e. more concentrated) for higher values of realised sailing times. This indicates a more right skewed distribution for cases where the waiting times are distributed such that the ship reaches anchorage more often. For two symmetric arcs, when ignoring the possible difference in the distributions of the waiting times at the two destination terminals, the ship ends up at anchorage more often for the arc directed away from anchorage than for the arc directed towards anchorage. This is because the departure terminal is closer to anchorage for the arc directed away from anchorage than for the symmetric arc.

How the skewness affects the variance of the distribution of the sailing times depends on several conditions. The higher the probability that the ship will end up at anchorage, the more right skewed will the distribution of realised sailing times be. When the mean is to the left of the concentrated, high realised sailing times, then the concentrated sailing times contribute to more variance than when the mean coincides with the values of the concentrated sailing times. In other words, *where* the frequent sailing times are concentrated relative to the mean affects the contribution that the Anchorage Effect has on the variance.

The Pile-Up Effect only applies to arcs directed towards anchorage, and contributes to left skewness, while the Anchorage Effect applies to arcs in both directions and contributes to right skewness. As such, for arcs directed towards anchorage, the Pile-Up Effect and the Anchorage Effect may interfere. If the two effects contribute to lower or higher variance for the sailing time distribution of an arc depends on the relative sizes of the distances between terminals, distances to anchorage and the distribution of the waiting times. Longer distances to anchorage mitigates the Anchorage Effect (i.e. less right skewness). Longer distances between terminals increases the Pile-Up Effect (i.e. more left skewness).

6.2.3 Simulation of Arc Distributions

To illustrate the Pile-Up Effect and the Anchorage Effect, we do a simulation of realised sailing times of the two artificial, symmetric arcs (A, B) and (B, A)from Figures 6.3 and 6.4. We let the direct sailing time between Terminals A and B be 1 hour. Terminal A is located closer to anchorage than Terminal B, and (A, B) is directed away from anchorage while (B, A) is directed towards anchorage. The direct sailing time from anchorage to Terminal A is 7 hours, and from anchorage to Terminal B the direct sailing time is 8 hours. The distributions of the waiting times at Terminals A and B are $w_A \sim N(3.0, 1.5^2)$ and $w_B \sim N(5.0, 2.5^2)$, respectively. These characteristics are summed up in Figure 6.5.



Figure 6.5: Characteristics of the port used to simulate and explain the distribution of arc sailing times.

From the distribution of the waiting time at the destination terminal, we draw 10,000 realised waiting times at random and calculate the resulting sailing times using Equations (6.1) and (6.2). As waiting times are normally distributed, there is a left tail where negative values may be drawn from. This is not realistic, thus if a sample is negative, it is set to zero. In other words, the distribution of waiting times is truncated. Figure 6.6 shows the resulting distributions of the realised sailing times for the two symmetric arcs, (A, B) and (B, A), and Table 6.6c summarizes the characteristics of the distributions. Note that the loading and cleaning times are not included as they are not a part of the stochastic sailing time.

From Figure 6.6 we see that both distributions are truncated, as predicted. The truncated distribution of (B, A), shown in Figure 6.6a, is partly due to





(a) Skewed distribution of the realised sailing times of arc (B, A), directed towards anchorage.

(b) Skewed distribution of the realised sailing times of arc (A, B), directed away from anchorage.

	(B,A)	(A,B)
Sample set μ	5.1 hours	10.7 hours
Sample set σ^2	2.76^{2}	4.46^{2}
Sample set skewness	0.4 (left skewed)	-0.3 (right skewed)

(c) Characteristics of the distributions of realised sailing times for arcs (B, A) and (A, B), directed towards and away from anchorage, respectively.

Figure 6.6: Distributions of realised sailing times for two symmetric arcs given realised waiting times at the destination terminals. The table summarizes the characteristics of the distributions.

the truncated distribution of the waiting time, but mostly due to the Pile-Up Effect explained above. The truncated distribution of (A, B), shown in Figure 6.6b, is only due to the truncated distribution of the waiting times, as the Pile-Up Effect does not apply to arcs directed away from anchorage.

As discussed, the Pile-Up Effect contributes to a more left skewed distribution of sailing times, and the Anchorage Effect contributes to a more right skewed distribution of sailing times. As the probability of waiting at anchorage is higher for an arc directed away from anchorage than for the corresponding symmetric arc (given equal waiting time distributions at terminals), arcs directed away from anchorage are typically right skewed. However, when distances to anchorage are small relative to the waiting times at terminals, arcs directed towards anchorage may also become right skewed. Whether arcs are left skewed or right skewed, and to what degree, depends on several conditions, and the Pile-Up Effect and the Anchorage Effect may interfere with each other.

From Table 6.6c we see that (B, A) is left skewed with a skewness of 0.4, while (A, B) is right skewed with a skewness of -0.3, implying that the Anchorage

Effect is at play for arc (A, B), directed away from anchorage. Figure 6.6b shows that the distribution of sailing times for arc (A, B) has a second peak to the right of the mean, at approximately 16 hours. It is for these high sailing times that the ship has waited at anchorage. Recall that when the ship waits at anchorage the realised sailing time only increases by one unit for each extra unit of waiting time, as opposed to an increase of two units of sailing time for each unit of waiting time when the ship sails toward anchorage. For arc (B, A), directed towards anchorage, the distribution is right skewed and it is the Pile-Up Effect that affects the variance the most. As (B, A) has a very low probability of ending up at anchorage, the Anchorage Effect does not contribute much to the variance.

Both the mean and the variance of the sailing time distribution of arc (A, B) is much higher than for (B, A), implying that for this example the Anchorage Effect contributes to more variance of sailing times distributions than the Pile-Up Effect does. Which of the effects contributes to more variance depends on the distribution of waiting times at the terminals and the relative distances between the terminals, and the anchorage and the terminals. Note that for other conditions, the stronger Pile-Up Effect or Anchorage Effect may give lower variance.

From Figure 6.6 we see that the sailing times of the arcs are not normally distributed, but parts of the distributions resemble normal distributions. Most relevant literature assume normal distributions and the contributions to the problem at hand when considering skewed and truncated distributions instead of normal distributions are minor relative to the complications it brings. For these reasons, we assume the distribution of the sailing time of each arc to follow an unskewed normal distribution. The mean and standard deviation of the normal distribution are set equal the mean and standard deviation of the real distribution of the arc's sailing time found by simulation. One must keep in mind that the distributions of the sailing times are found assuming normally distributed waiting times, which is not necessarily the case. However, due to incomplete data for real waiting times, and for simplicity, we assume normally distributed waiting times for the remainder of this thesis.

6.2.4 The Pile-Up Effect and Anchorage Effect for Various Port Characteristics

As explained, the Pile-Up Effect and the Anchorage Effect give rise to variations in the uncertainty of different arc sailing time distributions, which motivates the search for routes with lower variance than the deterministic route. To illustrate how the Pile-Up Effect and the Anchorage Effect may contribute to the variance of the distributions of arc sailing times, the realised sailing times for varying port characteristics that impose the two effects are simulated.

Table 6.1 shows the Pile-Up Effect for arc (B, A) in Figure 6.7 directed towards anchorage. The characteristics of the port are as shown in Figure 6.7, except for the direct sailing time between Terminal A and Terminal B, D_{AB} , which we give a different value for each distribution. All other characteristics are held constant. By changing D_{AB} , and holding the distribution of the waiting time at Terminal A constant, the probability that the ship sails directly to Terminal A changes. The sailing time for (B, A) is simulated with five different values of D_{AB} , and these values are 0.5 hours, 1 hours, 2 hours, 3 hours, and 5 hours. The waiting time at terminal A is normally distributed with a mean of 3 hours and and a standard deviation of 1.5 hours. As the waiting time at Terminal A is constant, a higher D_{AB} increases the probability for the ship to sail directly to Terminal A without passing the terminal. We see from the characteristics of the distributions in Table 6.1 that variance of the realised sailing time distributions decreases with increasing D_{AB} . This is because the size of the pile increases and the mean is shifted horizontally towards the pile. Note that for different port characteristics the Pile-Up Effect could behave differently, especially if the Anchorage Effect comes into play as well. For the examples in Table 6.1 the Anchorage Effect is minimal due to the high sailing time to anchorage.



Figure 6.7: Characteristics of the port used to simulate the Pile-Up and Anchorage Effect.

The Anchorage Effect is illustrated in Table 6.2. The distributions shown are for realised sailing times of arc (A, B) directed away from anchorage. Except from the direct sailing time between anchorage and Terminal A, D_{0A} , the port characteristics are equal for all distributions, and the characteristics are shown in Figure 6.7. When D_{0A} changes, the probability of to ship to wait at anchorage changes as well. The first distribution in Table 6.2 is for $D_{0A} = 10$ hours, where the probability of having to wait at anchorage is quite low as the waiting time at Terminal B is normally distributed with a mean and standard deviation of 5 and 2.5 hours, respectively. For the other distributions, D_{0A} decreases such that the probability of having to wait at anchorage increases, as everything else is held constant. From the table, we see how the variance of the sailing time distribution decreases for increased anchorage effect. As the arc is directed away form anchorage, the Pile-Up Effect is not at play. The pile of values at the lower limit is due to the fact that drawn waiting times less than zero are set equal to zero in the simulation. We see from the characteristics reported in Table 6.2 that the distributions are all right skewed. The skewness is first low, then high, and then low again. Whether the Anchorage Effect and the Pile-Up Effect give more or less variance depends on the specific combination of various conditions.



Table 6.1: The Pile-Up Effect for arc (B, A) directed towards anchorage.



Table 6.2: The Anchorage Effect for arc (A, B) directed away from anchorage.

6.3 Summary

The direction of an arc relative to anchorage as well a the relative sizes of the distances and the waiting times affect both the realised sailing times and what variance is associated with each arc. In reality, the distributions of the sailing times are skewed and truncated. The linear movement and the possibility of having to wait at anchorage give rise to the Pile-Up Effect and the Anchorage Effect. The Pile-Up Effect applies to arcs directed toward anchorage, while the Anchorage Effect applies both to arcs directed toward anchorage and to arcs directed away from anchorage. In spite of the skewed and truncated characteristics of the real distributions of arc sailing times for simplicity. To approximate reality as much as possible, we use the mean and standard deviation of the real distributions of arc sailing times, found by simulation as described in Section 6.2.3, as the mean and standard deviation of the normal distributions of arc sailing times used in the model.

Chapter 7

Input Data and Generation of Test Instances

In this chapter, we explain how input data is generated and test instances created. The input data used is presented in Section 7.1. The data is based on realistic data from the case port, Houston Ship Channel, and the case company, Odfjell SE. We also explain in more detail how some of the additional parameters necessary for the test instances are decided. In this thesis a set of 100 instances is used as a basis for the analysis and as a reference group for the sensitivity analysis. This set is referred to as the *Base Set*. A summary of the generated test instances is presented in Section 7.2 at the end of this chapter.

7.1 Input Data

The input data is based on numbers for Houston Ship Channel, Houston, Texas, which is the case port. In this port, Odfjell's chemical tankers can visit a number of terminals and may service several customers at each terminal. Figure 7.1 shows a map of Houston Ship Channel with the location of 11 terminals, here referred to as Terminals A-K. The anchorage is located at sea outside the port, 22 nautical miles from Terminal A.



Figure 7.1: Terminals in Houston Ship Channel used in the test instances (Arnesen and Gjestvang, 2015).

Direct sailing times

As can be seen from Figure 7.1, the terminals are located along the narrow port channel. To be able to analyse the effects this narrow geography gives rise to, we assume the terminals are distributed along a straight line, and all ship movement can be regarded as movement along this line. Figure 7.2 shows Terminals A-K the way they are modelled in the test instances. By using Google Maps Distance Calculator, we estimate the distance between the different terminals. Hence, the distances may deviate slightly from the true distances, but are close enough to the real values to service the purpose of this thesis. Approximated direct sailing times, i.e. the time it takes for a ship to sail the distance between the terminals, are then calculated using the estimated distances and a sailing speed of 4.3 knots, as used by Arnesen and Gjestvang (2015). All test instances generated in this thesis have six terminals. The direct sailing times between the terminals shown in Figure 7.1 are presented in Figure 7.2.



Figure 7.2: Terminals in Houston Ship Channel the way they are modelled in the test instances. The numbers show the approximated direct sailing time in hours between each pair of adjacent terminals.

Geography

Depending on where the customers require their cargo to be loaded/unloaded, the ship has to visit different terminals on different port calls. The geography of the routing problem depends on the location of the customer requests to be serviced. To illustrate this, we differentiate between three types of geographies, namely *Even*, *Far Away* and *Split*. For the Even geography, all terminals and the anchorage are spread out evenly along a line segment. This is illustrated in the upper figure in Figure 7.3. For the Far Away geography, the terminals are evenly distributed, but the anchorage is located further away than in the Even geography. This is illustrated in the middle figure in Figure 7.3. For the third geography, Split, at least one of the terminals is placed closer to the anchorage than to the closest of the remaining terminals, causing a gap/split. This is illustrated in the bottom figure in Figure 7.3.



Figure 7.3: The three different geographies, Even, Far Away, and Split.

The terminals in Houston Ship Channel do not allow us to obtain perfectly distributed locations of the terminals, so the set of terminals used for each of the geographies are picked such that the chosen set of terminals approximate the geographies mentioned above. For this thesis, instances with six terminals are used for testing. Hence, for each geography, six of the Terminals A-K are picked to approximate the geographies as good as possible. Terminals A, B, D, G, I and K are used for the Even geography, Terminals F-K for Far Away geography, and A, B, H, I, J and K for Split geography. In reality, the main anchorage used for Houston Ship Channel is located approximately 22 nautical miles away from the terminal labeled as terminal A in Figure 7.1. This corresponds to a sailing time of 5.12 hours for a speed of 4.3 knots. As seen from Figure 7.2, this direct sailing time is used. Note, however, that for the Far Away geography, terminal F is the terminal closest to anchorage. The sum of the sailing time along the arcs from anchorage to terminal F in Figure 7.2 is 7.79 hours, and this is the distance used for the Far Away geography. The information regarding which terminals are being used for the different geographies is summarized in Table 7.1.

		3				
Even	A	В	D	G	Ι	K
Far away	\mathbf{F}	G	Η	Ι	J	Κ
Split	А	В	Η	Ι	J	Κ

Table 7.1: Terminals from Figure 7.1 used to approximate the geographies Even, Far Away and Split.

For the Base Set used as a basis for analysis and as a reference group when analysing different variations in parameters, the Even geography is used.

