

Jørgen G. Rakke

Optimization Models and Methods for Maritime Routing and Scheduling Problems

Thesis for the degree of Philosophiae Doctor

Trondheim, May 2012

Norwegian University of Science and Technology
Faculty of Social Sciences and Technology Management
Department of Industrial Economics and Technology
Management



NTNU – Trondheim
Norwegian University of
Science and Technology

NTNU

Norwegian University of Science and Technology

Thesis for the degree of Philosophiae Doctor

Faculty of Social Sciences and Technology Management

Department of Industrial Economics and Technology Management

© Jørgen G. Rakke

ISBN 978-82-471-3564-8 (printed ver.)

ISBN 978-82-471-3565-5 (electronic ver.)

ISSN 1503-8181

Doctoral theses at NTNU, 2012:140

Printed by NTNU-trykk

Acknowledgements


This thesis is the result of three and a half years of intensive research of Operations Research in maritime transportation. During this period I have had the opportunity to work with some of the most skilled persons in our research area, which has been a really great experience. Thanks to these people I have managed to stay motivated, even when I felt total stagnation in my research.

First, I want to thank my supervisor Professor Marielle Christiansen. It has been a pleasure working with you. You have done a great job facilitating my learning during my research and introducing me to people who I learned so much from. With out your contributions, both as a supervisor and colleague, this thesis would not have been as good as it is today. Thanks are also due to my co-supervisors Bjørn Nygreen and Henrik Andersson. You both have greatly contributed to making this thesis what it is today. Few people have as detailed knowledge about how solvers work as you Bjørn, and you shared this knowledge with great enthusiasm. Henrik, the door to your office was always open, and you were always ready to help. Even when you yourself had way too much to do you took the time to help both with research questions and problems of a more technical nature. I would also like to thank my colleagues at the department at NTNU and at MARINTEK. In particular I would like to thank Kjetil Fagerholt, Matthias Nowak and Magnus Stålhane. No matter how trivial or interesting the question was, you were always there to answer it. For this I am deeply grateful.

I would also like to express my deepest gratitude to Professor Guy Desaulniers and Professor Gilbert Laporte for a very pleasant and helpful co-operation. Your guidance and knowledge was vital to solve the problems I was working on. I would also like to thank Professor Guy Desaulniers for the supervision, help, and support during my stays in Montréal. Your enthusiasm and dedication was overwhelming, and I hope I get to return some of this help and support someday.

My special thanks go to Ann-Mari who always was there to encourage and support me. Even when I get back from work to wake you in the middle of the night. Finally I would like to thank my family and friends for always supporting me and being there for me. Especially my parents for encouraging me to always follow my own path.

Trondheim, March, 2012



Jørgen G. Rakke

Contents

1 Introduction	1
1.1 Background	2
1.2 Purpose and Outline of Thesis	9
1.3 Contributions	13
1.4 Concluding Remarks	15
2 A Rolling Horizon Heuristic for Creating a Liquefied Natural Gas Annual Delivery Program	25
2.1 Introduction	25
2.2 The LNG Supply Chain and the Planning Levels within the LNG Business	28
2.3 Related Literature	30
2.4 Problem Description	33
2.5 The Mathematical Model	35
2.6 Solution Method	38
2.7 Computational Study	44
2.8 Concluding Remarks	51
3 A Construction and Improvement Heuristic for a Liquefied Natural Gas Inventory Routing Problem	61
3.1 Introduction	61
3.2 Related literature in maritime LNG transportation	63
3.3 Problem description	64
3.4 Mathematical model	66
3.5 Solution approach	69
3.6 Computational results	77
3.7 Concluding remarks	85
4 Creating Annual Delivery Programs of Liquefied Natural Gas	91
4.1 Introduction	91
4.2 Literature	95
4.3 Problem description	96
4.4 Model	99
4.5 Strengthening the formulation	101

4.6	Computational study	106
4.7	Concluding remarks	108
5	Branch-Price-and-Cut for Creating an Annual Delivery Program of Multi-Product Liquefied Natural Gas	115
5.1	Introduction	115
5.2	Problem definition and notation	118
5.3	Mathematical Formulations	120
5.4	Solution Method for the PBF	131
5.5	Computational Study	138
5.6	Concluding Remarks	144
6	The Traveling Salesman Problem with Draft Limits	151
6.1	Introduction	151
6.2	Mathematical models	153
6.3	Branch-and-cut algorithm	158
6.4	Computational study	158
6.5	Concluding remarks	163

Introduction

This thesis discusses how Operations Research may be used to offer support to ship operators dealing with maritime routing and scheduling problems. The main focus is on the modeling of the transportation network for the inventory and maritime transportation aspects of the supply chain. In particular, it focus on the routing and scheduling of large heterogeneous fleets of ships for producers of Liquefied Natural Gas (LNG). The thesis consists of five papers; four discussing the distribution of LNG, and one focusing on a subproblem in petroleum product distribution.

The research leading to this thesis is part of a larger research project at NTNU, OPTIMAR – Optimization in Maritime Transportation and Logistics, (OPTIMAR, 2005). This project’s aim was to develop models and efficient solution methods for challenges within ship routing and scheduling in order to increase ship utilization to reduce transportation and logistic costs, as well as reducing the environmental impact and improving customer service. A second key objective of the OPTIMAR project was to create a good foundation for commercializing optimization based decision support systems (DSS). Hence, the main objectives was to create the methodological basis and increase knowledge for solving complex ship routing and scheduling problems. This thesis is a part of creating this methodological basis by developing optimization models and methods. And through the discussions regarding these models and methods it helps expanding the knowledge we have of complex maritime routing and scheduling problems.

The first four papers in this thesis discuss an inventory routing problem from one of the worlds biggest producers of LNG. A thorough description of the problem and a mathematical model for creating an annual delivery program (ADP) is presented. In the first and second paper two different heuristics for the problem are developed. The third paper describes a branch-and-cut method. In the last paper a branch-price-and-cut approach for LNG inventory routing problems is described, introducing a new decomposition approach to get better bounds in order to evaluate the performance of the heuristics. The fifth paper introduces a new Traveling Salesman Problem (TSP) to the literature, namely the TSP with Draft Limits (TSPDL). This problem is a subproblem in a larger maritime supply chain for petroleum products.

1.1 Background

In this section an introduction to relevant theory for the topics of this thesis is given. Section 1.1.1 briefly presents definitions for maritime transportation problems, and relevant literature on maritime routing and scheduling problems. In this section the focus is on different problem types, not the solution methods. Section 1.1.2 gives a short overview of the exact solution methods used, while Section 1.1.3 gives a brief introduction to some of the most commonly used heuristics.

1.1.1 Maritime Routing and Scheduling

Traditionally maritime transportation is divided into three modes of operation; liner, tramp, and industrial shipping (Lawrence, 1972). In liner shipping the ships operate according to a published itinerary and schedule similar to a bus, and the operator tries to maximize their profit from cargoes lifted. Using the same type of analogy, tramp shipping resembles taxi services, where the ships follow the available cargoes. The objective is to maximize the profit from mandatory (cargoes with a contract of affreightment) and optional (spot) cargoes, with mandatory cargoes covering the daily operating costs and optional cargoes being the major contributor to the profit of the operator. Industrial shipping operators usually own and control both the cargoes and the ships, and the objective of the operator is usually to minimize transportation cost (Christiansen, Fagerholt, Nygreen, and Ronen, 2007). Today the boundaries between tramp and industrial shipping are less clear than they used to be. For instance, many vertically integrated companies, e.g. in the LNG industry, don't fit any of these classes. They operate all (most) vessels picking up LNG at their liquefaction terminal(s), but they only own part of the fleet. The main share of the volume of gas is delivered to long-term customers, but there is also the option of selling some LNG on the spot market to increase profit.

Maritime transportation problems can also be classified based on the length of the planning horizon. Traditionally this divides the problems into three classes; strategic, tactical, and operational. Strategic problems are long term planning problems. In shipping this usually means 5 to 20 years. For instance *fleet size and mix*, *network design*, *ship design*, and *contract evaluation* are all typical strategic planning problems. Problems with a medium planning horizon length are usually referred to as tactical planning problems. These problems usually have a planning horizon of a few week up to 18 months. Problems such as *ship routing and scheduling*, *fleet deployment*, *inventory ship routing* are important tactical planning problems. Short-term, or operational, planning is usually applied when the decisions only have a short-term impact (sometimes as short as only one sailing leg) or the operational environment is highly uncertain. Decisions such as *ship*

loading, environmental routing, speed optimization, and booking of single orders are all important operational planning problems. In shipping the boundaries between tactical and operational, and strategic and tactical can be somewhat fuzzy. For instance ship routing and scheduling can be used in fleet size and mix problems, but usually with less details, and speed optimization can be an integrated part of a routing and scheduling problem.

Given the scope of this thesis the rest of this introduction will focus on maritime routing and scheduling for tramp and industrial shipping. Liner shipping is omitted since it differs significantly from the two other modes of operation and the problems studied in this thesis. For a more comprehensive review on maritime transportation in general see Christiansen et al. (2007).

Al-Khayyal and Hwang (2007) proposed to separate maritime routing and scheduling problems for bulk shipping into two categories; *cargo routing* and *inventory routing* problems. The authors described cargo routing problems as problems "mainly constrained by the cargo, which is usually specified by the loading/discharge ports, and by the time windows of loading and unloading". Inventory routing problems are mostly constrained by the inventory levels in either loading or unloading ports, or both.

The basic cargo routing problems can be seen as a maritime variant of the land based *pickup and delivery problem with time windows* described in for instance Dumas, Desrosiers, and Soumis (1991) and Desrosiers, Dumas, Solomon, and Soumis (1995). What differentiates the maritime-PDPTW from the land based PDPTW is that the set of ships often is heterogeneous and fixed, and that no central depot exists. The planning period for maritime-PDPTWs is also usually longer, and the vessels sail around the clock. In the basic maritime-PDPTW only one product is considered, but several cargoes are allowed to be onboard the ship at the same time as long as the capacity of the ship is respected. One cargo consists of a designated number of units and have time windows both for pickup and delivery. The operator tries to minimize the total cost, while ensuring that all contractual cargoes are picked up. The basic maritime-PDPTW is rare in real world cargo routing problems. It is usually extended with one or several constraints either on the cargoes or the ports, or both.

In the following some examples of the different extensions to the maritime-PDPTW are presented. One common extension to the maritime-PDPTW is optional cargoes as in the pioneer work of Appelgren (1969, 1971). This problem also differs from the maritime-PDPTW since the ships are restricted to carry only one cargo at the time. Full ship loads are also discussed in Brown, Graves, and Ronen (1987) where the authors describe a model for scheduling ocean transportation of crude oil. Later, Fisher and Rosenwein (1989) extended this model to involve the designated amount of a cargo to be picked up at several loading ports and delivered to several unloading ports.

A variant of the maritime-PDPTW involving both optional cargoes and multiple cargoes onboard is discussed in Brønmo, Christiansen, Fagerholt, and Nygreen (2007a) and Korsvik, Fagerholt, and Laporte (2009). In this model any of the optional cargoes can be accepted or rejected as long as all the contractual cargoes are lifted. Fagerholt, Hvattum, Johnsen, and Korsvik (2011) extend this model by introducing groups of cargoes that have to be accepted or rejected. These groups of cargoes are defined as coupled cargoes. The authors also consider stowage constraints on the cargoes. This cargo coupling is extended further by Andersson, Duesund, and Fagerholt (2011b). The authors introduce additional synchronization constraints on the coupled cargoes, restricting all the coupled cargoes to be delivered within a certain time interval of the first cargo delivery. The stowage constraints from Fagerholt et al. (2011) are not considered in this paper. A cargo routing problem with flexible cargo sizes is discussed in Brønmo, Christiansen, and Nygreen (2007b), Brønmo, Nygreen, and Lysgaard (2010), and Korsvik and Fagerholt (2010). Instead of fixed cargo quantities, the shipping operator is able to adjust the cargo quantities within a given interval. Brønmo et al. (2007b) propose a set partitioning approach with pre-generated columns, while the same authors propose a heuristic column generation approach in order to solve larger instances in Brønmo et al. (2010), and a tabu search heuristic is developed by Korsvik and Fagerholt (2010). In the maritime-PDPTW with split loads presented by Korsvik, Fagerholt, and Laporte (2011), a cargo may be split over several ships. The same problem was also studied in Andersson, Christiansen, and Fagerholt (2011a).

Bausch, Brown, and Ronen (1998) discuss the short-term distribution shipping of multiple liquid bulk products using vessels with multiple compartments. A similar problem can also be found in e.g. Scott (1995) and Sherali, Al-Yakoob, and Hassan (1999). The cargo holds in the previous problem have fixed sizes, and the setup can not be reconfigured during the planning horizon. Fagerholt and Christiansen (2000a) and Fagerholt and Christiansen (2000b) extend this model to handle ships with flexible cargo holds.

One major cost component for shipping operators is fuel costs, and as a result of the high fuel prices in the 1970s Ronen (1982) presented a model for optimizing speed on a voyage leg to reduce cost. Lately, there has been an increased focus on making shipping more environmentally friendly, and in this context Fagerholt, Laporte, and Norstad (2009b) present a model and solution method for reducing fuel emissions by optimizing the speed of ship routes.

The *Inventory routing problem* (IRP) is defined by Dror and Ball (1987) as a distribution problem for a set of customers, where each customer maintains a local inventory of a product, such as heating oil or methane, and consumes a certain amount of that product every day. A central supplier (depot) tries to minimize the transportation costs, while ensuring that no customer runs out

of product at any time. An overview of maritime inventory routing problems by Christiansen and Fagerholt (2009) points out that there are usually several differences between road based IRPs and *maritime inventory routing problems* (MIRP)s. Road based IRPs often have one central depot and many customers consuming a product, while in MIRPs one might have several production and consumption ports. Further, the amount unloaded at each customer in road based IRPs is often small compared to the capacity of the vehicle, while in MIRPs the unloaded amount is often a full ship load.

The *basic* MIRP is defined as a problem of transporting a single product from one or more production (loading) port(s) to a set of customers (unloading ports). In this problem everything is loaded before starting to unload. The ship is then completely emptied before a new loading process can be started. We find many extensions to this basic problem. These extensions may be structured, in a non-exhaustive way, by e.g. *port structure* and *loading/unloading policy, production/consumption*, and *products*. In real world applications of maritime inventory routing problems several of these extensions may apply. Many different port structures and loading/unloading policies are discussed in the literature. An inventory routing problem with one production port, and several consumption ports are discussed in this thesis by Rakke, Stålhane, Moe, Christiansen, Andersson, Fagerholt, and Norstad (2011) and Stålhane, Rakke, Moe, Andersson, Christiansen, and Fagerholt (2012). All voyages are direct sailings and only one loading port and one unloading port are visited on each voyage, while Grønhaug, Christiansen, Desaulniers, and Desrosiers (2010) present a problem where the product is picked up at one out of several loading ports before unloading at one or multiple unloading ports. Even though the ship may unload at several ports in one voyage it is always completely emptied before starting a new loading operation. In Christiansen (1999) and Song and Furman (2010) the structure is even more complex; the ships may pick up and deliver product at multiple ports, i.e. the ships may still have product onboard when entering a loading port.

In the basic MIRP there are inventory limits both at the production and consumption ports which we see in Al-Khayyal and Hwang (2007), while in Rakke et al. (2011) inventory constraints are only considered at the production port.

Production and consumption can be considered in several different ways, e.g. as a fixed rate as in Christiansen (1999) where the production rate is the same throughout the planning horizon. It can also be a varying fixed rate as in Rakke et al. (2011) and Stålhane et al. (2012). Here the production is considered as fixed over the planning horizon, but it still might vary from one day to the next. When this assumption is too coarse, one needs to resort to variable production and/or consumption rates. Ronen (2002) describes an inventory routing problem for refinery products. Here, variable rates for both production and consumption are considered to give the necessary information about the true nature of the

problem. The production is also a decision variable in Grønhaug et al. (2010).

Most classical MIRPs consider only one product, while lately MIRPs with multiple products are given more attention. Examples of the former can be found in e.g. Christiansen and Nygreen (1998a), Christiansen and Nygreen (1998b), Christiansen (1999) and Song and Furman (2010). Song and Furman (2010) present a MIRP for a single bulk product and a flexible modeling framework for MIRPs that can handle several practical features of such problems. The increased focus on MIRPs with multiple products comes from the challenges faced by i.e. chemical and petroleum industries as in e.g. Al-Khayyal and Hwang (2007). Other industries facing the challenge of shipping multiple products is e.g. cement industry as presented in Christiansen, Fagerholt, Flatberg, Haugen, Kloster, and Lund (2011) and the pulp industry as presented in Bredström et al. (2005). Inventory routing of multiple products is also discussed in e.g. Persson and Göthe-Lundgren (2005), Halvorsen-Weare and Fagerholt (2010), and Rakke et al. (2011).

1.1.2 Exact Solution Methods

Several exact solution methods are used in order to solve maritime routing and scheduling problems, and in this section a short description of the most commonly used methods will be presented. It is beyond the scope of this thesis to give an exhaustive description of all exact solution methods and variants, so only the basic concepts or methods will be given, together with some references to implementations. Solving most modern ship routing and scheduling problems involves solving either an *integer program* (IP) or a *mixed integer program* (MIP). Theoretically many of the problems presented in Section 1.1.1 can be solved using a standard optimization software. However, for most maritime routing and scheduling problems only small instances of the problems are solvable in this straight forward way due to the complexity.

Usually one can divide models for maritime routing and scheduling problems into two categories based on the decision variable for ship activity. Models using decision variables for each sailing leg are often called *arc-flow* models, while models using decision variables for ship routes are usually referred to as *path-flow* models.

For arc-flow models the most common exact solution method for solving hard problems is *branch-and-cut*. Here, classes of valid inequalities are left out of the LP relaxation because they contain too many constraints to be handled efficiently and the fact that most of these constraints will not be binding in the optimal solution. Instead, the constraints are dynamically added if they are violated by the optimal solution to the LP relaxation. These violated constraints are identified in a subproblem, often called a separation problem. After adding the violated con-

straints the problem is re-optimized, and this is repeated until no more violated constraints can be found. When no such violated constraints can be identified, branching is performed. This process is repeated until the optimal solution is found. Practical implementations of branch-and-cut in maritime routing and scheduling can be found in e.g. Song and Furman (2010), Rakke, Christiansen, Fagerholt, and Laporte (2012), and Andersson, Christiansen, Desaulniers, and Rakke (2012).

Many maritime routing and scheduling problems have groups of common constraints affecting multiple ships and groups of constraints only affecting a single ship. This structure may be exploited to separate the problem into one masterproblem containing the common constraints, and one subproblem for each ship. Such a decomposition is usually referred to as *Dantzig-Wolfe decomposition*. For each ship the subproblem has to generate feasible paths through the transportation network with respect to for instance capacity constraints, routing constraints, and time windows on the cargoes. These paths can be generated either a priori or dynamically throughout the branch-and-bound tree. Generation of such paths is usually called column generation, as one path corresponds to a column in the master problem. When the paths are generated a priori a set of pregenerated paths, or columns, is added to the LP relaxation and the overall problem is solved using branch-and-bound. This approach can be found in for instance Brown et al. (1987), Fisher and Rosenwein (1989), Fagerholt and Christiansen (2000a), Fagerholt and Christiansen (2000b), and Brønmo et al. (2007b). For many problems the number of feasible columns is too large to be generated a priori, and the columns have to be generated through dynamic column generation. Here, a *restricted masterproblem* (RMP) containing only a sufficient number of columns to provide a feasible solution to the LP relaxation is solved. This initial solution provides dual variables to the subproblems in order to obtain new columns. Only columns with a negative reduced cost (minimization) are returned from the subproblems. The RMP is re-optimized, and new duals are provided to the subproblems. When no more negative reduced cost columns can be found, branching is performed. This procedure is repeated until an optimal solution to the problem is found. Different algorithms for solving the subproblems in dynamic column generation are discussed in Desaulniers, Desrosiers, and Solomon (2005). Dynamic column generation in a maritime transportation setting is used in e.g. Christiansen (1999) and Brønmo et al. (2010).

1.1.3 Heuristics

As with exact solution methods, a large number of different heuristics are used to solve maritime routing and scheduling problems. This thesis will not give an exhaustive description of all heuristic methods and variants. Only the basic

concepts of the most commonly used methods will be given, together with some references to implementations. The heuristics are divided into two major categories, heuristics based on mathematical programming and heuristics based on algorithm design.

Heuristics based on mathematical programming usually starts with an exact formulation of the problem, but makes certain adjustments in order to be able to provide good solutions within reasonable time. Such modifications can for instance be decomposition and solution space reduction. Decomposition refers to breaking the overall problem into smaller, more manageable, problems and solving these sequentially. The solution of one problem is used as input to the next. One heuristic used in several maritime routing and scheduling problems is the rolling horizon heuristic (RHH) presented by Baker (1977) and Baker and Peterson (1979). The idea of the RHH is to repeatedly solve a sequence of MIPs, each covering a short time horizon. The overall problem is solved when all days in the planning horizon have been considered in at least one MIP. Often the MIP is extended by adding an extra period, called the forecasting period, where the binary/integer constraints are relaxed. This heuristic can be found in Bredström and Rönnqvist (2006) and Rakke et al. (2011). In both these papers an improvement heuristic is used on top of the RHH. After fixing all routes using a RHH, a LP determining the flow is applied in Bredström and Rönnqvist (2006), while a MIP where the solution from the RHH is used to reduce the solution space, is solved in Rakke et al. (2011). A similar improvement heuristic is also applied in Stålhane et al. (2012). Reducing the solution space usually involves removing variables, as in Rakke et al. (2011), or introducing inequalities that are invalid for the exact problem, but only has minor impact on the objective function.

Most algorithm based heuristics use some kind of construction heuristic to generate one or more initial solution(s). An improvement heuristic is then used to permute the initial solution(s) in a logical way to obtain a better solution. Construction heuristics build a solution by adding individual components one by one until a feasible solution is obtained. Such components may be arcs, vessels, nodes, cargoes, etc. In Stålhane et al. (2012) the construction heuristic chronologically adds scheduled voyages to the solution, while keeping the inventory within its limits. The scheduled voyages are picked according to the availability of the ships, the remaining demand of the contracts, and the inventory status. After generating a feasible solution improvement heuristics are usually applied. Examples of commonly used improvement heuristics are; local search, tabu search, and large neighborhood search (LNS). Local search tries to improve the solution by moving from one feasible solution to a neighboring solution while improving the objective value. When no such moves can be identified a local optimum is found, and the search terminates. The success of a local search is very dependent on the

initial solution, and the quality of the local optima around this initial solution. In order to deal with this nature, local search is often used with multiple initial solutions. Therefore, this algorithm is often called a multi-start local search. Maritime applications of multi-start local search can be found in e.g. Brønmo et al. (2007a) and Fagerholt, Korsvik, and Løkketangen (2009a).

A local search heuristic only explores a small area around the initial solution, and can easily be trapped in a local optimum. To escape local optima, several meta-heuristics are proposed, and perhaps the most commonly used in transportation and scheduling are tabu search and LNS.

The LNS works by destroying and repairing a part of the current solution in each iteration. It starts from an initial solution and performs a search until a local optimum is reached. The new solution is accepted if it is better than the previously best known solution, if not it is rejected and the algorithm returns to the best known solution. This solution is then destroyed and repaired to obtain a new starting solution. The procedure is performed for a predetermined time or number of iterations. Korsvik et al. (2011) use this heuristic to obtain very good solutions to a maritime-PDPTW with split loads.

In tabu search local optima are escaped by allowing non-improving moves if no improving moves are found. If the local optimum is better than the best known solution the best known solution is updated, and the tabu search chooses a non-improving move according to a predetermined strategy. A tabu-list is kept in order to avoid revisiting the same local optimum in the next iterations. The tabu search continues until a limit is reached. This limit can be e.g. number of moves, time or number of local optima visited. Several papers apply tabu search to maritime routing and scheduling problems, for instance Korsvik et al. (2009), Korsvik and Fagerholt (2010), and Fagerholt et al. (2011).

1.2 Purpose and Outline of Thesis

First, Section 1.2.1 discusses the purpose of the thesis. Then, Sections 1.2.2-1.2.6 present a brief summary of each paper. The papers presented in Sections 1.2.2-1.2.5 are all based on the same problem, and hence, the problem will only be presented in Section 1.2.2.

1.2.1 Purpose of Thesis

The purpose of this thesis is to develop new optimization models and methods for the shipping industry. In particular it focuses on models and methods for the transportation of LNG and other petroleum products. These models and methods can be used to develop decision support systems (DSSs) that can help ship operators dealing with routing and scheduling problems. The models and

methods presented in the papers in Sections 1.2.2–1.2.5 are developed to solve a real-world problem, but they are generic and may also be applied to other problems with the same characteristics. Section 1.2.4 is an invited book chapter, but will be referred to as a paper as the book is a collection of papers.

The thesis contains five papers. The first four papers discuss a problem from the LNG industry, but present different methods to solve it. The papers in Sections 1.2.2 and 1.2.3 present different heuristics in order to produce good solutions, while the papers in Sections 1.2.4 and 1.2.5 are more theoretical. In these two papers we focus more on providing stronger formulations of the problem in order to strengthen the bound. The paper in Section 1.2.6 is based on a problem that might appear as part of a subproblem in an inventory routing problem, and the inspiration is from a real-world petroleum transportation problem.

In the process of writing these 5 papers several different co-authors and referees have influenced the writing style, notations, and model syntaxes. This has resulted in some differences. In the first four papers the mathematical models and problem descriptions are quite similar, and may seem repetitive, but this was done in order to make each paper possible to read by itself. For the purpose of this thesis the paper formats, fonts, and reference styles have been standardized, leading to some differences to the published papers.

1.2.2 Paper 1: A Rolling Horizon Heuristic for Creating a Liquefied Natural Gas Annual Delivery Program

In this paper a maritime inventory routing problem for one of the world’s largest producers of liquefied natural gas (LNG) is presented. The producer is responsible for the LNG inventories at the liquefaction plant, a loading port with a limited number of berths, and routing and scheduling of a heterogeneous fleet of LNG ships. Production may vary during the planning horizon, but these variations are assumed to be known. The limited number of berths restrict the number of ships that can be loaded in a given day. In addition, the producer has to fulfill a set of long-term contracts to customers all around the world.