Mean and Variance of the Stochastic Waiting Times

To be able to analyse the impact of the mean and the standard deviation of the distribution of the stochastic waiting times, we create test instances with varying values for these two variables. For a given set of data for realised waiting times for the terminals in Houston which Odfiell's chemical tankers can visit, the average waiting time is in reality approximately 36 hours. For this given set of data, the waiting times vary quite a lot, with a standard deviation of approximately 2.15 days (more than 51 hours). The waiting times of this data set range from 0 hours to approximately 366 hours (15.26 days). However, waiting times of two days and more is typically due to complications beyond regular waiting times, and we do not consider these data points as representative. About 80% of the realised waiting times for this data set are less than 1.75 days (42 hours), and only this part of the data set is used to set the parameters. 50% of these 80% of the data points lower than 42 hours are 0 hours or almost 0 hours, and the average of these results in a mean of approximately 10 hours. The standard deviation is 13.0 hours. Figure 7.4 shows the distribution of the realised waiting times as given by the data set provided by representatives for Odfjell.

We see that the waiting times are not normally distributed. However, this is only based on a limited amount of data. As relevant literature commonly assume normally distributed arc travel times, and for simplicity, we also assume normally distributed waiting times and sailing times in spite of the data shown in Figure 7.4.

We let the mean of the waiting times be drawn at random between 0 and 10 hours. For each instance, these are drawn for each terminal independently of other terminals used for the same instance. In other words, the various terminals used in the same instance have different distributions of the waiting



Figure 7.4: Distribution of the realised waiting times for a collection of terminals at Houston Ship Channel, based on a set of data received from Odfjell SE.

times. The standard deviation of the data less than 42 hours is, as mentioned, 13.0 hours, which corresponds to 130% of the mean of 10 hours. However, such a large standard deviation gives a too thick left tail for our purposes. This is because the left tail covers some negative values for waiting times, which is not possible in the real case. Instead, to assign variances for the waiting time for a given terminal, we let the standard deviation be set to 50% of the respective mean. In addition, the distributions of the waiting times are truncated at zero.

Ship and Cargo Information

The ship types used by the case company for chemical shipping have capacities ranging from 4,000 to 50,000 dead weight tonnes. For this thesis, the same total ship capacity applies for all test instances. However, we let draft limits and cargo sizes be expressed as percentages of the total ship capacity, implying that the results are representative for more than one type of ship. The way draft limits vary with ship capacity is explained in the next paragraph, and the following paragraph explains how cargo sizes vary depending on ship capacity.

All test instances generated for this thesis have 20 cargoes. The sizes of the cargoes are set such that the larger of the total pickup load and the total delivery load each corresponds to 80%-90% of a fully loaded ship. In other words, the total pickup load is drawn at random from 80% to 90% of total ship capacity, and the total delivery load is drawn at random from 80% to 90% of total ship capacity. This way, the ship capacity can be binding. In reality, the ship capacity is often a binding constraint, and using larger, non-binding ship

capacities is not interesting for practical applications.

For a given number of cargoes, some cargoes are to be picked up and others are to be delivered. The integer number of pickups is picked at random between $\frac{2}{5}$ and $\frac{3}{5}$ of the total number of cargoes, and the remaining cargoes are deliveries. The indexes are distributed randomly between pickup and delivery cargoes. This way, the difference between the amount of pickups and deliveries is not too big, but can vary a bit. The loads of the pickups are set at random, but such that their sum equals the total pickup load. The same goes for the delivery cargoes, whose sum equals the total delivery load.

We let each terminal used for an instance contain at least one cargo. The cargo indexed as 1 is always located in the first terminal, and the cargo with the highest index is always located in the last terminal (the terminal furthest away from anchorage). The remaining cargoes are also ordered with their index increasing along the terminals.

Draft Limits

In the real problem faced by the case company the draft limits are usually binding at some of the terminals. Thus, we wish to examine instances where draft limits may be restrictive. In this thesis we express the draft of the ship as the loaded weight on board the ship in tonnes. As such, draft limits at the terminals are also expressed in tonnes, and give the maximum load on board the ship for the ship to be able to enter and leave the terminal. At least one terminal must have large enough draft limit for the ship to be able to enter the port with total delivery load on board, and leave the port with total pickup load on board. For this reason, one or two terminals are drawn at random from the set of terminals, and their draft limits are set to 120% of total ship capacity. The draft limits for the remaining terminals are chosen randomly between 50% and 90% of total ship capacity. This does not, however, ensure that there exists feasible routes for a given instance. For the Base Set, we make sure that there exists at least one feasible route for each instance. A discussion of the impact of various draft limits follows in Chapter 8.

Assigning Means and Variances to the Distributions of the Stochastic Sailing Times

Recall from Chapter 6 that the distributions of the stochastic travel times must be approximated. The simulations used to construct the travel time distributions inputted to our model is explained here. The sailing time is the time it takes to sail from one cargo location to another excluding cargo locading and cleaning times. Adding cleaning and loading times is discussed in the next subsection, and only the stochastic sailing times are discussed here. As explained in Chapter 6, when the two cargoes associated with an arc are located in different terminals, the sailing time is stochastic. The stochastic sailing times depend on the stochastic waiting times at the destination terminal. The relation between the sailing times and the waiting times is also influenced by sailing directions relative to anchorage and the size of the direct sailing times relative to the size of waiting times. This relation is explained in Equations (6.1) and (6.2) in Chapter 6.

As explained above, the terminals are assigned independent normal waiting time distributions where the mean of the distributions are randomly drawn between 0 and 10 hours, and the standard deviation is set to 50% of the associated mean for a given distribution. For a test instance with six terminals, as used for the test instances in this thesis, there is a set of six independent waiting time distributions. To decide the distribution of the sailing times, we begin by drawing 10,000 sets of realised waiting times at random from the set of waiting time distributions. In other words, we get 10,000 sets where each set includes six realised waiting times for an instance with six terminals. For each set of realised waiting times, the corresponding realised sailing times can be calculated using Equations (6.1) and (6.2) for each arc. This results in 10,000 realised sailing time matrices. Each element (i, j) in a realised sailing time matrix represents the realised sailing time from cargo i to cargo j for a given set of realised waiting times. The mean of the approximated normal distribution of the sailing time for a given arc used as input data to our model equals the mean of the 10,000 sailing time samples for this given arc.

The same goes for the standard deviation used to construct the approximated normal distribution of an arc; to decide the standard deviation of arc (i, j), we use the standard deviation of the 10,000 sailing time samples for that arc as the standard deviation of the distribution assigned to the arc. As the means and standard deviations we get from this approach are used to construct the normal distributions of the sailing time of each arc, this is an approximation. As explained in Chapter 6, the real sailing times follow truncated, skewed distribution even if the waiting times follow unskewed normal distributions. This skewness is ignored when this approach is used. A discussion on and evaluation of this approximation is presented in Chapter 8.

Cleaning and Loading Time

To obtain the stochastic travel time, \tilde{t}_{ij} , the loading and cleaning times are added to the expected sailing time, $\tilde{s}(\tilde{w})_{ij}$, for each arc. How the distribution of $\tilde{s}(\tilde{w})_{ij}$ is constructed is explained above. As shown in Chapter 4, the travel time can be expressed as

$$\tilde{t}_{ij} = \tilde{s}(\tilde{w})_{ij} + l_j + c_j$$

The time it takes to load or unload a cargo depends on the type and size of the cargo. For the case company, loading or unloading of cargoes can take from 0.5 to 14 hours. Hence, each cargo, regardless of if it is a pickup or delivery cargo, is assigned a random loading time between 0.5 and 14 hours. The tanks used by the case company take from 3 to 15 hours to clean after a cargo has been discharged. Cleaning of tanks only applies after cargoes are delivered. The cargoes that are to be delivered are hence assigned an additional random cleaning time between 3 and 15 hours (Arnesen and Gjestvang, 2015). The varying cleaning times are used to illustrate different cleaning and disposal requirements for different types of cargoes.

Risk Profile of the Decision Maker

There is a one-to-one correspondence between the confidence level and threshold for a given normal distribution of a route's completion time. This implies that from a set of routes with different distributions, which one is optimal depends on both the threshold and the confidence level. In this thesis, the optimal route of a set of routes is defined as the route with the lowest threshold corresponding to a given confidence level. In a business context, reliability is important in customer relations, and for a business to perform well, planning ahead is crucial. As such, we assume a relatively risk averse decision maker, and set the required confidence level for a route to be optimal to 95%.

7.2 Summary of the Test Instances

In this thesis, a set of 100 test instances, the Base Set, is used as a basis for the analysis. The generated test instances are based on real values for the case company in Houston Ship Channel in the Port of Houston. All instances have 20 cargoes and six terminals. The true distances between terminals in Houston Ship Channel are approximated, and a speed of 4.3 knots is used to obtain the direct sailing time between each terminal. We differentiate between three types of geographies, Even, Far Away and Split, where different terminals in Houston Ship Channel are used to approximate each geography. The Base Set has an Even geography. The distribution of the waiting time at each terminal is independent of the distribution at other terminals, and is assumed to be normal. The means of the distributions of the waiting times are randomly drawn from 0 to 10 hours, and the standard deviation is set to 50% of the corresponding mean. The total load to be picked up and delivered are both randomly set between 80% and 90% of total ship capacity. The total pickup and delivery loads are distributed randomly among the pickup and delivery cargoes. The cargoes are randomly distributed among the terminals such that at least one cargo is located in each terminal. To make sure the ship can leave and enter the port, one or two terminals are randomly chosen and given draft limits of 120% of total ship capacity. The draft limits at the remaining terminals vary randomly from 50% to 90% of total ship capacity. The distribution of the stochastic sailing times are constructed by calculating the realised sailing times resulting from 10,000 realised waiting times drawn at random from the distribution of the waiting time at the corresponding terminal. The mean and standard deviation of these 10,000 realised sailing times of each arc are used to construct the normal distribution of each arc of the test instance. This is an approximation that ignores the skewness of the distribution of the true realised sailing times. To obtain the travel times, loading times and cleaning times are added to the sailing times. The loading times are randomly drawn between 0.5 and 14 hours and are assigned to both pickup cargoes and delivery cargoes. Cleaning times only apply for delivery cargoes, as cleaning must be conducted after discharge, and vary from 3 to 15 hours. Table 7.2 summarizes the most important input values used to generate the Base Set. The analysis presented in Chapter 8 is based on results from the Base Set. Variations in some of the values are also tested and the results are discussed Chapter 8.4.

Explanation of data input	Values			
Geography	Even			
Number of cargoes	20			
Number of terminals	6			
Range for total pickup/delivery as percentage of total ship capacity	80% - 90%			
Range for draft limits as percentage of total ship capacity	50% - $90%$			
Range for mean of waiting time	0 - $10~{\rm hours}$			
SD as percentage of mean	50%			
Range for loading time	0.5 - $14~\mathrm{hours}$			
Range for cleaning time	3 - $15~{\rm hours}$			
Confidence level for choice of optimal route	95%			

Table 7.2: Summary of input data used to generate the Base Set of test instances.

Chapter 8

Computational Study: Static Version of the Problem

In this chapter, we analyse the results of the static version of the problem and how the solutions are influenced by uncertainty. The results are obtained using the solution method explained in Chapter 5 and the Base Set of 100 test instances generated as described in Chapter 7. For each test instance, the solution method is used to identify a set of possibly optimal routes a priori. The routes are compared to each other and to the optimal deterministic route. For test instances where more than one route is identified, it is not always obvious what is the optimal route. Thus, the model output must be processed and analysed further for a decision maker to be able to use it as decision support.

Section 8.1 presents and explains the output of the solution method and gives some guidance on how to interpret the results. The results from testing the Base Set of 100 instances are aggregated and analysed in Section 8.2. In Section 8.3, some of the instances are analysed in more detail, and the effect of the aspects discussed in Chapter 6 are analysed. A sensitivity analysis is presented in Section 8.4, where variations in aspects such as geographic characteristics, draft limits, and risk profiles of the decision maker are examined. What is learnt about the behaviour of the stochastic variables is summarized Section 8.5. An evaluation of the approximation used to model the distributions of the stochastic sailing times is presented in section 8.6, before the practical implications and the value of accounting for uncertainty is assessed in Section 8.7. The solution algorithm is implemented in MATLAB R2015a. The optimization model is implemented using Xpress Mosel Version 3.8.0 as the modelling language, and is run using Xpress-IVE Version 1.24.06 64 bit with Xpress Optimizer Version 27.01.02. The computer used runs Windows 7 Enterprise 64-bit Operating System with Intel(R) Core(TM) i7-3770 CPU @ 3.40 GHz, and has installed memory of 16 GB. All test instances are created using MATLAB R2015a.

8.1 Model Output and Interpretation

In this section we present and explain the output we get when we use the solution method explained in Chapter 5. The aim is to give the reader an introduction to how information can be extracted from the employed analysis tools and model outputs.

By using the solution method explained in Chapter 5, a set of possibly optimal routes for each test instance is found. These routes can be characterized by the mean and standard deviation, or variance, of the completion time of the route, (μ, σ^2) .

The first route found for a given test instance, labeled as Route 1 in graphs and tables throughout this chapter, is always found by minimizing the *expected completion time*, and no attention is paid to the *variance*. Route 1 is thus always equal to the optimal *deterministic* route of the instance. As explained in Chapter 5, the rest of the set of possibly optimal routes are found by minimizing a linear combination of the mean and variance of the expected completion time. Thus, the routes labeled Route 2 and higher are referred to as *stochastic* solutions or routes.

For a given confidence level, there is one optimal route. The set of possibly optimal routes found for a test instance has to be further analysed to be able to find the optimal route for a given confidence level. What confidence level should be applied depends on the risk profile of the decision maker. For the Base Set of 100 instances, a confidence level of 95% is assumed. A discussion on different required confidence levels is presented in Section 8.4.

An example of a set of possibly optimal routes found using the chosen solution method is shown in Figure 8.1. The figure presents the (μ, σ^2) -characteristics of the routes found for one of the instances in the Base Set. The characteristics of the three routes identified using Nikolova's method are plotted in Figure 8.1a, each point representing a route. Figure 8.1b shows the cumulative distribution function (CDF) of the routes. Route 1 is, as always, the optimal deterministic route. Route 2 and Route 3 are stochastic routes. Figure 8.1b allows us to examine the prescribed level of confidence, α , by which



the completion time of the route is less than a threshold, H.

(a) The (μ, σ^2) -characteristics of the completion time of the three routes found for an instance.

(b) Cumulative distribution function (CDF) of the completion times of the three possibly optimal routes found for an instance.

Figure 8.1: Example of model output for a test instance from the Base Set. Three routes are identified as possibly optimal routes.

From Figure 8.1b we see that for every threshold shown in the graph, either Route 1 or Route 3 gives the highest confidence level, i.e. the highest probability of completing within the threshold. Route 2 is thus dominated by Routes 1 and 3 for the thresholds shown in the graph. Recall the one-to-one correspondence between the confidence level and the threshold for a given combination of route mean and variance. Thresholds, or deadlines, often represent requirements imposed by customers. Confidence levels often represent requirements set internally by the case company based on their risk profile.

Figure 8.1b can be read both ways. For a given required confidence level, we can see which of the routes is the optimal one and the associated threshold the route should complete within. Let us say a confidence level of 0.9 is required. Moving horizontally from left to right from $\alpha = 0.9$ on the y-axis, the first graph that intersects the line for $\alpha = 0.9$ is the graph representing Route 3 (the red graph). That is, for $\alpha = 0.9$, Route 3 has the lowest corresponding threshold. Moving downwards vertically from the point where the line representing $\alpha = 0.9$ intersects the graph representing Route 3, we hit the x-axis and find that for Route 3 to provide a confidence level of 0.9, the threshold can be no less than 460 hours. For lower thresholds than 460 hours, Route 3 offers lower confidence levels for route completion within the threshold than $\alpha = 0.9$.

Reading the figure the opposite way, we start with a given threshold or deadline that must be adhered to, and move upwards vertically to find the route that can provide us with the highest confidence level. For a deadline of e.g. 460 hours, Routes 1 and 3 both give confidence levels close to $\alpha = 0.9$, but Route 3 can provide a slightly higher confidence level for this threshold. Thus, Route 3 is the optimal choice of route for a threshold of 460 hours.

Figure 8.1a shows that the difference in the mean of Routes 1 and 3 is rather small, but the difference in variance is larger. Route 1 has a lower mean, but Route 3 has a lower variance. Knowing this, it is not surprising that Route 1 has a higher confidence level than Route 3 for some thresholds *lower* than approximately 455 hours, but Route 3 has the highest confidence level and is optimal for thresholds larger than 455 hours.

For all practical purposes, confidence lower than 0.5 are never interesting. For this reason, only the part of the figure showing the CDF representing confidence levels from $\alpha = 0.5$ and higher will be shown for the rest of the analysis presented in this chapter.

8.2 Aggregated Analysis of the Base Set

In this section we present the results and analysis of the Base Set. The 100 instances in the Base Set all have an even geography, randomly drawn draft limits between 50%-90% of total ship capacity, total pickup and delivery each randomly drawn between 80%-90% of total ship capacity, and the optimal route is found using a required confidence level of 95%. As such, we assume a relatively risk averse planner. For more information about the values used for the Base Set, see Chapter 7 and Table 7.2.

As explained in the previous section, for each test instance, we find both the optimal deterministic route and several stochastic routes. We then find which of these routes is the optimal route for the given confidence level. If the optimal route for the given confidence level differs from the optimal deterministic route, the test instance has an optimal *stochastic* route. On the contrary, if the deterministic route is the optimal route for the given confidence level, the test instance has an optimal *deterministic* route.

When analysing the results of a *set* of test instances, two measures are used. The first measure is the percentage of instances with a stochastic optimal route (i.e. Route 1 is not optimal) out of the total number of instances in the set. The instances where a stochastic route is optimal for the given confidence level are referred to as *interesting instances*. In other words, the interesting instances have an optimal solution which differs from the optimal deterministic solution. The measure is shown in Equation (8.1).

. .

$$\frac{\text{Percentage of interesting}}{\text{instances for a set}} = \frac{\frac{\text{Number of instances with}}{\text{Total number of instances in the set}} (8.1)$$

For the interesting instances, we observe that the optimal stochastic solution has a somewhat higher expected completion time than the optimal deterministic solution, but a lower standard deviation. Motivated by this, the second performance measure is the ratio of *decrease* in standard deviation to the *increase* in expected completion time between the optimal stochastic route and the optimal deterministic route. We refer to this ratio as the *SD/Mean ratio*, and the SD/Mean ratio only applies to the interesting instances. For an interesting instance,

SD/Mean ratio =
$$\frac{\text{decrease in standard deviation}}{\text{increase in mean}} = \frac{\sigma_D - \sigma_S}{|\mu_D - \mu_S|},$$
 (8.2)

where μ_D and σ_D are the mean and standard deviation of the completion time of the optimal deterministic route, and μ_S and σ_S are the mean and standard deviation of the completion time of the optimal stochastic route. The SD/Mean ratio tells us by how much we are able to reduce uncertainty by allowing a slightly higher mean. A higher SD/Mean ratio indicates more interesting instances, because taking uncertainty into account is more valuable the higher the SD/Mean ratio is. When analysing and comparing sets of instances we are interested in the *average* SD/Mean ratio of the interesting instances in the set.

Table 8.1 shows the aggregated results for the Base Set. The table includes the number of instances in each set, the average SD/mean ratio and the ratio of interesting instances. In addition, we have included the average mean and average standard deviation of the completion time of both the optimal stochastic solutions and the optimal deterministic solutions for the interesting instances.

From Table 8.1 we see that the percentage of interesting instances is 18%, i.e. 18 of the 100 instances in the Base Set have a stochastic solution as the optimal solution. This implies that for more than 80% of the instances, the optimal deterministic route performs better than the optimal stochastic route. Hence, taking uncertainty into account matters to some degree, but in most cases a deterministic approach would give equally good results. For the interesting instances we thus want to examine how much better the optimal stochastic solution performs than the optimal deterministic solution. We see from Table 8.1 that the SD/Mean ratio is 1.24 on average for the interesting instances. This implies that the standard deviation, and hence uncertainty,

	The Base Set
Number of instances in the set	100
Percentage of interesting instances	18%
The interesting instances of	the Base Set
Average SD/Mean ratio	1.24
Average mean of optimal stochastic routes	387.2 hours
Average standard deviation of optimal stochastic routes	10.0 hours
Average mean of optimal deterministic routes	386.7 hours
Average standard deviation of optimal deterministic routes	10.6 hours

 Table 8.1: Aggregated results for the Base Set.

is on average reduced 1.24 times per unit of increase of the mean from the optimal deterministic to the optimal stochastic solution for these instances. To decide if this is a large enough improvement to justify the use of a stochastic approach when solving the route planning problem we have to assess the possible gains from using the stochastic solution.

As further explained in Section 8.7, we find that by using a stochastic approach when solving the route planning problem for the instances in the Base Set we get less than 0.5% improvement of the threshold we can adhere to for a given confidence level. The improvement of confidence level for a given threshold is less than one percentage point. Solving an optimization problem requires more resources when using stochastic approaches compared to deterministic approaches, and the results form solving the Base Set give reason to question the necessity of using a stochastic approach when considering route planning in a port where waiting times at terminals are uncertain and the characteristics are similar to the instances in the Base Set. One can argue that a route planner will obtain good enough results even though uncertainty is not taken into account. But there could be ports with other characteristics which causes uncertainty to matter more. Some of these characteristics are examined in the following sections. In section 8.3, we take a closer look at the nature of arc and route variance, and in section 8.4 we take a closer look at what happens when certain parameters of the Base Set are changed.

8.3 Detailed Analysis

In this section, we take a closer look at the optimal routes of some test instances from the Base Set, and examine the differences between the optimal stochastic and optimal deterministic route. Reasons why routes have different variances are identified and discussed. The larger the decrease in the standard deviation of the route completion time from the optimal deterministic solution to the optimal stochastic solution relative to the increase in the mean of the route completion from the deterministic solution to the stochastic solution, the higher the SD/Mean ratio. Implying that the higher the SD/Mean ratio, the better the performance of the optimal stochastic solution, and the more important it is to take uncertainty into account for the given test instance. For this reason, we begin by examining the test instance from the Base Set with the highest SD/Mean ratio, Instance 46. To give insights in reasons to why routes have different variances beyond what Instance 46 can provide, one other instance is also examined.

8.3.1 Comparing Routes for Instance 46 from the Base Set

Figure 8.2 shows the geography and characteristics for Instance 46 in terms of draft limits, waiting times distributions, cargo locations and sizes, and direct sailing times.



Figure 8.2: Details of Instance 46 from the Base Set. The first row shows the draft limits for each terminal in tonnes. The second row shows the characteristics of the distribution of the waiting time at each terminal. The letters indicate what terminal from Figure 7.1 is used. The first integer in each cargo icon shows the cargo index, and the number inside is the load. The signs "+" and "-" indicate if it is a pickup or delivery cargo, respectively. The numbers under the black line indicate the approximated *direct* sailing time between the terminals in hours.

Figure 8.3 shows the cumulative distribution function and the plot of the (μ, σ^2) -characteristics of the completion times of the routes found for Instance



46 using the solution method described in Chapter 5.

(a) Cumulative distribution function (CDF) of the completion times of the four possibly optimal routes identified for Instance 46.

(b) Scatter plot of (μ, σ^2) -characteristics of the completion times of the four possibly optimal routes identified for Instance 46.

Figure 8.3: Model output for Instance 46 from the Base Set. Four routes are identified as possibly optimal routes.

Recall that we assume a relatively risk averse planner, and the confidence level, α , required for a route to be optimal is $\alpha = 0.95$ for the Base Set. Figure 8.3a shows that Route 3 is the route with the lowest threshold for $\alpha = 0.95$, and Route 3 is thus regarded as optimal. The black and green lines representing Routes 1 and 3, respectively, intersect at approximately threshold = 378 hours and confidence level $\alpha = 0.67$. This means that for thresholds, or deadlines, lower than 378 hours, Route 1 is optimal, while Route 3 is optimal for thresholds higher than 378 hours. For a risk averse planner, a confidence level of 0.67 is rarely interesting. Hence, for risk averse planners, Route 3 is always optimal.

Looking at Figure 8.3b we see that Route 3 has a slightly higher mean than Route 1, but a lower variance. Recall that the SD/Mean ratio is the decrease in standard deviation over the increase in mean for the optimal stochastic route relative to the optimal deterministic route. For instance 46, the SD/Mean ratio is 2.74, which gives a measure for how much we gain by allowing a slightly higher mean. For the instances of the Base Set with a different optimal route than the optimal deterministic route, the average of the SD/Mean ratio is 1.14. As such, for Instance 46, taking uncertainty into account gives a higher gain than for other instances. The two remaining routes, Routes 2 and 4, are completely dominated by the two other routes for all practical purposes.

Table 8.2 shows the sequence of cargo and terminal visits and the mean, standard deviation, and variance of the complete Routes 1 and 3 for Instance

46. The letters A, B, D, G, I and K indicate what terminal shown in Figure 8.3 is visited, and in what sequence. The routes begin and finish at anchorage, but this is not included in the table as it is true for all routes without exception.

Table 8.2: Cargo and terminal visiting sequence for Routes 1 and 3 of Instance 46 from the Base Set.

Route 1																				
Cargo seq.	4	3	1	20	18	19	11	12	13	15	17	16	14	7	6	8	10	9	2	5
Terminal seq. A				K		$\overline{\mathrm{D}}$	G		I				F		;	D	A	<u> </u>		
(μ,σ^2)	(3	373.2	24, 1	0.93^{2}) = (373.2	4,11	9.5)												
Route 3																				
Cargo seq.	4	3	1	11	20	18	19	12	13	15	17	16	14	7	6	8	10	9	2	5
Terminal seq.		А		D		Κ			G			Ι			В		Γ)	A	1
(μ,σ^2)	(3	373.4	49, 1	0.25^{2}) = (373.4	9,10	5.0)												

From Table 8.2 we can look more closely at what arcs are used in the two routes. To highlight the difference between the routes, Figure 8.4 illustrates the two routes graphically.

The normal distribution of the waiting time for each terminal is represented by the values for (μ, σ^2) above each terminal. The letters indicate the terminals corresponding to terminals from Figure 7.1. The draft limits are omitted here as they are not necessary for the following discussion. The cargo indexes and loads are presented inside each cargo icon, and the direct sailing times are included under each line segment. The terminal sequence of the optimal deterministic route, Route 1, and the optimal stochastic route, Route 3, are shown in the figure. The (μ, σ^2) -characteristics of the normal distribution of the sailing time of each arc is included above the arc. At each vertical line representing a terminal visit, the serviced cargoes are indicated by their index. The figure shows that the routes only differ for some of the arcs. The difference between these arcs is what is interesting when analysing the differences between the routes. For this reason, the rest of the route is dashed and the values shaded, and the parts of the routes that differ are accentuated.

It is interesting to examine the parts of the routes which differ, that is the part of the routes which are not dashed in Figure 8.4. Table 8.3 summarizes the distributions of the sailing times for arcs that differ for Routes 1 and 3.

The sum of the arc means from leaving Terminal A the first time until arriving at Terminal G is 28.09 hours for Route 1 and 28.33 hours for Route 3. Note that this is only the mean for the stochastic part of the route, namely the sailing times, and the loading and cleaning times are omitted here as the sum of these are the same for all routes. The difference in the mean of the sailing times for these *limited* parts of the routes is exactly equal to the difference in the means for the *complete* routes. This difference amounts to approximately



Figure 8.4: Cargo and terminal visiting sequence of the optimal deterministic route, (Route 1), and the optimal stochastic route (Route 3) of Instance 46 from the Base Set.

10,9

7,6,8

(5.1,0²) 2,5

0.25 hours (some rounding of numbers cause the minor deviation). The means of the complete routes are 373.24 for Route 1 and 373.49 for Route 3. We see

	Ro	ute 1	Route 3					
Arc	μ	σ^2	Arc	μ	σ^2			
(A,K)	9.716	$2.330^2 = 5.429$	(A,D)	8.195	$2.896^2 = 8.378$			
(K,D)	7.594	$4.355^2 = 18.966$	(D,K)	7.490	$2.341^2 = 5.480$			
(D,G)	10.776	$4.564^2 = 20.830$	(K,G)	12.647	$4.107^2 = 16.867$			
sum	28.086	$6.725^2 = 45.227$	sum	28.332	$5.544^2 = 30.736$			

Table 8.3: Distribution of the sailing times of the arcs which differ between Routes1 and 3 for Instance 46 from the Base Set.

that the difference between the means and variances of these limited parts of the routes is what gives the overall difference in mean and variance that we also saw from Table 8.2, with approximately $\Delta \mu = 0.25$ and $\Delta \sigma^2 = 14.49$.

Where Route 1 uses the arc from A to K, (A, K), Route 3 instead uses the arc from A to D, (A, D). Arc (A, D) used by Route 3 has a higher standard deviation than arc (A, K) used by Route 1, which is due to the different distributions of the waiting time at the destination Terminals D and K. In spite of this, the complete Route 3 has an over all lower variance than Route 1. This is due to the lower variance of arcs (K, G) and from (D, K) used by Route 3, instead of arcs (K, D) and (D, G), used by Route 1.