The goal is to create an annual delivery program (ADP) to fulfill the long-term contracts at minimum cost, while maximizing revenue from selling LNG in the spot market. An ADP is a complete schedule of every ship’s sailing plan for the coming year.

A mixed integer programming (MIP) formulation of the ADP planning problem is presented. The formulation is based on a priori generation of all possible scheduled voyages within the planning horizon. Due to the size and complexity of the problem, there was still a substantial gap even after running the MIP for 24h. In order to produce better integer solutions in less time a rolling horizon heuristic (RHH) is proposed. The RHH solves the problem by iteratively solving

subproblems with shorter planning horizons. Each subproblem is a MIP, and is solved using a commercial MIP solver. A solution space reduction method is also applied to limit the symmetry and make it easier to solve the subproblems. To further improve the initial solutions a MIP improvement heuristic is developed. The improvement heuristic takes the integer solution of the RHH as an input and creates a reduced MIP for the full problem. It allows swaps in the schedules and moves of the loading day within a restricted number of days from the loading day in the integer solution produced by the rolling horizon. The RHH produced better solutions to the problem than the MIP formulation in less time.

The main contributions in this paper is the introduction of a new problem to the literature together with a heuristic solution method that produces very good solutions in a limited amount of time. The paper was co-authored by Magnus Stålhane, Christian Rørholt Moe, Marielle Christiansen, Henrik Andersson, Kjetil Fagerholt, and Inge Norstad. It is published in *Transportation Research Part C* Vol. 19, 2011, pp. 896–911.

1.2.3 Paper 2: A Construction and Improvement Heuristic for a Liquefied Natural Gas Inventory Routing Problem

This paper discusses the same problem as the one in Section 1.2.2. The solution method used in the first paper produced very good solutions to the problem using less time than the MIP formulation, but it still used almost two hours for one of the instances. In this paper the focus was to create a faster heuristic that was easier to implement in a DSS than the RHH. A construction and improvement heuristic (CIH) was developed to accomplish this.

The CIH is a multi-start local search heuristic that constructs a set of solutions using a greedy insertion procedure. The solutions are then improved using either a first-descent neighborhood search, branch-and-bound on a mathematical formulation, or both. Tests on real-world instances show that the CIH provides good solutions in a short amount of time.

The main contributions in this paper is a heuristic solution method that produce very good solutions in a short amount of time. This makes it very suitable for inclusion in a DSS. The paper was co-authored by Magnus Stålhane, Christian Rørholt Moe, Henrik Andersson, Marielle Christiansen, and Kjetil Fagerholt. It is published in *Computers & Industrial Engineering* Vol. 62, No. 1, 2012, pp. 245–255.

1.2.4 Paper 3: A Branch-and-Cut Method for Creating a Liquefied Natural Gas Annual Delivery Program

A mixed integer programming (MIP) formulation of the ADP planning problem is presented and solved with a branch-and-cut algorithm. Several types of valid inequalities are developed that allow us to reduce the linear programming gap of the MIP formulation. The computational study shows that the problem is very complex, but that the valid inequalities are effective.

The main contribution of this paper is the new valid inequalities. These valid inequalities can be used for many other problems with similar characteristics. The article is co-authored by Henrik Andersson, Marielle Christiansen, and Guy Desaulniers. It is submitted as an invited chapter to a volume of Springer's International Series in Operations Research and Management Science tentatively titled *Optimization and Analytics in the Oil and Gas Industry* with K. Furman, J.-H. Song and A. El-Bakry as editors.

1.2.5 Paper 4: A Branch-Price-and-Cut Method for Creating a Liquefied Natural Gas Annual Delivery Program

This is the forth and final paper discussing the ADP planning problem described in Section 1.2.2. In this paper a new branch-price-and-cut method for the problem is presented. Two MIP formulations are presented and compared. One of these formulations is identical to the one presented in the three previous papers, while the other is new. In the new formulation a new decomposition approach based on delivery patterns at the customers is developed. This formulation is solved using a branch-price-and-cut method and the results are compared to the ones from the standard MIP. The results show that the new formulation drastically improves the initial lower bounds of the problem. In addition, the new formulation closes more than 56% of the remaining gap between the best known solution and the lower bound found by the standard MIP for the instances not solved to optimality by the MIP.

In this paper the main contribution is the new decomposition approach. It differs from other decomposition approaches presented in the literature for transportation networks by focusing on the pattern of deliveries in the demand ports rather than the sailing patterns of the ships. This decomposition can be an alternative for other problems where the commonly used decomposition methods do not work well. The article is co-authored by Henrik Andersson, Marielle Christiansen, and Guy Desaulniers. It is submitted to an international journal.

1.2.6 Paper 5: The Traveling Salesman Problem with Draft Limits

In maritime transportation, routing decisions are sometimes affected by draft limits in ports. The draft of a ship is the distance between the waterline and the bottom of the ship and is a function of the load onboard. Draft limits in ports can thus prevent ships from entering these ports fully loaded and may impose a constraint on the sequence of visits made by a ship. This paper introduces the *Traveling Salesman Problem with Draft Limits* (TSPDL), which is to determine an optimal sequence of port visits under draft limit constraints. We present two mathematical formulations for the TSPDL, and suggest valid inequalities and strengthened bounds. We also introduce a set of instances based on TSPLIB. A branch-and-cut algorithm is applied on both formulations for all these instances. Computational results show that introducing draft limits make the problem much harder to solve. They also indicate that the proposed valid inequalities and strengthened bounds significantly reduce both the number of branch-and-bound nodes and the solution times.

The main contributions in this paper are the introduction of a new problem class to the research community, two new mathematical formulations for this problem, and a set of new valid inequalities. The inspiration to the problem comes from a petroleum distribution problem in shipping, and the TSPDL may be found as a part of one of the subproblems. As shown in the paper, the TSP becomes much harder to solve when also considering draft limits, and we hope that the research made on this problem class may contribute to solve larger and more complex problems, including such subproblems. The paper also shows that the inclusion of subtour elimination constraints can efficiently be combined with two well known formulations of the TSP, and drastically improve the bound and performance of these models.

This article is co-authored by Marielle Christiansen, Kjetil Fagerholt, and Gilbert Laporte. It is published in *Computers & Operations Research* Vol. 39, No. 9, 2012, pp. 2161–2167.

1.3 Contributions

This thesis have contributed to both the research community and the industry. In this section these contributions are summarized together with an evaluation of my contributions to the papers. The rest of the section are organized in the following way. First, the contributions from this thesis to the research community are discussed, followed by a short description of the contributions this thesis has made to the industry. Then, the contributions I made to the papers in this thesis are discussed.

1.3.1 Contributions to the Research Community

Details of the contributions to the research community are presented in Sections 1.2.2–1.2.6. The main contributions are the introduction of a new maritime variant of the traveling salesman problem, introduction of a new type of inventory routing problem from the LNG industry, and the introduction of a new decomposition approach for inventory routing problems. New valid inequalities for both the MIRP and TSPDL have been presented, together with several mathematical models for these problems. Several previously presented valid inequalities are also strengthened. In addition to the publication of the papers, I have also presented the research at several scientific conferences around the world.

1.3.2 Contributions to the Industry

All four papers described in Sections 1.2.2–1.2.5 study an inventory routing problem based on a case study from the industry. Even though the research in these papers is focused on one problem, the solution methods, models and results are general, and thus applicable to a wide range of inventory routing problems. The heuristics in Section 1.2.2 and 1.2.3 together with some of the results from the papers in Section 1.2.4 and 1.2.5 can serve as a foundation for industry actors planning to develop a DSS. The models and solution methods in the papers in Section 1.2.4 and 1.2.5 can be used to provide valuable information of the quality of the solutions from a heuristic DSS. In addition, the research described has provided several valuable inputs to the DSS TurboRouter (MARINTEK).

1.3.3 Contributions to the Papers

Table 1.1 lists my level of contribution to the five papers in this thesis. The contributions are differentiated between *Intellectual Input*, *Implementation*, and *Writing*. Intellectual Input refers to the work related to identifying the problem, formulating the mathematical formulations and valid inequalities, and choosing and developing appropriate solution algorithms. Implementation covers the data handling, implementing the formulation to a runnable computer program, and the execution of the tests. It also covers the analyzes of the test results. Finally, the Writing refers to the actual writing of the paper. This also includes the creation of figures and tables, and the submission and referee process. The levels of contribution are ranked from 1 to 3, where 3 is highest. 1 means some contribution, 2 significant contribution, and 3 major contribution. The problem discussed in papers 1–4 was provided by Inge Norstad, MARINTEK. Christian Rørholt Moe, Magnus Stålhane, and I were the main contributors to the idea for the work done in papers 1 and 2. We worked together to identify the problem, formulate the mathematical formulation, and develop the heuristics. During this

Paper	Intellectual input	Implementation	Writing
1	3	3	3
2	3	2	1
3	3	3	1
4	3	3	3
5	3	3	3

Table 1.1: My contributions to the papers in thesis

process we got invaluable input from Henrik Andersson, Marielle Christiansen, and Kjetil Fagerholt. I was solely responsible for writing the first draft of Paper 1 with some input from Magnus Stålhane, but received very valuable help from the three co-authors mentioned to finalize and revise the paper. Paper 2 was written by Magnus Stålhane with some inputs from me, and valuable help from the other co-authors. The research presented in papers 3 and 4 was mainly performed during my research stay in Montreal, Canada. Here I worked in close co-operation with Professor Guy Desaulniers. I did most of the problem identification and model formulation for Paper 3, with a significant contribution from Guy Desaulniers and some input from the other co-authors. In addition, I did the full implementation. The paper was mostly written by Henrik Andersson and Marielle Christiansen. The model formulation and development of the solution method for Paper 4 were conducted in co-operation with Guy Desaulniers, who also provided invaluable help during the testing of the implementation. I was solely responsible for the implementation, data handling and testing, but received some help from Henrik Andersson. The writing of the draft for Paper 4 was done with input from Guy Desaulniers, Henrik Andersson, and Marielle Christiansen. Problem identification for Paper 5 was done as a shared work among the co-authors, while the initial mathematical formulations of the models was my responsibility with help from the co-authors. I did all the implementation and data handling for this paper. The draft of the paper was also written by me, with major contributions from Gilbert Laporte when it came to reviewing the paper. The two other co-authors also had significant inputs during the reviewing process.

1.4 Concluding Remarks

A substantial part of the worlds cargo transportation is sea-based. In spite of this, the amount of research within operations research on maritime transportation problems are relatively scarce compared to land-based transportation problems. Looking at inventory routing problems for maritime transportation there exists

even less literature. In the OPTIMAR project the goal was to provide efficient solution methods to challenging optimization problems, and many of the most challenging problems combine inventory management and routing. Hence, this thesis mainly focused on inventory routing problems or problems that are a part of such problems.

In papers 1 and 2 an exact solution method and two heuristics for a large scale inventory routing problem were developed. The heuristics provided good solutions to the problem in reasonable time. The exact method gives feasible solutions and lower bounds, but these bounds are not sufficiently strong for the method to be effective. Strong lower bounds are essential when determining the quality of the heuristic solutions. In papers 3 and 4 the aims were to develop cuts and model reformulations that could improve these bounds. The cuts improved the lower bounds for most instances tested. Combining these cuts with a new decomposition approach further increased the lower bounds and this approach solved almost all instances tested to optimality. Paper 5 introduced a new maritime variant of the traveling salesman problem (TSP); the TSP with draft limits (TSPDL). This problem is a common problem in maritime transportation, also in inventory routing problems. Two different models for the TSPDL were proposed, together with several new valid inequalities and methods for tightening bounds. The computational results showed that the branch-and-cut method developed, produced very good results for the instances tested. Combining previously known models and cuts in a new way was also shown to give significant improvements to both the TSP and TSPDL instances tested.

In recent years the short term market for LNG has had a strong increase. This creates several challenges for ship operators. In papers 1–4 the short term market (spot market) has not been modeled in detail. As the maturity and market share of this market increase, a more detailed modeling may be needed and this would be an interesting area of future research. The total amount of natural gas transported as LNG is also increasing, and this may impact the availability of ships for charter-in. In the models presented in this thesis one major assumption was that there would always be ships available if the producer was unable to lift the required amount of LNG with its own fleet. If the LNG market grows according to some of the predictions this might not be the case and further research on how to model charter-ins may be needed. Stochasticity in maritime inventory routing problems has not been discussed in this thesis as this was considered outside the scope of both the OPTIMAR project and the thesis. There are many interesting problems to investigate in this area. In addition to traditional uncertainties such as weather, breakdowns, and fuel prices, there are significant uncertainties associated with the spot market and the chartering of ships. Understanding the mechanisms and being able to model these uncertainties could prove to be of great value for ship operators.

Bibliography

- Al-Khayyal, F., Hwang, S., 2007. Inventory constrained maritime routing and scheduling for multi-commodity liquid bulk, Part I: Applications and model. *European Journal of Operational Research* 176 (1), 106–130.
- Andersson, H., Christiansen, M., Fagerholt, K., 2011a. The Maritime Pickup and Delivery Problem with Time Windows and Split Loads. *INFOR: Information Systems and Operational Research* 49 (2), 79–91.
- Andersson, H., Duesund, J. M., Fagerholt, K., 2011b. Ship routing and scheduling with cargo coupling and synchronization constraints. *Computers & Industrial Engineering* 61 (4), 1107–1116.
- Andersson, H., Christiansen, M., Desaulniers, G., Rakke J. G., 2012. Creating Annual Delivery Programs of Liquefied Natural Gas. *Working paper*, Norwegian University of Science and Technology, 19 pages.
- Appelgren, L., 1969. A column generation algorithm for a ship scheduling problem. *Transportation Science* 3 (1), 53–68.
- Appelgren, L., 1971. Integer programming methods for a vessel scheduling problem. *Transportation Science* 5 (1), 64–78.
- Baker, K., Peterson, D., 1979. An analytic framework for evaluating rolling schedules. *Management Science* 25 (4), 341–351.
- Baker, K. R., 1977. An experimental study of the effectiveness of rolling schedules in production planning. *Decision Sciences* 8 (1), 19–27.
- Bausch, D., Brown, G., Ronen, D., 1998. Scheduling short-term marine transport of bulk products. *Maritime Policy & Management* 25 (4), 335–348.
- Bredström, D., Carlsson, D., Rönnqvist, M., 2005. A Hybrid Algorithm for Distribution Problems. *IEEE Intelligent Systems* 20 (4), 19–25.
- Bredström, D., Rönnqvist, M., 2006. Supply chain optimization in pulp distribution using a rolling horizon solution approach. Discussion paper 2006:17, Norwegian School of Economics and Business Administration, Department of Finance and Management Science, 18 pages.

- Brønmo, G., Christiansen, M., Fagerholt, K., Nygreen, B., 2007a. A multi-start local search heuristic for ship scheduling - a computational study. *Computers & Operations Research* 34 (3), 900–917.
- Brønmo, G., Christiansen, M., Nygreen, B., 2007b. Ship routing and scheduling with flexible cargo sizes. *Journal of the Operational Research Society* 58 (9), 1167–1177.
- Brønmo, G., Nygreen, B., Lysgaard, J., 2010. Column generation approaches to ship scheduling with flexible cargo sizes. *European Journal of Operational Research* 200 (1), 139–150.
- Brown, G., Graves, G., Ronen, D., 1987. Scheduling ocean transportation of crude oil. *Management Science* 33 (3), 335–346.
- Cheng, L., 2004. Logistics for world-wide crude oil transportation using discrete event simulation and optimal control. *Computers & Chemical Engineering* 28 (6-7), 897–911.
- Christiansen, M., 1999. Decomposition of a Combined Inventory and Time Constrained Ship Routing Problem. *Transportation Science* 33 (1), 3–16.
- Christiansen, M., Fagerholt, K., 2009. Maritime Inventory Routing Problems. In: Floudas, C. A., Pardalos, P. M. (Eds.), *Encyclopedia of Optimization*. Springer US, Boston, MA, Ch. 13, 1947–1955.
- Christiansen, M., Fagerholt, K., Flatberg, T., Haugen, Ø., Kloster, O., Lund, E. H., 2011. Maritime inventory routing with multiple products: A case study from the cement industry. *European Journal of Operational Research* 208 (1), 86–94.
- Christiansen, M., Fagerholt, K., Nygreen, B., Ronen, D., 2007. Maritime transportation. In: Barnhart, C. and Laporte, G. (Eds.), *Handbooks in operations research and management science*. Vol. 14. Elsevier, Ch. 4, 189–284.
- Christiansen, M., Nygreen, B., 1998a. A method for solving ship routing problems with inventory constraints. *Annals of Operations Research* 81, 357 – 378.
- Christiansen, M., Nygreen, B., 1998b. Modelling path flows for a combined ship routing and inventory management problem. *Annals of Operations Research* 82, 391–413.
- Desaulniers, G., Desrosiers, J., Solomon, M. M. (Eds.), 2005. *Column Generation*. Springer-Verlag, New York.

- Desrosiers, J., Dumas, Y., Solomon, M., Soumis, F., 1995. Time constrained routing and scheduling. In: Ball, M. O., Magnanti, T. L., Monma, C. L., Nemhauser, G. L. (Eds.), *Handbooks in operations research and management science*. Vol. 8. Elsevier, Ch. 2, 35–139.
- Dror, M., Ball, M., 1987. Inventory/routing: Reduction from an annual to a short-period problem. *Naval Research Logistics* 34 (6), 891–905.
- Dumas, Y., Desrosiers, J., Soumis, F., 1991. The pickup and delivery problem with time windows. *European Journal of Operational Research* 54 (1), 7–22.
- Fagerholt, K., Christiansen, M., 2000a. A combined ship scheduling and allocation problem. *Journal of the Operational Research Society* 51 (7), 834–842.
- Fagerholt, K., Christiansen, M., 2000b. A travelling salesman problem with allocation, time window and precedence constraints - an application to ship scheduling. *International Transactions in Operational Research* 7 (3), 231–244.
- Fagerholt, K., Hvattum, L. M., Johnsen, T. V., Korsvik, J. E., 2011. Routing and scheduling in project shipping. *Annals of Operations Research*, 15 pages. DOI: 10.1007/s10479-011-0888-1
- Fagerholt, K., Korsvik, J. E., Løkketangen, A., 2009a. Ship Routing and Scheduling with Persistence and Distance Objectives. In: Nunen, J. A. E. E., Speranza, M. G., Bertazzi, L. (Eds.), *Innovations in Distribution Logistics. Lecture Notes in Economics and Mathematical Systems* Vol. 619. Springer, Berlin Heidelberg, Ch. 6, 89–107.
- Fagerholt, K., Laporte, G., Norstad, I., 2009b. Reducing fuel emissions by optimizing speed on shipping routes. *Journal of the Operational Research Society* 61 (3), 523–529.
- Fisher, M., Rosenwein, M., 1989. An interactive optimization system for bulk-cargo ship scheduling. *Naval Research Logistics* 36 (1), 27–42.
- Grønhaug, R., Christiansen, M., Desaulniers, G., Desrosiers, J., 2010. A Branch-and-Price Method for a Liquefied Natural Gas Inventory Routing Problem. *Transportation Science* 44 (3), 400–415.
- Halvorsen-Weare, E. E., Fagerholt, K., 2010. Routing and scheduling in a liquefied natural gas shipping problem with inventory and berth constraints. *Annals of Operations Research*. DOI: 10.1007/s10479-010-0794-y
- Korsvik, J. E., Fagerholt, K., 2010. A tabu search heuristic for ship routing and scheduling with flexible cargo quantities. *Journal of Heuristics* 16 (2), 117–137.

Bibliography

- Korsvik, J. E., Fagerholt, K., Laporte, G., 2009. A tabu search heuristic for ship routing and scheduling. *Journal of the Operational Research Society* 61 (4), 594–603.
- Korsvik, J. E., Fagerholt, K., Laporte, G., 2011. A large neighbourhood search heuristic for ship routing and scheduling with split loads. *Computers & Operations Research* 38 (2), 474–483.
- Lawrence, S., 1972. *International sea transport: the years ahead*. Lexington Books Lexington MA.
- Lejeune, M., Margot, F., 2008. Integer programming solution approach for inventory-production-distribution problems with direct shipments. *International Transactions in Operational Research* 15 (3), 259–281.
- OPTIMAR, 2005 *Optimization in Maritime Transportation and Logistics*. Project description, 11 pages.
- Persson, J. A., Göthe-Lundgren, M., 2005. Shipment planning at oil refineries using column generation and valid inequalities. *European Journal of Operational Research* 163 (3), 631–652.
- Rakke, J. G., Stålhane, M., Moe, C. R., Christiansen, M., Andersson, H., Fagerholt, K., Norstad, I., 2011. A rolling horizon heuristic for creating a liquefied natural gas annual delivery program. *Transportation Research Part C: Emerging Technologies* 19 (5), 896–911.
- Rakke, J. G., Christiansen, M., Fagerholt, K., Laporte, G., 2012. The Traveling Salesman Problem with Draft Limits. *Computers & Operations Research*, 39 (9), 2161–2167.
- Ronen, D., 1982. The effect of oil price on the optimal speed of ships. *Journal of the Operational Research Society* 33 (11), 1035–1040.
- Ronen, D., 2002. Marine inventory routing: shipments planning. *Journal of the Operational Research Society* 53 (1), 108–114.
- Scott, J., 1995. A transportation model, its development and application to a ship scheduling problem. *Asia Pacific Journal of Operational Research* 12, 111–128.
- Sherali, H. D., Al-Yakoob, S. M., Hassan, M. M., 1999. Fleet management models and algorithms for an oil-tanker routing and scheduling problem. *IIE Transactions* 31 (5), 395–406.

- Song, J.-H., Furman, K. C., 2010. A maritime inventory routing problem: Practical approach. *Computers & Operations Research*, 9 pages. DOI: 10.1016/j.cor.2010.10.031
- Stålhane, M., Rakke, J. G., Moe, C. R., Andersson, H., Christiansen, M., Fagerholt, K., 2012. A Construction and Improvement Heuristic for a Liquefied Natural Gas Inventory Routing Problem. *Computers & Industrial Engineering* 62 (1), 245–255.

Paper I

J. G. Rakke, M. Stålhane, C. R. Moe, M. Christiansen, H.
Andersson, K. Fagerholt and I. Norstad:

A Rolling Horizon Heuristic for Creating a Liquefied Natural Gas Annual Delivery Program

Transportation Research Part C: Emerging Technologies 19 (5),
896–911.

A Rolling Horizon Heuristic for Creating a Liquefied Natural Gas Annual Delivery Program

Abstract:

In this paper a maritime inventory routing problem for one of the world's largest producers of liquefied natural gas (LNG) is presented. The producer is responsible for the LNG inventories at the liquefaction plant, the loading port with a limited number of berths, and the routing and scheduling of a heterogeneous fleet of LNG ships. In addition, the producer has to fulfill a set of long-term contracts to customers all around the world.

The goal is to create an annual delivery program (ADP) to fulfill the long-term contracts at minimum cost, while maximizing revenue from selling LNG in the spot market. An ADP is a complete schedule of every ship's sailing plan for the coming year.

A mixed integer programming (MIP) formulation of the ADP planning problem is presented, and it is based on a priori generation of all possible scheduled voyages within the planning horizon. Due to the size and complexity of the problem, a rolling horizon heuristic (RHH) is proposed. The RHH solves the problem by iteratively solving sub-problems with shorter planning horizons. The RHH finds solutions to the problem within a reasonable amount of time, and creates very good ADPs according to the problem owner.

2.1 Introduction

Today the demand for energy is higher than ever, and it is forecast that it will increase even more in the near future. Natural gas (NG) is expected to be one of the most important sources of energy to fulfill this demand. Estimates predict that the world's demand for NG will increase with 52% from 2005 to 2030 according to EIA (2008). At the same time, traditional sources of energy such as oil and coal are becoming less attractive due to increased prices and environmental considerations. This, and new technology making extraction, refinement, and transportation of NG both cheaper and more efficient, has made it more attractive to invest in projects involving NG. As a result gas-rich areas which previously were considered too remote to be profitable are now being explored.

Traditionally, NG was transported through pipelines, but long distances from the gas fields to many high-consumption markets made it prohibitively expensive

to build pipelines. Other forms of transport were also too costly, e.g. the option of transporting the gas as liquefied natural gas (LNG) was undesirable due to the need for expensive specialized ships and equipment to handle the LNG at the loading and unloading ports. However, a combination of higher LNG prices, technical developments, lower production costs, rising demand, and the desire of producers to capitalize on their gas reserves, has set the stage for increased LNG trade and more use of ships as a mean to transport NG. As a result of this, IEA (2007) predicts that the total fleet of ships transporting LNG will increase to approximately 400 by year 2015, close to double of the total fleet in 2007.

In addition to the rapidly increasing demand for NG a spot market for LNG has arisen in recent years. Traditionally, long-term contracts for delivery have been negotiated in advance of starting projects with take-or-pay clauses shifting the volume risk to the buyer. Few LNG facilities were built until sales contracts were signed for the entire capacity. Contracts also contained "destination clauses" that prevented buyers from reselling the cargoes to third parties. Recently, some projects have gone forward with capacity unclaimed and long-term contracts have been growing increasingly flexible. Newer long-term contracts are designed to provide only a base supply of LNG, and can be supplemented by short-term contracts during periods of high demand. This spare capacity and more flexible contracts are expected to lead to increased short-term sales. Experts project that short-term trading will continue to grow, especially in the Atlantic Basin, and could reach 15–20% of the LNG market over the next decade (EIA, 2007).

With even larger fleets and more complex decisions to make, the rapid growth in the LNG trade and the increasingly complex market of NG have made the work much harder for the schedulers to accomplish manually. This has led to a growing need for decision support systems (DSS) in order to help schedulers make good decisions quickly.