Recall that for arcs directed towards anchorage, which is the case for arcs (K, D) and (K, G), the variance is affected by both the Pile-Up Effect and the Anchorage Effect, explained in Chapter 6. For (K, D), the ship sails directly to the destination terminal without passing it more often than for (K,G). This causes a larger pile of sailing times equal to the direct sailing distances for (K, D) than for (K, G). In fact, the probability for the ship to sail directly to the destination terminal for (K, D) ($\Pr(w_d \leq D_{KD})$) where $D_{KD} = 1.020 + 0.73 + 1.02$) is approximately 44%, while the corresponding probability for (K, G) is less than 3%. From Section 6.2.4 we know that, when the Anchorage Effect is marginal, the Pile-Up Effect gives a decrease in the variance relative to when there is no Pile-Up Effect, and the Pile-Up Effect is stronger the higher the probability of sailing directly to the next terminal. This argues that (K, D) should have less variance than (K, G), which is not the case. However, the Anchorage Effect is also at play here, and interferes with the Pile-Up Effect. From Section 6.2.4, we also know that the higher the probability of waiting at anchorage, i.e. stronger the Anchorage Effect, the less the resulting variance. The Anchorage Effect is stronger for (K, G) than (K, D) because of the higher expected waiting time at Terminal G than at Terminal D. When both effects are at play, their contribution to the variance is not necessarily by reduction, as the mean of the distributions is typically between the pile causing the Pile-Up Effect and the concentrated values giving the Anchorage Effect, such that the pile and the concentrated values might contribute to *more* variance instead of reducing the variance. This must be the explanation to why there can be less variance associated with (K, G) than with (K, D), even if the Pile-Up effect is stronger for (K, D).

Arc (D, G), used by Route 1, has a higher variance than arc (D, K), used by Route 3. It is more likely for the ship to end up waiting at anchorage for (D, G) than for (D, K) due to the different distributions of the waiting times at the two destination terminals. There is approximately a 17% probability for the ship to wait at anchorage for (D, G), while this probability is almost zero for (D, K). Hence, the Anchorage Effect is stronger for (D, G), which argues that (D, G) should have less variance associated with it according to the results of the simulations in Section 6.2.4. However, this is not the case, and the difference in the distributions of the waiting times at the destination terminals must be the reason why (D, K) has less variance associated with it.

Comparing the symmetric arcs (D, K) and (K, D) reveals how the conditions and relative sizes affect the resulting variance to a large degree, and the unique conditions are important for the outcome. With the small distances between terminals and the relatively large waiting times possible (mean of the distributions vary randomly between 0 and 10 hours) for all the instances in the Base Set, one might expect the Anchorage Effect to often overrule the Pile-Up Effect. The distance to anchorage is small *enough* for this to *often* be the case. When the Anchorage Effect generally overrules the Pile-Up Effect, arcs directed away from anchorage are typically more right-skewed than the symmetric arcs directed towards anchorage, and the results for the simulation in Section 6.2.4 suggest that increased Anchorage Effect reduces variance. The Anchorage Effect is stronger for (D, K), directed away from anchorage, than for (K, D), directed towards anchorage. This corresponds well with the lower variance associated with (D, K) than with (K, D). However, as described some paragraphs above, the Pile-Up Effect is rather strong for (K, D) in this specific example. The probability for the ship to sail directly to D when using (K, D) is much higher than the probability for the ship to wait at anchorage when using (D, K), which argues that (K, D) should have a lower variance. However, with the specific conditions here, the difference in variance must be due to the different distribution of the waiting times at the destination terminals.

Based on the discussion above, it is apparent that when there are several factors that affect the variance of arcs, it is not always straightforward what is the main reason why some routes have a lower variance than others. For this test instance, we see that the main reason why arcs have different variances associated with them is different distributions of the waiting times at terminals. However, the Pile-Up Effect and the Anchorage Effect can also contribute to making it possible to find a different route with a smaller variance than the optimal deterministic route.
8.3.2 Comparing Routes for Instance 1 from the Base Set

Instance 1 from the Base Set has a stochastic optimal route, and the reason why the optimal route has a lower variance than the optimal deterministic route can be explained differently than for Instance 46. Figure 8.5 shows the details of Instance 1 from the Base Set.



Figure 8.5: Details of Instance 1 from the Base Set. The first row shows the draft limits for each terminal in tonnes. The second row shows the characteristics of the distribution of the waiting time at each terminal. The letters indicate what terminal from Figure 7.1 is used. The first integer in each cargo icon shows the cargo index, and the number inside is the load. The signs "+" and "-" indicate if it is a pickup or delivery cargo, respectively. The numbers under the black line indicate the approximated *direct* sailing time between the terminals in hours.

Figure 8.6 shows the CDF and (μ, σ^2) -characteristics for the routes found for Instance 1. We see from the figure that Route 4 has a slightly higher mean, but a lower variance than Route 1. For thresholds lower than approximately 423 hours, Route 1 is optimal, while Route 4 is optimal for thresholds higher than 423 hours. The confidence level for which the graphs for Routes 1 and 4 intersect is approximately $\alpha = 0.88$. For confidence levels of 0.9977 and higher, corresponding to thresholds of 439.9 and higher, Route 3 becomes optimal. But these values are very high, and for all practical purposes, Route 4 is optimal for all thresholds equal to 423 hours and higher. This corresponds to required confidence levels of $\alpha = 0.88$ and higher. Route 2 and 5 are completely dominated by the other routes. Table 8.4 summarizes the cargo and terminal sequence of Routes 1 and 4.

As is done in the previous section for Instance 46, we can use Table 8.4 to



(a) Cumulative distribution function (CDF) of the completion times of the five possibly optimal routes identified for Instance 1.

(b) The (μ, σ^2) -characteristics of the completion times of the five possibly optimal routes identified for Instance 1.

Figure 8.6: Model output for Instance 1 from the Base Set. Five routes are identified as possibly optimal routes.

Table 8.4: Cargo and terminal visiting sequence for Routes 1 and 4 of Instance 1from the Base Set.

Route 1																				
Cargo seq.	11	10	9	13	12	3	1	2	6	4	14	15	19	20	18	16	17	5	8	7
Terminal seq.		D		(G		А		I	В		I			Κ			В	D	_
(μ,σ^2)	(40	9.63,	11.2	$20^2) =$	= (40	9.63,	125.4	14)												
Route 4																				
Cargo seq.	11	10	9	13	12	14	15	6	4	3	1	2	19	20	18	16	17	5	8	7
Terminal seq.		D		(G		I]	В		А				Κ			В	D	_
(μ,σ^2)	(41	0.63,	10.3	$34^2) =$	= (41	0.63,	106.9	(91)												

extract information about the different arcs that are used in Routes 1 and 4. We see that the routes differ by Route 1 using arcs (G, A), (A, B), (B, I) and (I, K) where Route 4 uses arcs (G, I), (I, B), (B, A) and (A, K). Figure 8.7 illustrates the routes graphically.

Figure 8.7 shows the terminal sequence of the optimal deterministic route, Route 1, and the optimal stochastic route, Route 4, of Instance 1 from the Base Set. The (μ, σ^2) -characteristics of the sailing time for each arc excluding loading and cleaning is included above each arcs, and the vertical lines representing terminal visits, the serviced cargoes are indicated by their index. The dashed and shaded parts of the figure represents the parts of the routes that are equal for Routes 1 and 4, and we are interested in the part of the route that is not dashed. Table 8.5 summarizes the characteristics of the arcs that differ between the routes.



Figure 8.7: Cargo and terminal visiting sequence of the optimal deterministic route (Route 1) and the optimal stochastic route (Route 4) of Instance 1 from the Base Set.

From Table 8.5 we see that the difference between the means corresponds to the difference we saw for the complete routes in Table 8.4, with approximately $\Delta \mu = 1$. The difference in variance between the two routes of approximately

	Ro	ute 1	Route 4			
Arc	μ	σ^2	Arc	μ	σ^2	
(G,A)	8.761	$3.640^2 = 13.253$	(G,I)	2.188	$0.720^2 = 0.519$	
(A,B)	10.042	$3.826^2 = 14.639$	(I,B)	10.188	$4.004^2 = 16.031$	
(B,I)	4.238	$0.720^2 = 0.519$	(B,A)	9.505	$2.924^2 = 8.551$	
(I,K)	14.002	35.528 = 35.528	(A,K)	16.161	$4.493^2 = 20.189$	
sum	37.043	$8.001^2 = 63.939$	sum	38.042	$6.730^2 = 45.290$	

Table 8.5: Distribution of the sailing times of the arcs which differ between Routes1 and 4 for Instance 1 from the Base Set.

 $\Delta \sigma^2 = 18.6$ also corresponds to the difference for the complete routes from Table 8.4 (minor deviations are due to rounding).

The most important difference between the routes, is that Route 4 avoids the large variance that applies when entering Terminal K from Terminal I. To do so, Route 4 enters Terminal K from Terminal A instead. To understand what gives arc (A, K) a lower variance than (I, K), even if the destination terminal is the same, we did a simulation of the realised sailing times. The resulting distributions of the sailing times for (I, K) and (A, K) are shown in Figure 8.8.



(a) Distribution of the realised sailing times of (b) Distribution of the realised sailing times of arc (I, K). arc (A, K)

Figure 8.8: Distributions of arcs (I, K) and (A, K) used for Routes 1 and 4, respectively, for Instance 1 from the Base Set.

Figure 8.8 shows that the Anchorage Effect is at play for both arcs (I, K) and (A, K). The high concentration of sailing times derived from waiting at anchorage is wider for (A, K) than for (I, K). This is not surprising, as Terminal A is located much closer to anchorage than Terminal I is, and a larger amount of the realised waiting times cause the ship to wait at anchorage. From

the simulations in Section 6.2.4 we know that when only the Anchorage Effect is at play, a stronger Anchorage Effect gives a lower variance, which is the case for these two arcs. We see that for arc (A, K), the mean of the arc, of 16.2 hours, falls within the range of sailing times where the ship has waited at anchorage for some amount of time. This, way, the concentration around the mean is high, causing a lower variance. For (I, K), the mean of 14.0 hours does not fall within the concentrated realised sailing times around 20 hours, and the sailing times where the ship has waited at anchorage contribute to a larger variance instead. This gives (I, K) a higher variance than (A, K).

We see from Figure 8.7 that Route 4 has one more arc directed towards anchorage than Route 1. When ignoring the difference in the waiting times distributions as a main reason why arcs have different variances, different conditions cause the Anchorage Effect and the Pile-Up Effect to work in various ways, and they often interfere. For shorter anchorage distances, or higher means of the waiting times, when the remaining distances are unchanged, ships will more often end up at anchorage. When the Anchorage Effect is strong relative to the Pile-Up Effect, arcs directed away from anchorage can typically have less variance associated with them. When the Pile-Up Effect is strong, arcs directed towards anchorage can typically have less variance associated with them. However, this ignores the possible differences of waiting times distributions and how the unique conditions largely affect the outcome when the two effects interfere. For the conditions used in the instances in the Base Set, with distances between terminals of approximately 1 hour, mean of the waiting time varying from 0 to 10 hours and a sailing time to anchorage of approximately 5 hours from the closest terminal, we see that for the 18 interesting instances from the Base Set, the optimal stochastic route always has equally many or more arcs directed towards anchorage compared to the optimal deterministic route. This must mean that the conditions in the Base Set in general cause the Pile-Up Effect to be stronger than the Anchorage Effect.

8.4 Sensitivity Analysis

In this section we analyse the impact of changing certain parameters in the input data explained in Chapter 7. For each variation we want to examine, a new set of 100 instances is generated from the Base Set, holding everything constant except from the parameter we want to analyse. The aggregated information of the new set is compared to the Base Set. We analyse variations in geography, draft limits, and the risk profile of the decision maker.

8.4.1 Geography

As explained earlier, the geography of the true routing problem changes based on which terminals the customer requests are located in. To investigate the impact of different geographies, two new sets of instances are generated. The instances in the new sets are equal to the Base Set except for the geography of the terminals. The instances in the Base Set have, as mentioned, Even geography, while the instances in the two new sets have Far Away and Split geographies, respectively. The geographies are explained in more detail in Chapter 7. When changing the geography, the direct sailing times change, and the expected travel times and variances of the arcs have to be recalculated. Every thing else, including the distribution of the waiting times at the terminals, is held constant.

Table 8.6 shows the aggregated results of the sets of test instances with Even, Far Away, and Split geographies. Recall that the *interesting instances* are the test instances where the optimal stochastic route performs better than the optimal deterministic route for the chosen confidence level of 95%. The table includes the number of instances in each set, the average SD/Mean ratio and the percentage of interesting instances. In addition, we have included the average mean and average standard deviation of the completion time of both the optimal stochastic solutions (opt. stoch. solution) and the optimal deterministic solutions (opt. det. solution) for the interesting instances.

	The Base Set	Variation Set 1	Variation Set 2
Geography	Even	Far Away	Split
Instances in the set	100	100	100
Percentage of interesting instances	18%	20%	22%
	Ι	nteresting instance	es
Average SD/Mean ratio	1.24	1.51	1.99
Average μ of opt. stoch. solution, $\bar{\mu}_s^i$	387.2 hours	396.7 hours	402.6 hours
Average σ of opt. stoch. solution, $\bar{\sigma}_s^i$	10.0 hours	10.7 hours	10.7 hours
Average μ of opt. det. solution, $\bar{\mu}_d^i$	386.7 hours	395.8 hours	401.9 hours
Average σ of opt. det. solution, $\bar{\sigma}_d^i$	10.6 hours	11.6 hours	11.4 hours

Table 8.6: Aggregated results from the set of test instances with Even, Far Away,and Split geographies.

From the table we see that both Far Away and Split have a somewhat higher percentage of interesting instances the Base Set, meaning that taking uncer-

tainty into account matters more often for these geographies. The physical split, or gap, of the Split geography gives longer arcs, and for arcs directed towards anchorage where the gap is crossed, the Pile-Up Effect can be stronger. When ignoring the Anchorage Effect, the Pile-Up Effect contributes to less variance, but when both effects are at play, they typically contribute to more variance, as the mean of the distribution of the sailing times typically falls somewhere between the pile associated with the Pile-Up Effect and the concentrated sailing times associated with the Anchorage Effect. This way, stronger Pile-Up Effect might cause more variance, which can be the reason why variance in general increases for the Split geography. The higher SD/Mean ratios and the higher percentages of interesting instances for Split and Far Awat geographies, taking uncertainty into account is somewhat more important than for the Even geography.

The value of the stochastic solution of the interesting instances of Variation Sets 1 and 2 should also be examined. In Section 8.2 it is mentioned that the value of the average stochastic solution for the interesting instances of the Base Set is very small. This is examined in more detail in Section 8.7. From Table 8.6 we see that $\bar{\mu}_s^{VS1} = 396.7$ hours and $\bar{\mu}_d^{VS1} = 395.8$ hours are the average means of the completion time of the optimal stochastic solutions and optimal deterministic solutions of the interesting instances of Variation Set 1. The difference in the expected value of these two routes is 0.9 hours, which is very little, considering that the routes last for around 16 days. This is also true for Variation Set 2, as $\bar{\mu}_s^{VS2} = 402.6$ hours to $\bar{\mu}_d^{VS2} = 401.9$ hours, a difference of 0.7 hours during almost 17 days. The small increase in expected completion time from the optimal deterministic route to the optimal stochastic route is something one could accept as long as the gain in decreased uncertainty makes up for it. But when looking at the decrease in standard deviation from the average optimal stochastic to average optimal deterministic routes, we see that this decrease is rather small for both Variation Set 1 and Variation Set 2. These small differences in the distributions of the completion times of the average optimal stochastic and deterministic routes, respectively, suggest that the value of the stochastic solution is rather small. As for the Base Set, uncertainty matters to some degree for Variation Sets 1 and 2, as there are instances where the stochastic optimal route performs better than the deterministic optimal route. But choosing the optimal stochastic route over the optimal deterministic route does in reality not result in significant improvements.

The reason for the relatively small differences between the values of Base Set, Variation Set 1, and Variation Set 2 could be that the direct sailing times between the terminals in Houston Ship Channel are short compared to the sizes of the waiting times and the loading and cleaning times. Load and clean times contribute to typically 60% to 80% of the total completion time of the route. Thus, changing the geography does not give rise to large changes of the problem network.

Distance Between Anchorage and the Closest Terminal

In the case port, Houston Ship Channel, the anchorage is located at sea 22 nautical miles away from the closest terminal. In other ports, the anchorage could be placed closer to shore. In Houston Ship Channel, there are in reality layberths closer to the terminals that can be used. Hence, we are interested in examining the impact that the distance between anchorage and the terminal closest to anchorage has on our results. For the instances in the Base Set, which uses realistic numbers from Houston Ship Channel, the direct sailing time from anchorage to the terminal closest to anchorage is 5.12 hours. For the 100 new instances generated the direct sailing time to the terminal closest to anchorage is set to 1.00 hour, while everything else is held constant. 1.00 hour is chosen because it equals the average direct sailing distance between the terminals in the Base Set. This means that for the new instances, the terminals *and* and the anchorage are close to evenly distributed along a horizontal line.

Table 8.7 shows the aggregated results of the Base Set and the new set of test instances with a shorter distance between anchorage and the closest terminal. This new set of test instances is referred to as *Variation Set 3*. Recall that the *interesting instances* are the instances where the optimal stochastic route performs better than the optimal deterministic route for the chosen confidence level of 95%. The table shows the sailing time between anchorage and the terminal closest to anchorage (*Sailing time to anchorage*) used for the two sets. includes the number of instances in each set, the average SD/Mean ratio and the percentage of interesting instances. In addition, we have included the average mean and average standard deviation of the completion time of both the optimal stochastic solutions and the optimal deterministic solutions for the interesting instances.

From Table 8.7 we see that the percentage of interesting instances for Variation Set 3 is much higher than for the Base Set. For Variation Set 3, more than half of the instances have a stochastic optimal solution, i.e. another route than the optimal deterministic route is the best choice. This suggests that uncertainty is more important when the distance to anchorage is closer to the closest terminal.

When anchorage is closer while the expected waiting times at terminals remain unchanged, the ship ends up waiting at anchorage more often, and the distribution of the sailing times of the arcs are on average more right skewed than for the corresponding arcs in instances from the Base Set. For a skewed distribution of the sailing time of an arc where the conditions are such that a large portion of the sailing times are concentrated close to the mean, the variance is typically smaller than when the mean of the distribution does not coincide with the concentrated sailing time values. For a short anchorage distance, the conditions typically cause the mean to coincide with the concentrated sailing time values. This happens more often for arcs in Variation set 3

	The Base Set	Variation Set 3
Sailing time to anchorage	5.12 hours	1.00 hours
Instances in the set	100	100
Percentage of interesting instances	18%	56%
	Interesting	g instances
Average SD/Mean ratio	1.24	2.29
Average μ of stochastic optimal routes	387.2 hours	378.8 hours
Average σ of stochastic optimal routes	10.0 hours	8.5 hours
Average μ of deterministic solution	386.7 hours	378.5 hours
Average σ of deterministic solution	10.6 hours	8.8 hours

Table 8.7: Aggregated results from the set of test instances with a shorter distancebetween anchorage and the terminal closest to anchorage.

than in the Base Set, because anchorage distance in Variation Set 3 is shorter. In other words, the Anchorage Effect is stronger and gives lower variance on average for the optimal routes found for instances from Variation Set 3, as seen from table 8.7.

The average SD/Mean ratio is higher for Variation Set 3 than for the Base Set. A reason for this could be that when there exists routes with lower variances, the possibility to improve from the deterministic solution increases. Ignoring differences in waiting times distributions, the increased Anchorage Effect is typically stronger for arcs directed away from anchorage, as the departure terminal is closer to anchorage. This way, the increased Anchorage Effect might give bigger differences between the variances associated with different arcs, even if the overall variance is lower. The higher percentage of interesting instances shows that taking uncertainty into account matters *more often*, while the higher SD/Mean ratio indicates that taking uncertainty into account is *more valuable* per interesting instance. However, the increase in SD/Mean ratio is marginal compared to the increase in the percentage of interesting instances.

If we take a closer look at how much better the stochastic solutions are on average compared to the average optimal deterministic solutions, we see that the differences in the distributions of the route completion times are very small. The expected completion time increases with 0.3 hours on average when choosing the stochastic optimal route instead of the average deterministic optimal routes for Variation Set 3. The standard deviation decreases with 0.3 hours. This implies that even though the stochastic route is optimal for almost half of the instances, i.e. more frequently than for the Base Set, the benefit from choosing this route instead of the optimal deterministic route is usually very small.

8.4.2 Draft Limits

The draft limits at the terminals can affect the optimal solution of a test instance as well as the run time of the model. When draft limits become tighter, the problem is restricted, and the solution space decreases. As such, the optimal route becomes the same or worse when draft limits become tighter. For many instances, the ship will have to visit the terminals more often than for less tight draft limits to be able to service all the cargo without violating the draft limits. The problem becomes harder to solve, and the solution time increases.

Recall from Chapter 7 that the draft limits of the test instance are expressed as a percentage of ship capacity and are randomly drawn for each terminal from a certain range, except for one or two terminals which get a draft limit of 120% of total ship capacity to ensure that the ship is allowed to leave and enter the port. For the Base Set, the draft limits are randomly drawn from the range of 50%-90% of total ship capacity. For the new set, the draft limits are randomly drawn from the range of 40%-80% of total ship capacity. Everything else remains unchanged. The new set generated with tighter draft limits is referred to as *Variation Set 4*.

Table 8.8 shows the aggregated results of the Base Set and of Variation Set 4. Recall that the *interesting instances* are the test instances where the optimal stochastic route performs better than the optimal deterministic route for the chosen confidence level of 95%. The table includes the number of instances in each set, the average SD/Mean ratio and the percentage of interesting instances. In addition, we have included the average mean and average standard deviation of the completion time of both the optimal stochastic solutions and the optimal deterministic solutions for the interesting instances.

When tightening the draft limits, almost half of the instances have no feasible solutions. In addition, the run time of the model increases drastically. From taking 20 minutes for solving the 100 instances in the Base Set, the model takes more than 2 hours for the 100 instances in Variation Set 4, where everything is equal to the Base Set except for the draft limits. It is interesting to observe that a small change in the draft limit range give rise to quite large changes in the results.

The rate of interesting instances is only 9% for Variation Set 4 compared to 18% of the Base Set, suggesting that the importance of including uncertainty is less when the draft limits are tighter. Again, we see that the difference

	The Base Set	Variation Set 4
Draft limit range:	50%-90%	40%-80%
Instances in the set	100	100
Interesting instances	18%	9%
Model run time	$20 \min$	2.3 hours
Instances with no solution	0%	44%
	Interesting	g instances
Average SD/Mean ratio	1.24	1.11
Average μ of optimal routes	387.2 hours	418.4 hours
Average σ of optimal routes	10.0 hours	11.6 hours
Average μ of deterministic solution	386.7 hours	417.6 hours
Average σ of deterministic solution	10.6 hours	12.4 hours

 Table 8.8:
 Aggregated results from the set of test instances with more binding draft limits.

between the average characteristics of the stochastic optimal routes and the deterministic optimal routes of the interesting instances of Variation Set 4 is rather small. This implies that there is not too much to gain by solving the route planning problem with a stochastic approach instead of a deterministic approach even though the draft limits are tighter.

In addition, we find that, on average, the expected completion time of the routes are higher. This is not surprising as the ship has to visit terminals more often for each terminal to be able to service all the cargoes located at the terminal.

8.4.3 Risk Profile of the Decision Maker

The risk profile of the decision maker, and hence the confidence level used to determine what route is the optimal route, can affect the results. Less risk averse decision makers will allow lower confidence levels when deciding what route is the optimal one. A less risk averse decision maker typically prefers routes with lower expected completion times at the expense of higher variances. A set of instances where the optimal route is identified as the route with the best threshold for the confidence level $\alpha = 0.6$ is generated, and the set is referred to as *Variation Set 5*.

Table 8.8 shows the aggregated results of the Base Set and the new set of test instances having a lower confidence level for choice of optimal route. Recall that the *interesting instances* are the test instances where the optimal stochastic route performs better than the optimal deterministic route for the chosen confidence level. The table includes the number of instances in each set, the average SD/Mean ratio and the percentage of interesting instances. In addition, we have included the average mean and average standard deviation of the completion time of both the optimal stochastic solutions and the optimal deterministic solutions for the interesting instances.

	The Base Set	Variation Set 5
Confidence level	95%	60%
Instances in the set	100	100
Interesting instances	18%	0%
	Interesting	g instances
Average SD/Mean ratio	1.24	N/A
Average μ of optimal routes	387.2 hours	N/A
Average σ of optimal routes	10.0 hours	N/A
Average μ of deterministic solution	386.7 hours	N/A
Average σ of deterministic solution	10.6 hours	N/A

 Table 8.9:
 Aggregated results from the set of test instances where a lower confidence level is require for routes to be optimal.

As reported in Table 8.9, none of the instances have a stochastic solution as the optimal, i.e. the deterministic solution is always the preferred one. This implies that less risk averse planners should not bother to use a stochastic approach when solving the route planning problem.

8.5 Summary of Reasons to Different Variances

From Sections 8.3 and 8.4, it is apparent that there are many conditions and relative sizes that must be taken into account when trying to understand the behaviour of the uncertain sailing times, and the interaction of the conditions is rather complex. From Section 8.3 we can learn that for the interesting instances in the Base Set, the different variances associated with arcs is often largely due to the differences between the distributions of waiting times at terminals. But the difference between arc variances can sometimes be explained by the Anchorage Effect and the Pile-Up Effect. The possibility of waiting at anchorage gives rise to a higher concentration of realised sailing times stemming from waiting times where the ship waits at anchorage. This gives the Anchorage Effect, which contributes to right-skewed distributions. When the Pile-Up effect is ignored or is marginal, stronger Anchorage Effect gives less variance, as seen from the simulations in Section 6.2.4. The strength of the Anchorage Effect depends the distances between terminals and anchorage, the distributions of the waiting times at the destination terminals and the means of the sailing time distributions relative to the value of more frequent and concentrated realised sailing times. For arcs directed towards anchorage, the possibility for the ship to sail *directly* to the destination terminal gives high frequency of realised sailing times to equal the direct sailing time, i.e. a high *pile* for values equal to the direct sailing time for the distribution of sailing times. This is referred to as the Pile-Up Effect, which gives left-skewed distributions. When the Anchorage Effect is ignored or is marginal, a stronger Pile-Up effect contributes to decreased variance of the distribution of the sailing times. However, when both effects are at play, they interfere. When they interfere, the mean of the distribution of the sailing time typically falls between the pile associated with the Pile-Up Effect and the concentrated realised sailing times associated with the Anchorage Effect. This might cause stronger Pile-Up Effect and the Anchorage Effect to increase the variance of the distribution of arc sailing times when they interfere.

Regarding the Pile-Up Effect, the contribution to decrease in variance can be two-sided. Ignoring interference with the Anchorage Effect, the strength of the Pile-Up Effect depends on the mean of the waiting time at the destination terminal relative to the direct sailing time between the terminals. The larger the sailing time between the terminals relative to the mean of the waiting time, the stronger the Pile-Up Effect, and the lower the arc variances, as shown in Section 6.2.4. But, as explained in Chapter 7, the distribution of the waiting times are modelled such that the standard deviation, and hence the variance, increases as the mean increases. A lower expected waiting time gives a stronger Pile-Up Effect, but a lower expected waiting time also gives a lower variance regardless of the Pile Up Effect. When the direct sailing time is increased rather than decreasing the mean of waiting time distribution, as is done in Section 6.2.4, it is the Pile-Up Effect alone that gives the decrease in variance. This means that the decrease in variance that we see when the Pile-Up Effect is the only effect at play, is in general not necessarily only due to the Pile-Up Effect, but can also be due to changes in the distribution of the waiting time.

In general, we see that the difference in the behaviour of the stochastic sailing times for different arcs is sometimes due to the possibility of sailing directly to the next terminal and/or the possibility of waiting at anchorage and the resulting Pile-Up Effect and Anchorage Effect, but more often the difference is due to the differences in the distributions of the waiting times at terminals. By choosing arcs with different variances, the variance associated with complete routes differ, naturally.

8.6 Evaluation of the Approximation

In this section, we evaluate the assumption that arcs and hence complete routes follow unskewed normal distributions. As explained in Chapter 6 and in the above sections, there are in reality several effects at play which cause arcs to follow skewed and truncated distributions. To approximate the distribution of the arc sailing times, we use the mean and variance of the skewed and truncated distributions of the sailing times we get from simulation to construct regular, unskewed normal distributions. These approximated, unskewed and non-truncated normal distributions are used when solving our model. To evaluate this approximation, we focus on Instance 46. Recall from Section 8.2 that Instance 46 is the instance with the highest SD/Mean ratio, meaning it is the instance from the Base Set where taking uncertainty into account matters the most.

To be able to evaluate the assumption of normally distributed arcs and routes and the applied approximation, a simulation is done. For this simulation we obtain a distribution of realised route completion times resulting from randomly drawn waiting times. 10,000 sets of realised waiting times are drawn at random from the distributions of waiting times at each terminal, and the corresponding realised arc sailing times are calculated. In other words, we get 10,000 sets of realised sailing times along each arc. The sailing times for the arcs used by the optimal deterministic and optimal stochastic routes, and the sum of the fixed loading and cleaning times, can then be summed to get 10,000 realised route completion times for these two routes. Note that unlike simulations done in earlier chapters and sections, we now simulate the entire route completion times, not just arcs. Recall from Section 8.3.1 that Route 1 is the optimal deterministic route for Instance 46, and Route 3 is the optimal stochastic route.