In this paper, a combined LNG ship routing and inventory management problem for one of the world's largest producers and distributors of LNG is considered. The producer has a single liquefaction plant with storage tanks and a connected loading port. Two types of LNG are produced at the liquefaction plant. The amount of LNG in the storage tanks at the liquefaction plant has to stay within certain limits, and the size of the loading port restricts the number of ships that can load simultaneously. From the loading port, the LNG is shipped to customers world-wide, see Figure 2.1.

The producer has to fulfill a set of long-term customer contracts that either outline monthly demands, or state that a certain volume of LNG is to be delivered fairly evenly spread throughout the year to a given regasification terminal. Over- and under-deliveries are accepted within reasonable limits. In addition, the producer has the opportunity to sell LNG in the spot market using short-term contracts. Besides extraction and refinement of NG, the producer operates



Figure 2.1: Map showing some of the main sailing lanes.

a large fleet of heterogeneous LNG ships. In peak periods it is possible to charter in additional ships.

To plan for the coming year, an annual delivery program (ADP) which is a list of scheduled voyages is created. Each scheduled voyage includes information about the ship sailing, the day of loading at the loading port, and the contract served. The objective is to create an ADP that abides by the long-term contractual agreements at lowest possible cost, while maximizing the expected revenue from spot contracts. Henceforth, this will be referred to as the ADP planning problem.

The purpose of this paper is twofold: (1) Presenting the LNG supply chain and a case study from one of the world's largest producers and distributors of LNG, where the objective is to create an ADP. (2) Describing a rolling horizon heuristic (RHH) for the ADP planning problem that solves instances in reasonable time and with good solution quality. In addition, a modified version of the underlying mathematical programming formulation is used to further improve a feasible ADP created by the RHH. The heuristic is compared with an exact solution method and a heuristic by evaluating the solution quality and solution time. The RHH is designed for use in a DSS for the producer in order to facilitate the creation of ADPs. With a fleet of 46 ships, 17 long-term contracts and one year planning horizon the need for a DSS is great, and the possible gain is substantial.

The rest of the paper is organized as follows: Section 2.2 describes the LNG supply chain and the typical planning problems at various levels within the LNG business. An overview of the recent literature within maritime inventory routing and use of rolling horizon heuristics is presented in Section 2.3. Section 2.4 is devoted to a detailed description of the planning problem, while it is formulated as a mixed integer programming problem in Section 2.5. The RHH is described in Section 2.6. Real-world cases are described in Section 2.7 and computational results for these cases are also reported. Finally, concluding remarks and suggestions for future research follow in Section 2.8.

2.2 The LNG Supply Chain and the Planning Levels within the LNG Business

In order to get a better understanding of the LNG business and the considered problem of the producer, a presentation of the LNG supply chain is given in Section 2.2.1 and the various planning levels are described in Section 2.2.2.

2.2.1 The LNG Supply Chain

The LNG supply chain, illustrated in Figure 2, begins with the NG being extracted from underground reservoirs and sent through pipelines to a liquefaction facility. At the liquefaction plant, impurities are removed from the gas, and the gas goes through three cooling processes where it reaches its boiling temperature of approximately -160°C . A liquefaction plant may consist of several parallel units ("trains"). By liquefying the gas, its volume is reduced by a factor of 600. In the supply chain considered, the NG is converted into two types of LNG, rich (RLNG) and lean (LLNG).




Figure 2.2: The LNG supply chain. The elements considered in this paper are highlighted.

The liquefaction plant constitutes a major cost element in the LNG supply chain. LNG plant costs are high primarily due to strict design and safety standards, considerable quantity of cryogenic materials required, and a historical

2.2 The LNG Supply Chain and the Planning Levels within the LNG Business

inclination to over-design in order to guarantee supply. In recent years, major economies of scale have been achieved by increasing the size of liquefaction trains, requiring fewer trains to achieve the same output (Chidinma, 2004).

After the refinement and liquefaction processes, the LNG is either temporarily stored in specially constructed tanks, or directly loaded onto ships where it is kept at its boiling temperature for the duration of the voyage to the regasification terminals. The ships can be owned by the producer, one or many customers or a distributor. If not owned by the producer, the ships may be committed to only serving specific contracts.

Upon a ship's arrival at the destination port, the LNG is fed into a regasification facility. There it is pumped into a storage tank where it is stored until needed. At that time, the gas is warmed in a controlled environment, and the revaporized natural gas then enters the pipeline system as methane used by consumers, power plants, and industrial customers.

2.2.2 The LNG Planning Levels

It is common to distinguish between three planning levels with different time horizons when planning the LNG supply chain (Stremersch, Michalek, and Hecq, 2008). The strategic, tactical, and operational planning levels will be discussed in the following, but focus is on the tactical level for creating an ADP.

Strategic Planning

Strategic decisions are long-term decisions important for both tactical and operational planning. In the LNG business, strategic decisions cover a wide spectrum and have an impact many years ahead. Investment decisions include physical assets associated with the liquefaction plant and the regasification terminals, where key issues are production, storage and berth capacities. In addition, the fleet composition is an important long-term decision for an LNG transporter. Up to now, most of the sales of LNG are tied up in long-term contracts which may have a duration of 20–30 years. The strategic decision concerning the contract portfolio is closely related to the long-term decisions associated with the physical assets.

Tactical Planning

One of the main tasks at the tactical level is to create a new ADP. The planning horizon is typical one year, but could be extended up to 18 months. When creating the ADP, the aim is to determine an optimal fleet schedule, including the delivery dates at the different customers' terminals. This fleet schedule must also satisfy inventory considerations, as well as contractual requirements.

In practice, the process of creating an ADP starts by creating an initial ADP in-office. This initial ADP is used when the producer negotiates with its cus-

tomers on delivery dates of LNG for the next year. As some customers may not accept the delivery dates suggested in the initial ADP, it is reconstructed with certain delivery dates changed. This process continues until all customers get a sufficient supply of LNG at acceptable delivery dates. The problem discussed in this paper is how to create the first in-office ADP, and the ADP for this producer is a table with the following information for each ship and voyage: the type of LNG transported, the contract served, and the dates for loading, unloading and returning to the loading port.

Operational Planning

The operational planning level deals with updating fleet schedules due to various logistical, economic, or contractual reasons. Examples of logistical reasons can be rescheduling due to unplanned events, such as equipment breakdown or ship delays. An example of an economic reason is when spot market prices change; new sales or purchase opportunities may create needs for rescheduling. The typical length of the operational planning horizon is 3 months.

2.3 Related Literature

In this section, the presentation is limited to consider literature related to the problem and solution approach considered. Hence, maritime inventory routing problems are described in Section 2.3.1, while rolling horizon heuristics are reviewed in Section 2.3.2.

2.3.1 Maritime Inventory Routing Literature

Combined routing and inventory management within maritime transportation have only been considered in the literature during the last 15 years. Christiansen and Fagerholt (2009) give an overview of maritime inventory routing problems, while Christiansen, Fagerholt, Nygreen, and Ronen (2007) present a comprehensive review within maritime transportation in general. Within LNG inventory routing problems (LNG-IRP), the literature is scarce. The most general study of the LNG supply chain, including some of its main characteristics, is presented in Andersson, Christiansen, and Fagerholt (2010). The authors consider two problems, one for a producer of LNG and one for a vertically integrated company. Mathematical models are presented for each problem and solution methods to both models are discussed, but no computational results are given. The main differences are that the problem studied in this paper considers more than one LNG product and takes into account the possibility of selling LNG in the spot market. For the vertically integrated company Andersson et al. (2010) consider

inventory management at the regasification terminals, which is not considered in this paper.

The first study reported in the literature on optimization of LNG-IRPs is done by Grønhaug, Christiansen, Desaulniers, and Desrosiers (2010) and Grønhaug and Christiansen (2009). The problem has a more complex structure than the problem considered here, as it involves more than one loading port, inventory constraints both at the loading and delivery ports, and each ship can visit more than one regasification terminal on a given voyage. However, the problem has fewer ships and a much shorter planning horizon than the problem presented in this paper. Grønhaug and Christiansen (2009) give both an arc flow and a path flow formulation for the problem. The paths are pre-generated, while Grønhaug et al. (2010) use a branch-and-price-and-cut algorithm to solve the problem. The authors also introduce the LNG-IRP as a problem class in the literature, a problem class which also includes the problem studied in this paper.

The ADP planning problem presented in this paper is also studied by Stålhane, Rakke, Moe, Andersson, Christiansen, and Fagerholt (2010), who developed a construction and improvement heuristic to create good ADPs. The heuristic solves the problem by iteratively constructing an initial ADP. This ADP is then improved using two different heuristics, one local search and one based on a reduced MIP formulation. Halvorsen-Weare and Fagerholt (2009) study a simplified version of the problem where cargoes for each long-term contract are pre-generated with defined time windows, and the fleet of ships can be divided into disjoint groups. The problem is decomposed into a routing sub problem and a scheduling master problem where berth, inventory and scheduling decisions are handled in the master problem, while routing decisions are dealt with in the sub-problem. Unlike branch-and-price, the sub-problems and master problem are solved only once.

The LNG-IRP is a special case of the inventory ship routing problem (ISRP), defined by Christiansen et al. (2007) and is a maritime adaption of the well known inventory routing problem (IRP). The first study of an ISRP was published by Christiansen and Nygreen (1998a,b), and Christiansen (1999). The authors consider a supply chain for ammonia consisting of several locations that either produce or consume ammonia and the transportation network between these locations. The problem has inventory constraints at both the production and consumption locations and the goal is to minimize the transportation cost. The overall problem is solved by a branch-and-price method in Christiansen (1999) and by a heuristic in (Flatberg, Haavardtun, Kloster, and Løkketangen, 2000). Al-Khayyal and Hwang (2007) present a MIP formulation for a multi-commodity liquid bulk ISRP. The problem considered is to find a minimum cost routing solution for a heterogeneous fleet of ships engaged in pick-up and delivery of several liquid bulk products, while also deciding the volumes of each

product carried. Small problems are solved to optimality, but due to the exponential growth in computational time as the number of port visits increases, the solution method seems to be unsuitable for solving large problem instances. The maritime inventory routing problem described in (Ronen, 2002) also includes multiple products. The underlying model focuses on the inventory management and not the routing part of the problem, as the model solution suggests shipment sizes that are assumed to be input for a cargo routing problem at a later stage.

2.3.2 Literature on Rolling Horizon Heuristics

Most contributions describing rolling horizon heuristics in the literature are found within manufacturing scheduling. Here, the practice of rolling horizons or rolling schedules is to routinely update or revise schedules, taking into consideration more reliable and recent data as they become available. Baker (1977) and Baker and Peterson (1979) were among the first to test the effectiveness of schedules obtained from a rolling horizon planning model using limited information about future demand. The authors considered the effects of the length of the horizons, the cost structure, and the demand patterns. The conclusion was that rolling schedules produced low-cost results, but their efficiency depended on the length of the forecasting period.

There exist several algorithms for solving rolling horizon problems that have proved to be very powerful, see for instance (Stauffer and Liebling, 1997), (Mercé and Fontan, 2003), (Dimitriadis, Shah, and Pantelides, 1997) and (Araujo, Arenales, and Clark, 2007). These algorithms were constructed to be applied in particular manufacturing settings and are not directly applicable to ship scheduling problems.

The only contribution we know of that applies a rolling horizon approach in a maritime setting is (Bredström and Rönnqvist, 2006). Here, a combined supply chain and ship routing problem for a large Scandinavian producer of pulp is considered. It is an operative planning problem with daily ship routing decisions over a 40 day planning horizon. The authors develop a MIP model with binary variables for combinations of ships and routes. The model is solved with a heuristic solution method based on a rolling horizon. The idea is to repeatedly solve a MIP, each covering a short time horizon. The overall problem is solved when all days in the planning horizon have been considered in at least one MIP. First an initial set of days are selected and then an extension serving as a forecasting period is added. In the forecasting period, the binary restrictions on the variables are relaxed. The approach is applied to real-world problem instances, and a comparison with an exact branch-and-bound on reduced problem instances is provided. The computational studies indicate that real-world instances are solvable with the proposed solution method and that it is very efficient in many

cases. The ship scheduling part of the problem shares some similarities with the ADP planning problem, but is considerably smaller. The fleet consists of three ships, and the planning horizon is shorter. However, each scheduled voyage is more complex, as it includes several port calls and the cargoes might be partial shiploads.

2.4 Problem Description

The problem studied in this paper is a combined ship routing and inventory management problem for a producer of LNG. It considers inventory management at the producer, ship routing and scheduling, and contract management between the producer and its customers.

2.4.1 Storage and Liquefaction

Two types of LNG are being produced at the producer's liquefaction facility, RLNG and LLNG. Production rates at the facility are always such that the equipment available for gas refinement is 100% utilized. Fluctuations in production rates occur due to planned maintenance or unforeseen events, meaning that production rates may vary during a year. The different types of LNG are stored in separate tanks with given upper and lower limits, and a limited number of berths for loading LLNG and RLNG are available.

2.4.2 Shipping

The producer operates a heterogeneous fleet of ships. These ships are either owned by the producer or by one, or a group of customers. However, all ships are included in the fleet scheduled by the producer.

Several factors influence the availability of the ships in the fleet. Since there is no depot, some of the ships may become available for operation after the start of the planning horizon, when they return to the loading port. The ships may also be unavailable due to certain pre-allocated activities, e.g. maintenance. The maintenance is not fixed to certain dates, but is required to be carried out within a given time interval. It is always scheduled to a dry-dock along the sailing lane, thus minimizing the time the ship is out of commission. After the maintenance has been performed, the ship has to go through a purge and cool-down procedure at a berth in the loading port before starting its next voyage. The length of this procedure takes on average 24 h.

A shipload must contain a single type of LNG, but a ship can carry different types of LNG on consecutive voyages without any intermediate preparations. Though it is technically possible for a ship to sail between the loading port and

a regasification terminal with only some of the tanks filled, this is not considered economically feasible by the producer and is never done in practice. Due to this, and in order to avoid sloshing during transportation, the ship tanks are always filled up to their capacities at the loading port. The duration of the voyage influences the amount of LNG delivered to the customer due to boil-off, which causes a fixed ship-dependent percentage (normally between 0.11% and 0.13%) of the total capacity of a tank to evaporate each day. To keep the LNG tanks on board a ship cool during the return leg, a certain amount of LNG must be kept in the tanks. If boil-off causes the ship tanks to be completely empty, a 24-h cool-down period is required at the loading port before the next loading operation can take place. However, the producer assumes that a ship will never be completely empty, meaning that ships will never have to go through a cool-down process, except after a maintenance period. In total, the operations executed between the arrival and departure at the loading port, excluding cool-down, takes approximately 24 h.

All ships have individual cruising speed, making the travel times dependent on the ship. The duration of a voyage may also depend on the time of year, as sailing conditions may vary between the summer and the winter months. Not all ships can visit all regasification terminals. This is due to vessel acceptance policies at the ports, and that some ships are limited to only visit their owners' regasification terminals.

The transportation costs consist of several components. The fixed costs are the time charter rates, while the variable costs are port and canal fees that are determined by the ship type and the contract served. A third variable component is the bunker cost, which is dependent on the ship size, the load on board and the duration of the voyage. Since the fixed costs cannot be changed during the time horizon, the cost of sailing a scheduled voyage is assumed to only be dependent on the capacity class of the ship, the duration of the voyage and the regasification terminal visited.

If the producer does not have enough ships available at a given time, additional ships may be chartered in for one-off deliveries. A daily charter rate defines the cost of these ships.

2.4.3 Long-term Contracts and Sale of Spot Cargoes

The producer has a set of long-term contracts to which it is obliged to deliver a certain amount of LNG to specified regasification terminals each year. These contracts have time-frames of 20–30 years. A long-term contract either outlines a monthly demand, or simply states that the LNG is to be delivered fairly evenly spread throughout the year. Due to the contracts being long-term, there is some flexibility in the delivered volumes in a given year.

In addition to serve the long-term contracts, the producer has the opportunity to sell LNG in the spot market. The spot market for LNG has until recently been limited. However, this is now changing and it is expected that the market will grow into a trade-market similar to that of oil in the near future. The spot sales are short-term contracts for one-off deliveries from the producer to buyers on agreed terms, and represent a possibility for increased revenue.

When including the option of selling LNG in the spot market, a new dimension is added to the ADP planning problem. First, the commercial framework shifts; in addition to fulfill the long-term contracts at minimum cost, the producer now seeks to maximize the marginal contribution from selling LNG in the spot market. Second, it becomes necessary to make assumptions about the expected spot price and the times at which spot contracts can be expected to occur in the market.

2.5 The Mathematical Model

In this section a MIP formulation for the ADP planning problem will be given. In the formulation a *contract* represents a combination of customer, regasification terminal, and the type of LNG delivered.

The whole planning horizon is divided into a number of *time partitions*. The time intervals in the time partitions cover the planning horizon and may be overlapping. If the planning horizon is a year, months and quarters could be time intervals, where months and quarters clearly overlap. The whole planning horizon is also considered a time interval. Demand in a given time interval is assumed to be known and is used as a parameter in the model independent of if a long-term contract states a monthly demand or if it states that the LNG should be delivered fairly evenly spread. The demand in each time interval need not to be met exactly, but there exists penalty costs associated with over- and under-deliveries.

A *scheduled voyage* serves a given contract using a specific ship starting from the loading port on a given day. For each feasible combination of ship, contract and start time, a scheduled voyage is generated. A combination is infeasible if the starting time is prior to the ship being available for operation, or if the vessel is not allowed to visit the regasification terminal associated with the contract. The sailing time of a scheduled voyage might vary during the year, and this is implemented by having the sailing time dependent on the starting day.

In contrast to vehicle routing where the trucks usually return to one or more central depots there are usually no central depots in ship routing. This means that ships may be anywhere at the start of a planning horizon. This model assumes that the starting positions of all ships are known. All ships are obligated to dry-dock for maintenance with an interval of a given number of years. After maintenance the ships are warm and completely empty and will need to go

through a purge and cool-down procedure at a berth in the loading port before starting the next voyage. Maintenance is modeled using one contract for each type of LNG, defining which berth is used for the purge and cool-down procedure.

Sometimes, the producer does not have sufficient capacity to lift all cargoes with its own fleet and needs to charter in ships. It is assumed that at most two ships can be chartered in each day, one for each type of LNG.

The production at the loading port may vary, but this is due to planned maintenance or unforeseen events. In this deterministic model unforeseen events are not handled, hence, the production rates are assumed to be known at the time of planning.

The storage tanks at the liquefaction plant have time dependent upper and lower limits, and the inventory levels must always be between these limits. There may be an initial inventory, and there may be LNG left in the tanks at the end of the planning horizon.

In the model, no revenue for long-term contracts are considered, only sailing costs and penalties. For spot cargoes there is a revenue of selling these on the spot market. In order to implement forecasts of the spot market price there is a time dependent revenue per m^3 based on which day the scheduled voyage starts. Since one year's ADP is dependent on the previous', there are introduced revenues on the end inventories in order to reduce end effects.

Sets

\mathcal{G}	The set of types of LNG, indexed by g .
\mathcal{C}	The set of contracts, indexed by c . $\mathcal{C} = \mathcal{C}^{LT} \cup \mathcal{C}^S \cup \mathcal{C}^M$.
\mathcal{C}^{LT}	The set of long-term contracts, $\mathcal{C}^{LT} \subset \mathcal{C}$.
\mathcal{C}^S	The set of spot contracts, $\mathcal{C}^S \subset \mathcal{C}$.
\mathcal{C}^M	The set of maintenance contracts, $\mathcal{C}^M \subset \mathcal{C}$.
\mathcal{C}_g	The set of contracts with demand for gas of type g , $\mathcal{C}_g \subset \mathcal{C}$.
\mathcal{C}_v	The set of contracts ship v may serve.
\mathcal{V}	The set of available ships, indexed by v . $\mathcal{V} = \mathcal{V}^P \cup \mathcal{V}^S$.
\mathcal{V}^P	The set of ships operated by the producer, $\mathcal{V}^P \subset \mathcal{V}$.
\mathcal{V}^S	The spot ships available for chartering, $\mathcal{V}^S \subset \mathcal{V}$.
\mathcal{V}^M	The set of ships that require maintenance during the planning horizon, $\mathcal{V}^M \subset \mathcal{V}^P$.
\mathcal{V}_c	The set of ships that may serve contract c , $\mathcal{V}_c \subset \mathcal{V}$.
\mathcal{T}	The set of days in the planning horizon, indexed by t . $\mathcal{T} = \{1, 2, 3, \dots, \bar{T}\}$, where \bar{T} is the last day in the planning horizon.
\mathcal{P}	The set containing elements of partitions of \mathcal{T} , indexed by p .
\mathcal{T}_p	The set of days in time interval p , $\mathcal{T}_p \subset \mathcal{T}$.
\mathcal{T}_v	The set of days where ship v is available to be scheduled, $\mathcal{T}_v \subset \mathcal{T}$.

\mathcal{T}_v^M The set of days where maintenance of ship v may start, $\mathcal{T}_v^M \subset \mathcal{T}$.

Constants

C_{cv}^T The transportation cost of serving contract c using ship v .
 C_{cp}^{D+} The penalty cost per unit of over-delivery of LNG to a contract c in time interval p .
 C_{cp}^{D-} The penalty cost per unit of under-delivery of LNG to a contract c in time interval p .
 R_{ct}^S The revenue of selling one unit of LNG under spot contract c at day t .
 R_g^I The per unit value of having LNG of type g left in the tank at day \bar{T} .
 L_v The loading capacity of ship v .
 T_{cvt} The time it takes to serve contract c using ship v when starting on day t .
 D_{cp} The demand of contract c in time interval p .
 P_{gt} Accumulated production of LNG of type g up to and including day t , plus the inventory level at the start of the planning horizon.
 B_g The number of berths available at the loading port for loading LNG of type g .
 \bar{I}_{gt} The maximum amount of LNG of type g that can be stored at the loading port at day t .
 \underline{I}_{gt} The minimum amount of LNG of type g that should be available at the loading port at day t .

Variables

x_{cvt} 1 if ship v starts a scheduled voyage for contract c on day t , 0 otherwise.
 y_{cp}^+ The over-delivery to long-term contract c in time interval p .
 y_{cp}^- The under-delivery to long-term contract c in time interval p .

$$\begin{aligned} \min z = & \sum_{c \in \mathcal{C}} \sum_{v \in \mathcal{V}_c} \sum_{t \in \mathcal{T}_v} C_{cv}^T x_{cvt} + \sum_{c \in \mathcal{C} \setminus \mathcal{C}^M} \sum_{p \in \mathcal{P}} C_{cp}^{D+} y_{cp}^+ + \sum_{c \in \mathcal{C} \setminus \mathcal{C}^M} \sum_{p \in \mathcal{P}} C_{cp}^{D-} y_{cp}^- - \\ & \sum_{c \in \mathcal{C}^S} \sum_{v \in \mathcal{V}_c} \sum_{t \in \mathcal{T}_v} R_{ct}^S L_v x_{cvt} - \sum_{g \in \mathcal{G}} R_g^I (P_{g\bar{T}} - \sum_{c \in \mathcal{C}_g \setminus \mathcal{C}^M} \sum_{v \in \mathcal{V}_c} \sum_{t \in \mathcal{T}_v} L_v x_{cvt}), \end{aligned} \quad (2.1)$$

$$\sum_{c \in \mathcal{C}_g \setminus \mathcal{C}^M} \sum_{v \in \mathcal{V}_c} x_{cvt} + \sum_{c \in \mathcal{C}_g \cap \mathcal{C}^M} \sum_{v \in \mathcal{V}^M} \sum_{\tau \in \mathcal{T} | \tau + T_{cv\tau} - 1 = t} x_{cv\tau} \leq B_g, \forall g \in \mathcal{G}, t \in \mathcal{T}, \quad (2.2)$$

$$\underline{I}_{gt} \leq P_{gt} - \sum_{c \in \mathcal{C}_g \setminus \mathcal{C}^M} \sum_{v \in \mathcal{V}_c} \sum_{\tau \in \mathcal{T}_v | \tau \leq t} L_v x_{cv\tau} \leq \bar{I}_{gt}, \forall g \in \mathcal{G}, t \in \mathcal{T}, \quad (2.3)$$

$$\sum_{v \in \mathcal{V}_c} \sum_{t \in \mathcal{T}_p} L_v x_{cvt} + y_{cp}^- - y_{cp}^+ = D_{cp}, \forall c \in \mathcal{C} \setminus \mathcal{C}^M, p \in \mathcal{P}, \quad (2.4)$$

$$\sum_{c \in \mathcal{C}} \sum_{\tau \in \mathcal{T} | t - T_{cv\tau} < \tau \leq t} x_{cv\tau} \leq 1, \forall v \in \mathcal{V}^P, t \in \mathcal{T}_v, \quad (2.5)$$

$$\sum_{c \in \mathcal{C}^M} \sum_{t \in \mathcal{T}_v^M} x_{cvt} = 1, \forall v \in \mathcal{V}^M, \quad (2.6)$$

$$x_{cvt} \in \{0, 1\}, \forall c \in \mathcal{C}, v \in \mathcal{V}_c, t \in \mathcal{T}_v, \quad (2.7)$$

$$y_{cp}^+, y_{cp}^- \geq 0, \forall c \in \mathcal{C} \setminus \mathcal{C}^M, p \in \mathcal{P}. \quad (2.8)$$

The objective function (2.1) minimizes the sum of the transportation and penalty costs minus the revenue from LNG sold in the spot market and the LNG left in the tanks at the end of the planning horizon. The berth restrictions are stated in constraints (2.2). Since a ship goes through a purge and cool-down procedure at one of the berths after maintenance, the variables corresponding to these scheduled voyages are added to the constraints on the return day of the voyages. Constraints (2.3) ensure that the inventory levels are always between their upper and lower bounds. The contractual demands are handled in constraints (2.4), relating the amount delivered in a time interval with the demand through the over- and under-delivery. The routing constraints (2.5) specify that a ship is only allowed to operate one scheduled voyage each day. Maintenance is handled by constraints (2.6), and constraints (2.7) and (2.8) are the variable restrictions.