For the optimal stochastic route, Route 3, of Instance 46, Figure 8.9 shows the distribution of route completion times we get from the simulation. The probability density function (PDF) of the normal distribution of completion times based on the output from our model presented in Chapters 4 and 5 (i.e. the (μ, σ^2) -characteristics we get using Nikolova's method) is illustrated by the orange line (scaled up as the y-axis of a PDF graph never exceeds 1).

From Figure 8.9 we see that the realised route completion time sample set is not far from normally distributed. However, one must keep in mind that this is based on the assumption of underlying normally distributed waiting times. The assumption of normally distributed waiting times is difficult to evaluate given the limited access to real data for waiting times, and this assumption is



Figure 8.9: Sample distribution and distribution outputted by our model for route completion times for the optimal stochastic route, Route 3, for Instance 46. The y-axis is occurrences and the x-axis is values in hours for route completion times.

not evaluated in this thesis. Table 8.10 summarizes some interesting characteristics of the sample set and the characteristics outputted by our model for Route 3 of Instance 46. The table shows that the distribution of the samples is only slightly skewed. This indicates that the approximation we do when ignoring skewness and assuming normally distributed route completion times is not that far off.

Table 8.10: Characteristics of the sample set of realised route completion times and characteristics outputted by the model for the optimal stochastic solution, Route 3, for Instance 46.

	The Sample Set	Our Model
Mean, μ	373.58 hours	373.49 hours
Variance, σ^2	11.41^2	10.25^{2}
Skewness	-0.21	N/A

Table 8.10 also shows that the means of the distributions are approximately the same. This is not surprising as the arc means are constructed by sampling 10,000 realised sailing times based on waiting times drawn from the same distributions as used for this simulation. However, the variance is higher for the sampled set of realised route completion times than what is suggested by our model. Kenyon and Morton (2003) found that the more skewed the underlying arc travel time distribution is, the higher the variance of the sum (i.e. the variance of the complete route) became. This conforms with the fact that our model, assuming non-skewed distributions, suggest a lower variance than the true variance. Note that it is the *sum* of skewed distributions that has more variance than the sum of corresponding but unskewed distributions. The underlying arc travel times that are summed in our model are not completely normally distributed *in reality*, as we saw from Sections 6.2.4 and 6.2.3, while the approximated distributions used by our optimization models are unskewed normal distributions. This explains why our model suggests a lower variance than the distribution found by simulation suggests.

Figure 8.10 shows the set of route completion times for the optimal deterministic route, Route 1. The sampled waiting times, and hence resulting sailing times, are the same as for Figure 8.9. Table 8.11 summarizes some interesting characteristics of the sample set of route completion times for Route 1 and the distribution characteristics outputted by our model for Route 1 of Instance 46 (scaled up as the y-axis for a PDF graph never exceeds 1).



Figure 8.10: Sample distribution and distribution outputted by our model for route completion times for the optimal deterministic route, Route 1, for Instance 46. The y-axis is occurrences and the x-axis is route completion time values.

 Table 8.11:
 Characteristics of the sample set of realised route completion times and characteristics outputted by the model for the optimal deterministic solution, Route 1, for Instance 46.

	The Sample Set	Our Model
Mean, μ	373.37 hours	373.24 hours
Variance, σ^2	12.26^{2}	10.93^{2}
Skewness	-0.14	N/A

Table 8.11 also shows a mean of the sample set approximately equal to the mean suggested by the model, but a higher variance. Not surprisingly, the variance for the samples for Route 1 is higher than the variance of the samples

for Route 3, and the mean is somewhat smaller. The SD/Mean ratio we get based on the samples for the two routes is 4.05, compared to 2.74 based on the model outputs. This *might* indicate that in reality, taking uncertainty into account is more important than what the results from our model indicate. The fact that our model suggests a lower variance than the real variance of the distributions supports this argument. When there is more uncertainty related to the problem than what is suggested by the model, the gain of accounting for uncertainty might be higher than suggested by our results.

8.7 Practical Implications

In this section, the practical implications of the results found in the above sections are discussed. To be able to assess the gain from accounting for uncertainty and choosing the stochastic solution over the deterministic solution, improvements in the threshold one must adhere to and in the achieved confidence levels are analysed. These improvements say something about the value of the stochastic solution.

As mentioned, Instance 46 is the instance of the Base Set with the highest SD/Mean ratio, implying that it is the instance where one gains the most by choosing the stochastic route over the deterministic route, as the increase in mean is small compared to the decrease in standard deviation. The analysis in Section 8.3.1 examines in details why the optimal stochastic route has a lower variance than the optimal deterministic route for Instance 46. In this section we look into how much the shipping company really gain by choosing the optimal stochastic route over the optimal deterministic route.

Figure 8.3a shows that for a confidence level of 95% the optimal stochastic route of Instance 46 is better than the optimal deterministic route, because it can provide a lower threshold for the completion time. The optimal stochastic route has a 95% probability of completing within 390.3 hours. The optimal deterministic route, on the other hand, can with 95% probability complete within 391.2 hours. This gives an improvement in the threshold of less than one hour. As both routes take approximately 16 days to finish, a difference in one hour is not much. It equals an improvement of 0.22% of the threshold.

For a confidence level of 99.9% the optimal stochastic route will complete within a threshold of 405.2 hours while the optimal deterministic route will complete within 407.0 hours. For this high confidence level the difference is about two hours, implying an improvement of confidence level of 0.46%. As the difference in completion time of the two routes increases with increasing confidence levels, we can assume that the improvement of the threshold from choosing the stochastic route over the deterministic route will be less than 0.5% for Instance 46. As Instance 46 is the instance where including uncertainty matters the most, we can conclude that for all the instances in the Base Set, we gain less than 0.5% in improved threshold by choosing the stochastic optimal route over the deterministic optimal route.

Recall the one-to-on correspondence between the threshold and the confidence level for a given route with given mean and variance. In the above paragraphs, the confidence level is fixed at some level (95%), and the threshold one must adhere to for this given confidence level depend on the route characteristics (i.e. mean and variance). If, on the other hand, the threshold is the fixed, required parameter which the company must adhere to, we can calculate the corresponding improvements of confidence level one gets by choosing the stochastic solution over the deterministic solution. If the threshold the company must adhere to is 390 hours, the stochastic route will complete within this threshold with 94.6% certainty. The deterministic route will complete within this threshold of 390 hours with 93.7% certainty. This implies an increase of confidence level by 0.9 percentage points. In real life an improvement of the confidence level of 0.9 percentage points does not give significant gains. For all practical purposes this can be considered as a negligible improvement. One might argue that as the fixed loading and cleaning times constitute 60%-80% of the sailing times the change as percentage of the remaining 20%-40%is larger, and it is changes to these 20%-40% which are of interest. This means that the value of using the stochastic solution over the deterministic can be considered as larger. However, the cleaning and loading times are necessary and an unavoidable part of the complete route, so it makes more sense to compare the change to the complete routes and not just the stochastic part of the routes.

In Chapter 1, a profitability indicator, the $\mathrm{T/C}$ result, is presented. The measure is expressed as

T/C result = $\frac{\text{freight income - voyage costs}}{\text{voyage duration}}$.

From the above paragraphs when the optimal stochastic solution is chosen over the deterministic solution, the threshold one must adhere to for a given confidence level is slightly reduced, but the expected route completion time this entails is slightly *increased*. The increase of route completion time is done at gain of higher confidence levels for given thresholds. This means that the denominator in the T/C result increases, which actually results in worse T/C results. However, the gain in increased predictability might be advantageous in business arrangements. If on-time arrival contributes to decrease in lost income or lower costs, or if schedules can be made tighter and different port-visits can be made more back-to-back, the nominator might improve. This depends on the extent to which income or costs depend on predictability. However, the increase of 0.9 percentage points does not suggest a high improvement of the nominator.

From the sensitivity analysis in Section 8.4 we see that the distance to anchorage is what influences the percentage of interesting instances the most. For the set of instances with shorter anchorage distance the stochastic solution performs better than the optimal deterministic solution for almost half of the instances. Even so, the average SD/Mean ratio is 2.29, just a little higher than the SD/Mean ratio of Instance 46 of 2.74, implying that even though the optimal stochastic solution is the better, there is not too much to gain by choosing it.

From Section 8.6, we saw that the variance of the distribution of route completion times outputted by our model is lower than the true variance. This might indicated that taking uncertainty into account is more valuable than what our results suggests. This is not necessarily the case, as this means that *both* the optimal stochastic and deterministic solutions are subject to more variance than our model says. The reduction in variance from choosing the optimal stochastic solution over the optimal stochastic solution is not necessarily better, but one might expect uncertainty to matter more *often*. In other words, the percentage of interesting instances might be higher than suggested by our model, but the SD/Mean ratio is not necessarily higher.

Chapter 9

Computational Study: Dynamic Version of the Problem

The real-life problem faced by the case company is dynamic by nature; the case company uses information about waiting times and travel times revealed during route execution when deciding what terminal to tender to next. For this reason, it is interesting to examine a dynamic version of the problem, where the revealed information is incorporated in the planning process. In this chapter, the results from solving the dynamic version of the problem are presented. The same mathematical model as presented in Chapter 4 is solved iteratively as new cargoes are serviced, and each iteration is referred to as a stage. The input data used to solve the mathematical model changes slightly between each stage, and the modifications are presented in Section 9.1. Two different planning strategies used to solve the dynamic problem are tested, and the results are presented in Section 9.2. Section 9.3 summarizes the value of using a stochastic dynamic approach.

To solve the dynamic version of the problem, the method and mathematical model used to solve the static version of the problem are used as a basis. In each stage, Nikolova's method is used to identify the optimal *next node* by solving a reduced version of the problem. The problem is reduced based on the chosen next node to be able to find the successive node, and so fourth. The reduced problem solved in a given stage is referred to as a subproblem. This way, the solution to the dynamic version of the problem is found by simulation of the method used to solve the static version of the problem.

Two different planning strategies are used to identify the next node in each

stage. One strategy is to define the optimal route as the route that maximizes the probability of route completion within the given threshold from some given *current node*. The mathematical model used when this planning strategy is applied is exactly the same as presented in Chapter 4, but the input data changes dynamically between each subproblem. This strategy is referred to as the *stochastic planning strategy*. When the second strategy is used, the optimal route is defined as the route that minimizes the expected completion time, which corresponds to solving the deterministic problem. This strategy is referred to as the *deterministic planning strategy*. For both strategies, the next node is the subsequent node from the current node in the optimal route. The routes that are identified from a given stage contain a subset of the initial set of nodes, and all routes in that stage start in the current node and end at anchorage. *Nodes* and *cargoes* are used interchangeably throughout this section.

9.1 Modifications

The same mathematical formulation as presented in Chapter 4 is also used to handle the dynamic version of the problem, but some modifications of the input data between each stage is necessary. Let us, for this explanation, denote two consecutive stages by stages s and s + 1, and the current nodes in the two stages as nodes γ_s and γ_{s+1} , respectively. At stage 0, the current node is anchorage. The set of nodes, $N = \{0, ..., n + 1\}$, from Chapter 4, changes from one stage to another as the problem is reduced. N_s is the set of nodes used when solving the subproblem for stage s. The sets of cargo nodes, $N^C = \{1, ..., n\}$, pickup nodes, $N^+ \subset N^C$, and delivery nodes, $N^- \subset N^C$, from Chapter 4 also change from one subproblem to another, and for this explanation, we use an index to indicate what subproblem the sets apply to. We let N_s^C , N_s^+ and N_s^- be the sets of cargo nodes, pickup nodes and delivery nodes used to solve the subproblem for stage s. The total delivery load from Chapter 4, Q^- , also changes between each stage. We let Q_0^- be the initial total load to be delivered during the entire route, and this is used for the problem solved in stage 0. For the subproblem to be solved in stage 1, the load of the first cargo is added to Q_0^- . This way, Q_s^- is the accumulated load on board the ship, including the load of the current node at stage s, and equals the load when leaving the current node for stage s. Recall that the load belonging to delivery cargoes has negative values, and Q_s^- may either increase or decrease from one stage to the next.

9.1.1 The Stochastic Planning Strategy

For the stochastic planning strategy, the route that maximizes the probability of completing within the given threshold is regarded as optimal. When the current stage is stage s, and the current node is γ_s , Nikolova's method is used to identify an optimal route. γ_s is the first node in the optimal route found for stage s, and the subsequent node after γ_s is referred to as the next node for stage s. The next node is denoted as γ_{s+1} , and γ_{s+1} becomes the current node for stage s + 1. Several modifications of the subproblem from stage s are necessary to obtain the subproblem for stage s + 1; To get the set of nodes used in the subproblem for stage s + 1, node γ_s is removed from N_s . If node γ_{s+1} is a pickup node, it is removed from the set of pickup nodes from the previous stage, N_s^+ , and from the set of delivery nodes from the previous stage, N_s^- , if it is a delivery node. It is, naturally, also removed from the set of cargo nodes, N_s^C . The total delivery load, Q_{s+1}^- , becomes $Q_s^- + Q_{\gamma_{s+1}}$, where $Q_{\gamma_{s+1}}$ is the load of cargo γ_{s+1} . When the next node from stage s has been identified, the realised waiting time is randomly drawn from the distribution of the waiting times at the destination terminal for the corresponding arc, (γ_s, γ_{s+1}) . Equations (6.1) and (6.2) from Chapter 6 are used to calculate the realised sailing time, and the fixed loading and cleaning times of node γ_{s+1} are added to the realised sailing time to get the realised travel time. Recall from Chapter 4 that the stochastic travel time along arc (i, j) can be expressed as

$$\tilde{t}_{ij} = \tilde{s}(\tilde{w})_{ij} + l_j + c_j,$$

where $\tilde{s}(\tilde{w})$ is the stochastic sailing time, l_j the fixed loading time of cargo j and c_j the cleaning time for cargo j. The *realised* travel time along arc (i, j) can be denoted as t_{ij} . The realised travel time along (γ_s, γ_{s+1}) becomes $t_{\gamma_s, \gamma_{s+1}}$. The threshold can be regarded as the remaining time until deadline. From stage s to stage s + 1, the threshold decreases by $t_{\gamma_s, \gamma_{s+1}}$, and we get $H_{s+1} = H_s - t_{\gamma_s, \gamma_{s+1}}$. This way, the threshold from Chapter 4, H, used for the stochastic approach, also changes from one stage to another.

When $N = \{0, ..., n + 1\}$, and the *next node* from stage n - 2, γ_{n-1} , is identified, there is only one remaining cargo node to visit in addition to the anchorage node, n + 1. In other words, we do not need to solve the reduced problems for stages n - 1 and out, as the sequence of the remaining nodes are given by the results from solving the problem in stage n - 2. To get the completion time of the entire route, the realised travel times of the remaining arcs must be found as described above.