There are some differences between this model and the model presented in (Stålthane et al., 2010). A cost of over-delivery to a contract has been added to the objective function (2.1) to make the model more general. In addition the revenue of selling LNG in the spot market and the sailing times are dependent on the starting day of a scheduled voyage. Inventory constraints (2.3) are also modeled as knapsack constraints to enable the solver used to generate more and stronger valid inequalities. Extensive testing shows that this improves the performance of the algorithm.

2.6 Solution Method

The ADP planning problem is too large to be solved in reasonable time as a regular MIP using a commercial solver, and even finding a feasible solution is

very time-consuming. This motivates the use of an alternative solution method that produces good results in less time. In the literature, rolling horizons have successfully been applied to several problems in order to deal with large MIPs.

Even though the RHH produces good ADPs, it has some limitations. The biggest drawback is that it does not explore the full planning horizon, which means that it is in some sense myopic. To deal with this a MIP-based improvement heuristic has been developed. It solves the problem for the full planning horizon, but the number of binary variables is reduced based on the given solution from the RHH.

The rest of this section presents the RHH and the improvement heuristic developed for the ADP planning problem. Section 2.6.1 outlines the RHH while Section 2.6.2 describes the improvement heuristic.

2.6.1 The Application of a Rolling Horizon

The general idea of the RHH is to repeatedly solve the MIP for shorter sub-horizons of the planning horizon using branch-and-bound. Each sub-horizon, in the following referred to as time window (TW), consists of two time periods (TP). The two TP s are TP_k^C , which is the central period that will be implemented in iteration k , and the forecasting period TP_k^F . In the forecasting period the model is simplified according to a simplification strategy. Before solving the problem for TW_{k+1} , the variables in TP_k^C are frozen, and the sub-horizon is shifted so that the whole, or the first part of, TP_k^F becomes TP_{k+1}^C . If the central period and the forecasting period have equal length, the forecasting period becomes the new central period. Otherwise only the first part of the forecasting period defines the new central period. The deviations from demand in TP_k^C are also transferred to TP_{k+1}^C . If there is under-delivery in TP_k^C , the demand in TP_{k+1}^C is increased in order to even out the demand deviations over the full planning horizon. Algorithm 1 gives the pseudo code for the RHH, while Figure 2.3 illustrates the RHH graphically. A similar figure is also given in (Mercé and Fontan, 2003).

The Central Period

The central period contains binary variables representing the scheduled voyages that will be frozen in the next iteration. When freezing the variables in the central period, the deviations from contracted demands are transferred to the next central period. One reason for this is that the predicted demands represent the yearly demands evenly distributed over the year, with special constraints for some contracts. These predicted demands seldom fit to an integer number of deliveries since they are a percentile of the yearly demand. To exemplify this, assume a long-term contract with an annual demand of 1000 m^3 of LNG and




Figure 2.3: The rolling horizon heuristic

Algorithm 1 RHH(K : Number of time windows)

```

 $k = 0$ 
while  $(k = k + 1) \leq K$  do
    Solve the mathematical model for the problem defined by  $TP_k^C$  and  $TP_k^F$ 
    Freeze the binary variables  $x_{cvt}$  in the central period  $TP_k^C$ 
    if  $k < K$  then
        for all  $c \in \mathcal{C}^{LT}$  do
            Transfer the deviations from demand in  $TP_k^C$  to  $TP_{k+1}^C$ 
        end for
    end if
end while

```

that the producer has one ship with a capacity of 100 m^3 . If the demand is to be evenly distributed throughout the year the monthly demand is 83.33 m^3 . Since all ships leaving the loading port have to be fully loaded it will be impossible to deliver 83.33 m^3 of LNG each month. The solution is then to over-deliver in 10 of the months, and deliver nothing at all in the two remaining months. This shows that the initially predicted monthly demands will depend on the deliveries in the previous months.

The Frozen Period

Once the computational time limit is reached or the optimal solution is found for a given time window, the variables in the central period are frozen according to a freezing strategy before the next iteration. Several freezing strategies are discussed in the literature. For the multi-item capacitated lot-sizing problem in (Mercé and Fontan, 2003) two different strategies were presented; freezing all

decisions from the previous central periods, or freezing only the periods when to produce but not the production quantities. In their case the latter strategy proved to be most efficient, but for the ADP planning problem this freezing strategy is not possible to implement because of the problem size. We therefore use the first strategy. By freezing all variables, the previous periods do not introduce any variables for the current period. In fact, this actually reduces the feasible space since many ships will not be available from the start of the time window.

The Forecasting Period

It is desirable to include a forecasting period to use information about a larger part of the planning horizon so that solutions that are clearly sub-optimal outside the central period hopefully can be avoided. There are two main issues that will be addressed in this section. First, the restrictions on the forecasting variables will be discussed. Second, it is necessary to determine the length of the forecasting period. The combination of variable restrictions and a given length constitutes a simplification strategy for the forecasting period.

The Variables in the Forecasting Period

There are two options for the variables representing scheduled voyages (x_{cvt}), $0 \leq x_{cvt} \leq 1$ or $x_{cvt} \in \{0, 1\}$. In the RHH the first option is chosen. Using continuous variables will reduce the computational effort required to solve the problem for the forecasting period. The drawback is that the variables can take fractional values in the final solution and, hence, give less accurate forecasting information than binary variables. Still, the role of the forecasting period is to forecast the implications of the decisions made in the central period, hence, less accurate information is acceptable in order to reduce the computational effort. From one iteration to the next, binary restrictions are imposed on the variables in the new central period before re-solving the problem.

The Length of the Forecasting Period

In principle, extending the planning horizon by including additional future periods will provide some incremental benefit simply because a greater amount of relevant information is brought into the analysis. Hence, it is expected that the RHH yields better results as the length of the forecasting period increases. Ideally the full remaining period should be included in order to get information about the demand over the entire planning horizon. However, this approach is not practically feasible due to the problem size. Consequently, the forecasting period needs to be relatively short for the heuristic to be efficient. At the same time, it should be long enough to be affected by, and affect, the decisions made in the central period.

Solution Space Reduction

When applying the principles described in the preceding sections, each sub-problem should be of manageable size. Still, due to the size of the feasible region and many almost symmetrical solutions, high solution times for each time window are expected. To deal with this, a solution space reduction (SSR) has been developed. This is a common strategy in scheduling and inventory problems where it has shown both to improve the solution time and quality (Zanakis, Evans, and Vazacopoulos, 1989). The general idea is to reduce the number of almost symmetrical solutions without cutting off many good solutions. By limiting the number of contracts a ship can deliver to on a given day, the number of scheduled voyages can be decreased, while the flexibility with respect to the LNG volumes delivered is kept. Hence, only variables to a subset of the contracts are generated for each day in the planning horizon. The SSR takes an integer N as input, and creates $1/N$ of all possible variables associated with each contract. The algorithm spreads the variables to different days by taking the sum of the day and the contract number modulo N and only creates variables when this is zero. For contract number 3, only schedules voyages starting on day 1, 5, 9, ... are generated if $N = 4$.

2.6.2 Improvement Heuristic

This section will describe how a modified version of the mathematical model in Section 2.5 is used to improve the ADPs produced by the RHH by generating only a small subset of the variables. By using a feasible ADP as a starting point, the number of variables in the formulation can be limited and, hence, focus the search effort to a more restricted area of the solution space. A similar improvement heuristic is used by Stålhané et al. (2010) with success. The section starts with a description of how a reduced number of variables is created in Section 2.6.2, and then additional constraints added to the formulation are explained in Section 2.6.2. The section ends with a short description of the improvement heuristic.

Variable Generation

In the full mathematical program described in Section 2.5, a binary variable x_{cvt} for every feasible combination of c, v, t is generated. In the improvement heuristic, only a subset of these binary variables is generated. The set \mathcal{S}^* of all scheduled voyages sailed in the ADP is given as input to the improvement heuristic. The sub-sets $\mathcal{S}_{ct}^* \subset \mathcal{S}^*$ are the sets of scheduled voyages serving contract c on day t . For each scheduled voyage in \mathcal{S}^* serving a long-term contract, the contract c and day t are kept fixed, but all ships $v \in \mathcal{V}_c$ are allowed to serve the contract. Thus

the following variables are generated:

$$x_{cvt} \in \{0, 1\}, \quad \forall c \in \mathcal{C}^{LT}, v \in \mathcal{V}_c, t \in \mathcal{T}_v \text{ such that } \mathcal{S}_{ct}^* \neq \emptyset. \quad (2.9)$$

Let \mathcal{C}^{U*} be the set of all long-term contracts c in \mathcal{S}^* that has an under-delivery in the ADP given to the improvement heuristic, and let \mathcal{S}_{gt}^{S*} be the set of all scheduled voyages carrying gas of type g serving the spot contracts on day t . Then, ships sailing scheduled voyages in \mathcal{S}_{gt}^{S*} can be re-routed to a scheduled voyage serving a long-term contract in \mathcal{C}^{U*} . This gives the following variables:

$$x_{cvt} \in \{0, 1\}, \quad \forall g \in \mathcal{G}, c \in \mathcal{C}_g \cap \{\mathcal{C}^{U*} \cup \mathcal{C}^S\}, v \in \mathcal{V}_c, t \in \mathcal{T}_v \text{ such that } \mathcal{S}_{gt}^{S*} \neq \emptyset. \quad (2.10)$$

Maintenance of ships is an area where there is a potential for improvement since taking a ship out of commission for a period of time may have great impact on the ADP. On average maintenance only concerns a fifth of the ships and the part of the planning horizon in which a ship is obligated to start maintenance is very limited. All variables involving the maintenance contract are generated as follows.

$$x_{cvt} \in \{0, 1\}, \quad \forall c \in \mathcal{C}^M, v \in \mathcal{V}_c^M, t \in \mathcal{T}_v^M. \quad (2.11)$$

By only creating the variables mentioned above, the size of the problem in the improvement heuristic is limited. Thus it might be easier to find good feasible solutions to the formulation given in Section 2.5 by using traditional branch-and-bound.

Additional Constraints

For some instances the problem may still be hard to solve. To ease the solution process further, additional constraints are added to the mathematical programming formulation. For each scheduled voyage, only one x_{cvt} variable can be equal to one. The reasoning behind this is that the number of scheduled voyages should be kept constant. To enforce this, the following constraints are added:

$$\sum_{v \in \mathcal{V}_c} x_{cvt} = |\mathcal{S}_{ct}^*|, \quad \forall c \in \mathcal{C}^{LT}, t \in \mathcal{T}. \quad (2.12)$$

Overview of the Improvement Heuristic

The improvement heuristic solves the mathematical model presented in Section 2.5 with some additional constraints. The objective function (2.1) and the constraints (2.2)-(2.6), (2.8) are kept and constraints (2.12) are added. The problem

is reduced in terms of number of variables based on a feasible ADP generated by the RHH. This is accomplished by only generating variables based on (2.9)-(2.11) instead of generating all possible variables as in (2.7). The effect of the improvement heuristic will be discussed in Section 2.7.

2.7 Computational Study

This section presents the tests performed to evaluate the RHH. Section 2.7.1 describes the cases used when testing. In Section 2.7.2 results from solving the full model in Section 2.5 using commercial software are presented. The parameter testing in Section 2.7.3 includes tests on the effect of the improvement heuristic and SSR for a large number of penalty settings. Section 2.7.4 outlines computational results from the best setting found during the parameter tests and compares them with the results for the full model and the construction and improvement heuristic in (Stålhane et al., 2010).

The RHH was implemented in Xpress Mosel ver 2.4.0, using Xpress ver. 19.00.00. The RHH is designed to run on a standalone computer, but due to the number of tests it was decided to run all tests on a cluster after confirming that the performance was similar to that of a standard laptop (A Dell Latitude 630 Intel(R) Core(TM)2 Duo, 2.5 GHz, 3.5GB of RAM, Windows XP). The nodes in the cluster are HP DL 160 G5 computers with an Intel Xeon QuadCore E5472 3.0 GHz processor and 16 GB of RAM running on a Linux operating system. It should be noted that even though the processor used to run these tests have multiple cores, only single thread versions of the programs have been used in order to give running times comparable to using a single core computer.

2.7.1 Description of the Cases

The RHH has been tested on four different cases, A–D. Cases A and C are based on real case data provided by the producer, while cases B and D are reduced versions of A and C, respectively. In these cases some of the ships, contracts, and production volumes are removed. The 16 instances created from these cases are listed in Table 2.1. For each instance the number of ships, number of contracts for each gas type, the length of the planning horizon, the number of berths for each gas type, and a range for the number of scheduled voyages needed are given. The number of scheduled voyages needed is calculated based on the amount of LNG to be shipped and the size of the ships available in the instance.

The instances in Table 2.1 are organized such that A1, B1, C1, and D1 represent the four cases with a one year planning horizon. The remaining instances represent the original cases with a planning horizon of 8, 6 and 4 months, respectively. The 4 and 8 months instances are used for testing the parameters, while

all instances are compared in Section 2.7.4.