For the tests, the applied initial threshold is the threshold that corresponds to a confidence level of $\alpha = 0.95$ for the optimal deterministic route. Recall that the optimal deterministic route is found by solving a static version of the problem as explained in Chapters 5 and 8. This relatively high threshold is used due to the behaviour of the model when the time used during route execution exceeds the threshold. If the threshold is passed at some given stage where all cargoes have not been serviced yet, the optimal route for the remaining stages will have higher variances than desired. When the threshold is passed, and as the threshold is the time remaining until deadline, the threshold becomes negative. As shown in Chapter 4, the objective function of the mathematical model is

$$max \qquad \alpha = P(T \le H) = \Phi(\frac{H - \mu}{\sigma}), \qquad T \sim \mathcal{N}(\mu, \sigma^2)$$

where H is the threshold, T the route completion time, μ the expected route completion time, σ the standard deviation, $\Phi(z)$ the cumulative distribution function (CDF) of the standard normal distribution and $z = \frac{H-\mu}{\sigma}$ in accordance with statistics theory. For H < 0, α increases with increasing σ , implying that routes with higher variances can be favoured. For this reason, we wish to use a high threshold when solving the dynamic problem with the stochastic solution method.

9.1.2 The Deterministic Planning Strategy

For the deterministic planning strategy, the optimal route is defined as the route that minimizes the expected completion time. For the deterministic problem, the travel time of each arc takes the value of the expected completion time of the arc. The expected completion times do not change as information is revealed, and the optimal deterministic route of the dynamic problem is expected to be the same for all simulations. The route will always equal the optimal deterministic route found by solving the static version of the problem. Nevertheless, the realised travel times vary from simulation to simulation.

This planning strategy is also tested to be able to compare the results with the results from the stochastic approach. The modifications done for the stochastic planning strategy also apply to the deterministic planning strategy. The mathematical model of the subproblem to solve at a given stage is very similar to the one described for the static dynamic problem, but with some minor differences. The objective function minimizes the expected completion time, so the threshold is not included. The variance constraint is not necessary.

As for the stochastic planning strategy, when the next node from stage s has been identified, the realised waiting time is randomly drawn from the distribution of the waiting times at the destination terminal for the corresponding arc, (γ_s, γ_{s+1}) . Equations (6.1) and (6.2) from Chapter 6 are used to calculate the realised sailing time, and the fixed loading and cleaning times of node γ_{s+1} are added to the realised realised sailing time to get the realised travel time.

As the route we get when solving the dynamic problem with the deterministic planning strategy always equals the optimal deterministic route found by solving the static version of the problem, we could have found the distribution by a simpler simulation. For the simpler simulation, the realised waiting times are drawn from the distribution of the waiting times at the destination terminal, and the realised sailing time is then calculated according to Equations (6.1) and (6.2). Lastly, the fixed cleaning and loading times for each arc in the given route are added. If this is done 500 times, a distribution of 500 route completion times is obtained, and it is not really necessary to solve the dynamic model. In other words, solving the dynamic problem with the deterministic planning strategy corresponds to doing a simulation similar to what is done in Section 8.6. Solving the dynamic problem 500 times is much more time consuming than to find the distribution by the simpler simulation. But, to verify that the route always equals the optimal deterministic route when solving the dynamic problem with the deterministic solution method, the dynamic problem is solved using the deterministic planning strategy.

9.2 Results

Figures 9.1-9.3 show the distribution of realised route completion times for three different instances where both planning strategies have been applied to solve the dynamic problem. Each distribution is found by simulating the given instance with one of the planning strategies 500 times. In other words, the dynamic problem is solved 500 times for each instance for a given planning strategy. This is done once for the stochastic planning strategy and once for the deterministic planning strategy for each instance, resulting in two distributions for each instance.

Three different instances are presented in this section. The chosen instances are Instance 46 from the Base Set, Instance 26 from Variation Set 2 with Split geography and Instance 5 from the Base Set. Recall from Chapter 8 that the instances in the Base Set have an Even geography, while the instances in Variation Set 2 have a Split geography. The geographies are further explained in Chapter 7. Instance 46 is included in the testing because it is the instance from the Base Set with the highest SD/Mean ratio. Correspondingly, Instance 26 is the instance from Variation Set 2 with the highest SD/Mean ratio out of the instances in this set. Instance 5 from the Base Set is tested because we wish to examine the results from solving the dynamic problem for an instance where the static problem version suggests that the *deterministic* solution is always optimal. In other words, Instance 5 is randomly chosen from the non-interesting instances from the Base Set.



9.2.1 Instance 46 from the Base Set

(c) Information about the two distributions.

Figure 9.1: Distribution of route completion times for the dynamic problem with deterministic and stochastic planning strategies for Instance 46 from the Base Set

Subfigures 9.1a and 9.1b in Figure 9.1 show the distributions of the route completion times we get when solving the dynamic problem with both the deterministic and the stochastic planning strategy for Instance 46 from the Base Set. Table 9.1c summarizes some characteristics and information about the distributions.

Figure 9.1a shows the distribution of the realised route completion times found using the deterministic planning strategy. As the route is the same for each simulation when using the deterministic planning strategy, the variation in route completion times is due to the realised waiting times being randomly drawn.

Figure 9.1b shows the distribution of route completion times using the stochastic planning strategy. For this planning strategy, different routes may be used in different simulations. In addition, as for the deterministic planning strategy, waiting times are randomly drawn from the corresponding distributions, which gives variation in sailing times along each arc from one simulation to another. In other words, the variation of route completion times are both due to different routes being used and different waiting times being drawn.

Figure 9.1 shows that the distribution of route completion times we get from solving the dynamic problem with the deterministic planning strategy has a mean of 373.09 hours and a standard deviation of 10.98 hours. The stochastic planning strategy gives a mean of 373.11 hours with a variance of 95.93 hours. From Section 8.3, we know that our model suggests an expected route completion time of 373.24 hours with a standard deviation of 9.79 hours for the optimal deterministic route, Route 1, for Instance 46. The expected route completion time of the optimal stochastic route, Route 3, outputted by our model for the static version of the problem is 373.49 hours with a standard deviation of 10.25 hours for the same instance. We see that the means of both distributions are almost the same as outputted by the model, and the distribution we get from using the stochastic planning strategy has a slightly higher mean than the distribution we get from using the deterministic planning strategy, but the difference is minor. The stochastic planning strategy gives a lower variance. Note however, that the standard deviations only differ by slightly more than an hour, and when the means are approximately 373 hours, the difference in the standard deviation is not very high. The results from solving the static version of the problem also suggest that the stochastic solution has a slightly higher mean and a somewhat lower variance than the deterministic solution.

The results from solving the dynamic problem for Instance 46, shown in Figure 9.1, correspond with the results from solving the static version of the problem. They suggest that the preferred route depends on the risk profile of the decision maker, but that the deterministic solution, for most practical purposes, performs good enough.

9.2.2 Instance 26 from Variation Set 2 with Split Geography

Subfigures 9.2a and 9.2b in Figure 9.2 show the distributions of route completion times we get when solving the dynamic problem with both the deterministic and stochastic planning strategy for Instance 26 from Variation Set 2 with Split geography. Table 9.2c summarizes some characteristics and information about the distributions.

As for Instance 46, the stochastic planning strategy gives a distribution with a slightly higher mean and a lower variance than the distribution found using the deterministic planning strategy. The difference in the standard deviation is minor. The results from solving the static version of Instance 26 in Variation Set 2 suggest that the optimal deterministic route has a mean of 429.97 hours



(c) Information about the two distributions.

Figure 9.2: Distribution of route completion times for the dynamic problem with deterministic and stochastic planning strategies for Instance 26 from Variation Set 2.

and a standard deviation of 12.68 hours, which is not far from the results from solving the dynamic problem with the deterministic planning strategy, shown in Subfigure 9.2a. The optimal stochastic route found by solving the static problem has a mean of 430 hours and a standard deviation of 12.17 hours. The results from solving the dynamic model with the stochastic planning strategy, shown in Subfigure 9.2b, also correspond well with the results from solving the static problem. The results from both versions of the problem suggest that the optimal stochastic route has a slightly higher mean and lower variance than the deterministic solution. What planning strategy to use depends on the risk profile of the decision maker, bur for most practical purposes, the deterministic solution performs good enough.

9.2.3 Instance 5 from the Base Set

Subfigures 9.3a and 9.3b in Figure 9.3 show the distributions of the route completion times we get when solving the dynamic problem with both the



(c) Information about the two distributions.

Figure 9.3: Distribution of route completion times for the dynamic problem with deterministic and stochastic planning strategies for Instance 5 from the Base Set. This instance is not among the interesting instances, meaning the results from solving the static equivalent suggests the deterministic solution is optimal.

deterministic and stochastic planning strategy for Instance 5 from the Base Set. Table 9.3c summarizes some characteristics and information about the distributions.

Recall that Instance 5 is one of the non-interesting instances from the Base Set, meaning that the deterministic solution is optimal based on the results from solving the static version of the problem. The results from solving the static version of the problem show that the optimal deterministic route has a mean of 382.26 hours and a standard deviation of 11.36 hours. The mean outputted by the model from solving the static problem corresponds well with the mean of the distributions in Figure 9.3, and the difference in the standard deviations are minor. Table 9.3c shows that using the stochastic planning strategy gives a slightly higher variance of the distribution than when using the deterministic planning strategy. This is opposite of what we see for Instance 46 and Instance 26 from Variation Set 2. But, as results from solving the static version of the problem shows that the deterministic solution is also optimal when using a stochastic approach, this is not surprising. Again, we see that the results from solving the static version of the problem correspond well with the results from solving the dynamic version of the problem.

9.3 Practical Implications

The results from solving the dynamic version of the problem correspond well with the results presented in Chapter 8, which are obtained by solving the static version of the problem. For instances where the optimal solution of the static stochastic problem differs from the optimal deterministic solution (i.e. the interesting instances), the results show that using the stochastic planning strategy when solving the dynamic problem gives a slightly higher mean, but lower threshold than when using the deterministic planning strategy. However, the differences are minor. What solution approach is preferred depends on the risk profile of the decision maker, but our results suggest that, for most practical purposes, the deterministic solution performs good enough. The optimal deterministic route for the static version of the problem equals the optimal route we obtain when solving the dynamic version of the problem with the deterministic planing strategy. Identifying the optimal deterministic route using the static version of the problem only takes a few seconds, so does the simple simulation of 500 route completion times for a given deterministic route. Solving the dynamic problem with the stochastic planning strategy 500 times took 11.2 hours for Instance 46 from the Base Set, 7.9 hours for Instance 26 from Variation Set 2 with Split geography, and 7.6 hours for Instance 5 from the Base Set. This suggests that the gain from using a stochastic approach does not justify the long solution time.

Chapter 10

Concluding Remarks

This thesis addresses the in-port routing of a chemical tanker. We consider a single ship which has to pick up and deliver a given number of cargoes located at different terminals while complying with capacity and draft limit constraints. The problem to be solved is deciding the sequence for servicing cargoes. The time it takes before terminals are ready to accommodate the ship is uncertain, and the waiting times at the port's terminals are hence stochastic. When the waiting times at terminals are stochastic, the travel times between terminals also become stochastic. This makes the problem a pickup and delivery problem subject to constraining draft limits and stochastic travel times (PDP-DLST). To our knowledge, this problem has not been explicitly studied before. The aim is to investigate the possible benefits of including uncertainty in the in-port route-planning problem, and to understand the behaviour of the stochastic variables.

As the case port is particularly long and narrow, the movement of the ship can be considered as movement along a straight line. This means that when the ship sails between terminals, it either sails towards anchorage or away from it. When the ship is finished servicing cargoes in a terminal, it is not allowed to wait by the terminal, and must sail towards anchorage until the next terminal is ready to accommodate it. This may either involve that the ship turns around on the spot before reaching anchorage, or that the ship waits at anchorage for some amount of time.

Both a static and dynamic version of the problem are solved, and a mathematical model is developed to find the route that maximizes the probability of route completion within a given threshold. Travel times between terminals depend directly on the waiting time at the destination terminal, arc direction relative to anchorage and distances between the terminals and to anchorage. Waiting times are assumed normally distributed, but are truncated at zero as negative values for waiting times is not realistic. The distributions of the stochastic travel times are approximated by normal distributions to be able to solve the optimization problem. The chosen solution method identifies a set of possibly optimal routes which have to be further analysed to be able to identify the optimal route for a given threshold. What level of required confidence level is used to identify the optimal route depends on the risk profile of the decision maker. We assume a relatively risk averse decision maker.

10.1 Summary and Conclusion

In this section, we summarize the key findings from the results presented in Chapters 8 and 9, and conclude on the importance of accounting for uncertainty when solving the problem presented in Chapter 2. The results for the static version of the problem are found by solving sets of 100 test instances with 20 cargoes and six terminals. The behaviour of the uncertain variables has been examined, and reasons to why routes have different variances are identified. For the dynamic version of the problem, a few chosen test instances are used to examine the possible gains from accounting for uncertainty in the dynamic planning process.

The possibility of sailing directly to the destination terminal for arcs directed towards anchorage, and the possibility for the ship to wait at anchorage gives special distributions of the arc travel times. The special distributions give rise to different effects on arcs' variances, depending on the unique combination of arc direction, distribution of the waiting time at the destination terminal, distances between the terminals and the distance from anchorage to the terminal closest to anchorage. Differences between variances associated with different arcs sometimes stem from these effects, but are more often due to different distributions of the stochastic waiting times at the destination terminals. The approximation used to model the distribution of the stochastic travel times is evaluated in Section 8.6, and our results show that the distribution of route completion times resemble normal distributions when waiting times are normally distributed.

The computational study shows that including uncertainty in the planning problem results in a different solution than the optimal deterministic solution for 18% of the tested instances. This is based on a required confidence level of $\alpha = 95\%$ for a route to be regarded as optimal, which represents a relatively risk averse decision maker. For a required confidence level of 60%, i.e. for less risk averse decision makers, the deterministic solution is always optimal.

The sensitivity analysis shows that when specific port characteristics are changed, taking uncertainty into account becomes more important. Decreasing the distance to anchorage especially increases the percentage of instances where the optimal stochastic solution performs better than the optimal deterministic solution, meaning taking uncertainty into account matters more often. Changing the draft limits also affects the results. The sensitivity analysis shows that only a slight change of the draft limits affects the results to a rather large extent. Almost 50% of the instances from the Base Set become infeasible when the range of the draft limits is changed from 50%-90% of total ship capacity to 40%-80% of total ship capacity. The same change causes the solution time (for solving a set of 100 instances with 20 cargoes) to increase from 20 minutes to more than two hours.

For the applied confidence level, the optimal route is the route which, for this confidence level, has the lowest corresponding threshold. For instances where the optimal stochastic route differs from the optimal deterministic route, our results suggest that the improvement of the threshold is less than 0.5% when choosing the optimal stochastic solution over the optimal deterministic solution. This does not represent a significant gain from using a stochastic approach instead of a deterministic approach. But, even if using a stochastic approach does not result in significant economical benefits for the shipping company in Houston Ship Channel, the stochastic model provides information about the risk levels associated with the different route options. From a business perspective, this information can be valuable input in risk management. It also allows decision makers to better plan for uncertainty, and lets them know what level of risk they are operating at.

From Sections 8.6 and 8.7, we see that the variance associated with the distribution of the route completion times is *in reality* larger than suggested by our model. This difference in variance is due to the applied approximation of the distribution of the stochastic sailing times. This suggests that, in reality, there is more variance associated with the problem studied than what is suggested by our model. As such, it *could* be that the value of accounting for uncertainty when solving the studied problem is higher than what our results suggest. As the variance for both the optimal stochastic and deterministic solutions are higher in reality, one might expect uncertainty to matter more *often*, but not necessarily to a greater extent per instance.

The results from solving the dynamic version of the problem correspond well with the results from solving the static version of the problem. Our results suggest that the gains from using a stochastic approach do not justify the high solution times. Theoretically, a very risk a version decision maker might prefer using a stochastic approach, but for all practical purposes, the results from solving both the static and the dynamic problem indicate that the deterministic solution performs good enough.

10.2 Future Research

Several possible improvements of the model and solution approach have been identified. In this section, suggestions for future research are presented. The first thing to notice is that many simplifications of reality are made, and extensions of the model can be done to represent the reality more accurately. Such extensions can be to include soft time windows for customer servicing and to include tank allocation on the ship. The choice of which cargoes to service, i.e. including spot cargoes, will also extend the models to represent reality more accurately. In reality, the ship is allowed to send NORs to more than one terminal at the time, and out of the terminals that have been tendered to, the ship must sail to the first available terminal. This entails a strategic element, and tendering strategies could also have been included to model the problem more realistically.

The distributions of the stochastic waiting times at terminals are modelled as normal distributions. The means are drawn at random from the range from 0 to 10 hours, and the standard deviation is set to 50% of the corresponding mean. Even if the range from 0 to 10 hours is based on a set of real data, the distribution of the waiting times can be made more realistic. This might involve larger sets of historical data and forecasting. The approximation of the distributions of the sailing times can also be improved or handled differently in the model. One suggestion is to better manage the skewness of the distributions, and to include this characteristic when summing the distributions of arc travel times to obtain the accumulated distribution of the complete route.

From the computational study, we see that the differences between waiting time distributions is often the source to different variances of arc travel times. We suggest modelling the problem with equal distributions of the waiting times at terminals, just to better extract information about other sources to different variances of arc travel times.

Future work may also focus more on technical analysis and improvement. For this thesis, the focus has not been on technical performance of the models, but rather on practical implications. With technical improvements, the models may be able to solve larger instances.

More sophistical analysis of economical and environmental effects of modelling with, and planning for, uncertainty also remains for future work. By including demurrage costs, port charges, commissions, fees and other voyage related costs, the economical implications can be analyzed and used to better support decision making.

Bibliography

- Adams, W. J. (2009). The life and times of the central limit theorem. Vol. 35. American Mathematical Soc.
- Anily, S. and G. Mosheiov (1994). "The traveling salesman problem with delivery and backhauls". In: Operations Research Letters 16 (1), pp. 11– 18.
- Arnesen, M. J. and M. Gjestvang (2015). "Routing and tank allocation for a chemical tanker". MA thesis. Norwegian University of Science and Technology.
- Baldacci, R., E. Hadjiconstantinou, and A. Mingozzi (2003). "An exact algorithm for the traveling salesman problem with deliveries and collections". In: Networks 42 (1), pp. 26–41.
- Berbeglia, G., J. F. Cordeau, I. Gribkovskaia, and G. Laporte (2007). "Static pickup and delivery problems: a classification scheme and survey". In: *TOP* 15, pp. 1–31.
- Carraway, R. L., T. L. Morin, and H. Moskowitz (1989). "Generalized dynamic programming for stochastic combinatorial optimization". In: Operation Research 37, pp. 819–829.
- Charnes, A. and W. W. Cooper (1959). "Chance-Constrained Programming". In: Management Science 6 (1), pp. 73–79.
- Christiansen, M., K. Fagerholt, B. Nygreen, and D. Ronen (2007). "Maritime transportation". In: *Handbooks in operations research and management science.* Ed. by C. Bernhart and G. Laporte. Vol. 14.4, pp. 189–284.
- Dantzig, G. B. and J. H. Ramser (1959). "The truck dispatching problem". In: Management Science 6, pp. 80–91.
- Dudley, R. M. (1999). Uniform central limit theorems. Vol. 23. Cambridge Univ Press.
- Ehmke, J. F., A. M. Campbell, and T. L. Urban (2015). "Ensuring service levels in routing problems with time windows and stochastic travel times". In: *European Journal of Operations Research* 240 (2), pp. 539–550.
- Fan, Y. Y., R. E. Kalaba, and J. E. Moore II (2005a). "Arriving on time". In: Journal of Optimization Theory and Applications 127.3, pp. 497–513.
- Fan, Y. Y., R. E. Kalaba, and J. E. Moore (2005b). "Arriving on Time". In: Journal of Optimization Theory and Applications 3 (127), pp. 497–513.

- Frank, H. (1969). "Shortest paths in probabilistic graphs". In: Operations Research 17.4, pp. 583–599.
- Fu, L. and L. R. Rilett (1998). "Expected shortest paths in dynamic and stochastic traffic networks". In: *Transportation Research Part B: Method*ological 32.7, pp. 499–516.
- Gendreau, M., G. Laporte, and R. Séguin (1996). "Stochastic vehicle routing". In: European Journal of Operational Research 88.1, pp. 3–12.
- Gendreau, M., G. Laporte, and D. Vigo (1999). "Heuristics for the travelling salesman problem with pickup and deliver". In: Computer and Operations Research 26, pp. 699–714.
- Kao, E. P. C. (1978). "A preference order dynamic dynamic program for a stochastic traveling salesman problem". In: *Operations Research* 26, pp. 1033–1045.
- Kaparias, I., M. Bell, and H. Belzner (2008). "A new measure of travel time reliability for in-vehicle navigation systems". In: *Journal of Intelligent Transportation Systems* 12 (4), pp. 202–211.
- Kenyon, A. S. and D. P. Morton (2003). "Stochastic Vehicle Routing with Random Travel Times". In: *Transportation Science*, pp. 69–82.
- Kruse, J. C. (2015). Analysis of Chemical Tanker Transits on the Houston Ship Channel. Tech. rep. Texas A & M Transportation Institute.
- Laporte, G., F. Louveaux, and H. Mercure (1992). "The vehicle routing problem with stochastic travel times". In: *Transportation science* 26.3, pp. 161– 170.
- Lawrence, S. A. (1972). *International sea transport: the years ahead*. Lexington Books.
- Li, X., P. Tian, and S. C. H. Leung (2010). "Vehicle routing problems with time windows and stochastic travel times: Models and algorithms". In: *International Journal of Production Economics* 125 (1), pp. 137–145.
- Loui, R. P. (1983). "Optimal paths in graphs with stochastic or multidimensional weights". In: Communications of the ACM 26.9, pp. 670–676.
- Mazmanyan, L. and D. Trietsch (2013). "Stochastic travelling salesperson and shortest route models with safety time". In: International Journal of Planning and Scheduling 2 (1).
- Meng, Q. and T. Wang (2010). "A chance constrained programming model for short-term liner ship fleet planning problems". In: Maritime Policy & Management: The flagship journal of international shipping and port research 37 (4), pp. 329–346.
- Miller-Hooks, E. D. and H. S. Mahmassani (2000). "Least expected time paths in stochastic, time-varying transportation networks". In: *Transportation Science* 34.2, pp. 198–215.
- Mirchandani, P. B. (1976). "Shortest distance and reliability of probabilistic networks". In: Computers & Operations Research 3.4, pp. 347–355.
- Mosheiov, G. (1994). "The Travelling Salesman Problem with pick-up and delivery". In: European Journal of Operations Research 72 (2), pp. 299– 310.
- Nahum, O. E. and Y. Hadas (2009). "Developing a Model for the Stochastic Time-dependent Vehicle-routing Problem". In: Computers & Industrial Engineering, 2009. CIE 2009, pp. 118–123.
- Nie, Y. and X. Wu (2009). "Shortest path problem considering on-time arrival probability". In: Transportation Research Part B: Methodological 43.6, pp. 597–613.
- Nikolova, E. (2009). "High-performance heuristics for optimization in stochastic traffic engineering problems". In: Large-Scale Scientific Computing. Springer, pp. 352–360.
- Nikolova, E., J. A. Kelner, M. Brand, and M. Mitzenmacher (2006). "Stochastic shortest paths via quasi-convex maximization". In: Algorithms-ESA 2006. Springer, pp. 552–563.
- Odfjell (2003). The Chemical Tanker Market. Tech. rep. Odfjell.
- (2014). OdfjellTankers. Internal Document, Power Point Presentation.
- Pillac, V., M. Gendreau, C. Guéret, and A. L. Medaglia (2013). "A review of dynamic vehicle routing problems". In: *European Journal of Operational Research* 225.1, pp. 1–11.
- Psaraftis, H. N. (1995). "Dynamic vehicle routing: Status and prospects". In: annals of Operations Research 61.1, pp. 143–164.
- Rakke, J. G., M. Christiansen, K. Fagerholt, and G. Laporte (2012). "The traveling salesman problem with draft limits". In: Computers & Operations Research 39.9, pp. 2161–2167.
- Russell, R. A. and T. L. Urban (2008). "Vehicle routing with soft time windows and Erlang travel times". In: *Journal of Operations Research Society* 59, pp. 1220–1228.
- Samaranayake, S., S. Blandin, and A. Bayen (2011). "A tractable class of algorithms for reliable routing in stochastic networks". In: *Procedia-Social* and Behavioral Sciences 17, pp. 341–363.
- Sniedovich, M. (1981). "Analysis of a preference order traveling salesman problem". In: Operations Research 29, pp. 1234–1237.
- Süral, H. and J.H. Bookbinder (2003). "The single-vehicle routing problem with unrestricted backhauls". In: *Networks* 41.3, pp. 127–136.
- Swihart, M. R. and J. D. Papastavrou (1999). "A stochastic and dynamic model for the single-vehicle pick-up and delivery problem". In: *European Journal of Operational Research* 114.3, pp. 447–464.
- Taş, D., N. Dellaert, T. van Woensel, and T. de Kok (2013). "Vehicle routing problem with stochastic travel times including soft time windows and service costs". In: Computers & Operations Research 40 (1), pp. 214–224.
- UNCTAD (2015). *Review of Maritime Transport 2015*. Chap. 1: Developments in International Seaborne Trade.
- Verweij, B., S. Ahmed, A. J. Kleywegt, G. Nemhauser, and A. Shapiro (2003). "The sample average approximation method applied to stochastic routing problems: a computational study". In: *Computational Optimization and Applications* 24.2-3, pp. 289–333.

- Waller, S. T. and A. K. Ziliaskopoulos (2002). "On the online shortest path problem with limited arc cost dependencies". In: *Networks* 40.4, pp. 216–227.
- Williams, H. P. (2013). Model building in mathematical programming. John Wiley & Sons.

Appendix A

Model Formulation

Sets and Indices

- Nset of nodes, indexed by i and j
- N^C set of all cargoes, indexed by i and j
- N^+ set of pickup nodes, indexed by i and j
- N^{-} set of delivery nodes, indexed by i and j
- Α set of feasible arcs (i, j) between nodes i and j

Parameters

- Q^+ the total load that is going to be picked up
- $Q^$ the total load that is going to be delivered
- Q_i load in node i
- Ktotal capacity of the ship in tonnes
- D_i draft limit, in tonnes, of the terminal associated with cargo i
- Hthreshold for the completion time of the route
- expected value of the travel time of arc (i, j) $\begin{array}{c} \mu_{ij} \\ \sigma_{ij}^2 \end{array}$
- variance of the travel time of arc (i, j)

Variables

- = 1 if the ship sails directly from node *i* to *j*, 0 otherwise x_{ij}
- load on board the ship on arc (i, j) y_{ij}
- expected value of the route completion time μ
- σ^2 variance of the route completion time
- Т route completion time

Objective function

maximize
$$\alpha = P(T \le H) = \Phi(\frac{H-\mu}{\sigma}), \quad T \sim \mathcal{N}(\mu, \sigma^2)$$

Flow Constraints

$$\sum_{\substack{j \in N \\ j \in N | (i,j) \in A}} x_{0j} = 1$$

$$\sum_{\substack{j \in N | (i,j) \in A \\ i \in N | (i,j) \in A}} x_{ij} = 1, \quad i \in N^C$$

$$\sum_{\substack{i \in N \\ i \in N}} x_{i,n+1} = 1$$

Cargo Constraints

$$\sum_{\substack{j \in N^C \\ (i,j) \in A}} y_{0j} = Q^-$$
$$\sum_{\substack{(j,j) \in A}} y_{ji} = -Q_j, \qquad j \in N^C$$

Capacity Constraints

$$\begin{array}{ll} 0 \leq y_{0j} \leq Q^{-}x_{0j}, & (0,j) \in A \mid j \in N \setminus \{0\} \\ 0 \leq y_{i,n+1} \leq Q^{+}x_{i,n+1}, & (i,n+1) \in A \mid i \in N \setminus \{n+1\} \\ Q_{i}x_{ij} \leq y_{ij} \leq (K-Q_{j})x_{ij}, & (i,j) \in A \mid i,j \in N^{+} \\ (Q_{i}-Q_{j})x_{ij} \leq y_{ij} \leq Kx_{ij}, & (i,j) \in A \mid i \in N^{+}, j \in N^{-} \\ -Q_{j}x_{ij} \leq y_{ij} \leq (K+Q_{i})x_{ij}, & (i,j) \in A \mid i,j \in N^{-} \\ 0 \leq y_{ij} \leq (K-Q_{j})x_{ij}, & (i,j) \in A \mid i \in N^{-}, j \in N^{+} \end{array}$$

Draft Limit Constraints

$$\begin{aligned} 0 &\leq y_{ij} \leq D_j x_{ij}, \\ 0 &\leq y_{ij} \leq D_i x_{ij}, \end{aligned} \qquad (i,j) \in A \mid j \in N^- \\ (i,j) \in A \mid i \in N^+ \end{aligned}$$

Time Constraint

$$\mu = \sum_{(i,j) \in A} \mu_{ij} x_{ij}$$

Variance Constraint

$$\sum_{(i,j)\in A} \sigma_{ij}^2 x_{ij} = \sigma^2$$

Integer Constraint

$$x_{ij} \in \{0, 1\},$$
 $(i, j) \in A$

Subtour Eliminating Constraints

$$x_{ij} + x_{ji} \le 1, \qquad (i,j) \in A$$