Table 2.1: Instances

Instance	Number of Ships	Number of Contracts		Planning Horizon Days	Number of Berths		Number of Voyages
		LLNG	RLNG		LLNG	RLNG	
A1	34	5	3	366	4	1	310 - 350
A2	34	5	3	244	4	1	200 - 230
A3	34	5	3	182	4	1	140 - 170
A4	34	5	3	121	4	1	100 - 120
B1	16	1	3	366	2	1	110 - 140
B2	16	1	3	244	2	1	75 - 90
B3	16	1	3	182	2	1	50 - 70
B4	16	1	3	121	2	1	40 - 50
C1	46	6	11	365	4	1	310 - 350
C2	46	6	11	243	4	1	200 - 240
C3	46	6	11	181	4	1	160 - 200
C4	46	6	11	120	4	1	100 - 130
D1	30	4	4	365	4	1	200 - 240
D2	30	4	4	243	4	1	130 - 160
D3	30	4	4	181	4	1	100 - 130
D4	30	4	4	120	4	1	70 - 100

The cases can be divided into two main scenario classes where cases A and B reflect the current situation, and C and D reflect a future situation when all planned production-trains, ships and loading docks have been completed. The main differences between the cases are: Cases A and B have fewer and smaller ships and the production rates are lower, but there is more excess LNG available to sell in the spot market. Cases A and B also include seasonal variations, the demands are of similar size, and the destinations are, on average, situated closer to the producer's loading port.

The model in Section 2.5 makes it possible to have different sailing times for different periods of the year, but in the data provided a constant sailing time is assumed throughout the year.

To model the possibility to sell excess LNG in the spot market, one contract for each type of LNG is defined. In all cases it is assumed that LNG can be sold on the spot market at a fixed price of USD 150 per m^3 at any time during the planning horizon. Further, at the time the ADP is created, the producer will only know the volume and the delivery times of the LNG to be sold on the spot market, but where the LNG is to be shipped is considered unknown. To make the ADP more robust, it is assumed that all spot cargoes are shipped to the regasification terminal lying furthest away from the loading port.

2.7.2 Exact Solution Method

Here the computational results of solving the mathematical model given in Section 2.5 by commercial optimization software are presented. The model is flexible in terms of partitioning the planning horizon and penalizing the different time intervals. It was decided to penalize monthly and annual deviations from demand.

Since some of the long-term contracts state monthly demands and others a fairly evenly spread of LNG deliveries, penalizing monthly deviations will capture both these aspects. A penalty cost of USD 25 per m^3 for deviations from monthly demands and USD 175 per m^3 for deviations from the annual demand has been used during testing. Over- and under-deliveries to the spot contracts are not penalized. The main reason of testing the exact solution method (ESM) is to get lower bounds in order to evaluate the quality of the RHH. Therefore each instance was run for a maximum of 24 h (86,400 s), using heuristics in every branch-and-bound node in order to produce good solutions. Table 2.2 shows the computational time (CPU(s)), the relative difference between the upper and lower bounds (Gap) and the absolute difference between the upper and lower bounds (UB-LB) given in 1000 USD. The ESM could not solve any instance to optimality within 24 h, but it was very close for some of the 4 month instances, and provided fairly good results on the 6 and 8 months instances as well. However, for the 12 months instances of cases A, C, and D the ESM is far from closing the gap.

Table 2.2: Results from the ESM

Instance	CPU(s)	Gap (%)	UB - LB
A1	86400	396.64	54641
A2	86400	20.62	20141
A3	86400	22.28	8700
A4	86400	3.44	2492
B1	86400	12.24	2985
B2	86400	6.23	3279
B3	86400	6.06	1387
B4	86400	10.63	2371
C1	86400	7.77	112224
C2	86400	3.66	37209
C3	86400	0.68	5514
C4	86400	0.45	2476
D1	86400	8.66	74957
D2	86400	3.39	21499
D3	86400	1.53	7581
D4	86400	0.51	1677

The poor results on the larger instances motivate the development of a heuristic solution method. Even though optimality cannot be proven by the heuristic, it will hopefully produce good solutions in less time than the ESM.

2.7.3 Parameter Testing and Settings for the Heuristic

As described in Section 2.6 the RHH solves the ADP planning problem by iteratively solving shorter time windows of the full problem. The solution time and performance greatly depend on the length of the central period and the forecasting period. For the ADP planning problem the length of the central and forecasting period should be an integer number of months since the demands of the long-term contracts are given on a monthly basis if any seasonal variations

are present. After extensive a priori testing it was chosen to set the central period to 1 month and the forecasting period to 2 months in order to keep the computational effort within reasonable limits. This means that the time window for each sub problem is 3 months. The time used to solve each time window and the improvement heuristic is also limited. If we are unable to prove optimality for a time window the best known solution is accepted as the solution for the time window. If the reduced problem in the improvement heuristic is not solved to optimality, the best known solution when reaching the time limit is returned as the solution of the instance. The time limit is set to 500 s for each time window and to 1000 s for the improvement heuristic. Preliminary testing shows that the results improve for almost all cases when time is increased, but these limits are set in order to have a total running time of less than 2 h for the largest instances, which was a specification provided by the producer.

Monthly and annual deviations from demand are penalized in the mathematical model in Section 2.5. Since the RHH only solves the problem for shorter time windows, the penalties in the full model, C_{cp}^{D+} and C_{cp}^{D-} , cannot be directly translated to penalties used by the RHH. In order to capture the penalties in the full model, penalties are introduced in each sub problem in the RHH. The over- and under-delivery in each time interval in the time periods and the total deviations in the time window are penalized. Tests of the RHH both with and without the SSR were run with penalties on both time windows and time intervals in the time periods ranging from 100 to 200, with an increment of 25. In later discussions a penalty setting of e.g. [100, 200], means a penalty of 100 for the time windows and a penalty of 200 for the time intervals in the time periods.

The parameter testing is done in three steps. First, the effect of the proposed improvement heuristic is studied. Second, the results of the different penalty settings for the RHH with and without SSR are analyzed separately. To rank the settings, the ADP created by the RHH for a given setting is evaluated using the same objective function as the ESM. This means that the penalties introduced in each sub problem are not used during the evaluation. The settings are ranked on the basis of the aggregated objective value for each of the original cases, i.e. the sum of the objective values of A2, A4, B2, and B4 or C2, C4, D2, and D4. An overall ranking of the penalty settings is then made by summing the rankings from the two original cases. Third, a comparison is made between the ten best penalty settings regardless of algorithm used. The best setting will then be used during further testing.

The Effect of the Improvement Heuristic

Section 2.6.2 describes an improvement heuristic which takes the ADP created by the RHH and improves it by solving a reduced version of the mathematical model in Section 2.5. In order to validate that the improvement could justify

the increase in the total running time of the RHH, tests with and without the improvement heuristic were made. All penalty settings for the RHH both with and without SSR were used in the tests. Results from these tests showed average improvements over all eight instances ranging from 2.44% to 10.20%. These results show a significant impact of using the improvement heuristic, and it is used in all succeeding tests.

Tests of Solution Space Reduction

In Table 2.3 the different penalty settings for the parameter tests of the RHH with and without SSR are ranked according to the objective value using objective function (2.1). For each setting the first two columns provide the penalty setting for a given run where (tw) is the penalty used for the time window and (tp) is the penalty used for each time interval in the time periods. Columns 3 and 4 show the ranking of the penalty setting with respect to the total objective value in cases A+B and C+D. The final column gives the total ranking of a particular setting. The overall ranking is based on summing the rankings of columns 3 and 4.

Table 2.3: Parameter tests for RHH with and without SSR

Penalty		Whitout SSR Aggregated objective rank			Penalty		With SSR Aggregated objective rank		
tw	tp	A+B	C+D	Overall	tw	tp	A+B	C+D	Overall
100	100	2	1	1	100	100	1	1	1
150	150	3	5	2	125	200	3	3	2
100	125	7	3	3	100	125	5	2	3
125	125	5	6	4	125	100	2	7	4
125	100	11	2	5	150	100	10	6	5
125	100	12	4	6	175	100	12	4	6
100	175	8	8	7	150	125	11	5	7
100	150	10	7	8	150	150	7	9	8
150	125	9	10	9	125	125	4	13	9
175	125	6	14	10	200	125	14	10	10
175	100	4	18	11	100	150	8	16	11
200	100	1	21	12	175	125	9	15	12
125	150	16	11	13	200	100	13	12	13
100	200	15	17	14	100	175	6	21	14
125	175	13	20	15	200	150	22	8	15
150	175	19	15	16	175	150	19	11	16
200	125	22	12	17	150	175	15	17	17
200	150	25	9	18	125	150	18	14	18
150	200	14	23	19	175	200	17	20	19
175	150	24	13	20	125	175	16	23	20
175	175	23	16	21	100	200	21	19	21
200	175	18	22	22	175	175	23	18	22
125	200	21	19	23	200	200	20	25	23
175	200	17	25	24	150	200	25	22	24
200	200	20	24	25	200	175	24	24	25

For the tests without SSR studying the columns A+B and C+D separately, there is no penalty setting that is non-dominated, i.e. best in all four cases. For example, the penalty setting of [200, 100] is best for the cases A+B, but it ends up as number 21 for C+D. The best penalty setting overall for the four cases is clearly [100, 100]. This setting ranks as number 2 in A+B and number 1 in C+D, with a total ranking sum of 3. For the tests with SSR the best penalty setting

for the four cases is clearly $[100, 100]$.

The Ten Best Penalty Settings from the Parameter Tests

In Table 2.4 the different penalty settings for all the parameter tests of the RHH are ranked. The table has an extra column (Alg) compared to Table 2.3. This column states if SSR was used (ssr) or if the full subproblems were solved (full). When it comes to the results of the RHH with and without SSR the former seems to be slightly better than the latter. On the basis of these results it was decided to run the final tests with penalty setting $[100, 100]$ and SSR. The overall ranking is based on the sum of rankings in columns 4 and 5.

Table 2.4: The ten best penalty settings from the parameter tests

Penalty		Alg	Aggregated objective rank		
tw	tp		A+B	C+D	Overall
100	100	ssr	1	2	1
100	100	full	3	1	2
150	150	full	4	7	3
125	125	full	6	8	4
100	125	full	11	4	5
100	125	ssr	13	3	6
125	200	ssr	9	11	7
100	175	full	12	9	8
125	100	ssr	8	14	9
175	125	full	7	17	10

2.7.4 Computational Results

In this section, the RHH is compared with the ESM and the construction and improvement heuristic (CIH) developed by Stålhane et al. (2010). When comparing with the CIH reported in (Stålhane et al., 2010) the setting CON-LS-2-MIP has been chosen since it gives the lowest average absolute difference between the upper and lower bounds. Table 2.5 presents the results for all instances tested. The relative and absolute gaps between the upper bound and the best lower bound found by the ESM are given for each algorithm.

Only focusing on the relative difference, the performance of all algorithms seems significantly worse for cases A and B compared to C and D. However, this is not the case when looking at the absolute difference between the bounds. The absolute differences are small on almost all instances from cases A and B. The explanation is that in the cases A and B there is a lot of excess LNG to sell, causing the revenue from selling LNG in the spot market to be approximately the same as the transportation and penalty costs giving a lower bound close to zero. This makes the relative differences large in these cases. Cases C and D are much more tightly constrained when it comes to excess LNG. Hence, the objective value is much closer to the sum of transportation and penalty costs, giving relatively smaller differences.

Table 2.5: Comparison of solution quality

Instance	ESM		RHH		CIH	
	Gap (%)	UB-LB	Gap (%)	UB-LB	Gap (%)	UB-LB
A1	396.64	54641	33.07	17002	51.91	23380
A2	20.62	20140	9.35	10069	19.41	19148
A3	22.28	8699	13.02	5501	37.70	13071
A4	3.44	2492	3.98	2866	9.35	6403
B1	12.24	2985	37.27	7431	18.92	4355
B2	6.23	3279	13.18	6507	7.82	4051
B3	6.06	1386	12.08	2613	16.64	3460
B4	10.63	2371	13.51	2938	17.73	3718
C1	7.77	112224	6.07	87694	1.84	26643
C2	3.66	37209	3.78	38379	2.56	25992
C3	0.68	5514	2.17	17674	2.82	22918
C4	0.47	2575	1.77	9703	2.40	13159
D1	8.66	74957	3.95	34216	1.59	13731
D2	3.39	21499	3.50	22238	2.99	18985
D3	1.53	7581	5.21	25839	5.20	25767
D4	0.51	1677	3.44	11289	0.81	2676
Average	31.55	22452	10.33	18873	12.48	14217

Starting with the overall picture and comparing the average performance, both the average percentage and absolute gaps of the ESM are inferior to the gaps of the heuristics. This is mainly due to a very large gap with the ESM for the current ADP planning problem, A1. In general the gaps in solution quality are more stable for the heuristics than the ESM. When comparing the average performance of the heuristics, the percentage gap is slightly better for the RHH while the absolute gap is slightly worse for the RHH.

The ESM produces best solutions on nine of the 16 instances, something that can be explained by the time limits used. The ESM is run for 86,400 s, compared with less than 7200 s for the RHH and even shorter time for the CIH. On all 12 months instances except B1, the heuristics are better, indicating that a heuristic approach is preferable when solving real size instances. Comparing the RHH and the CIH is harder. On the 16 instances tested, the RHH is better on eight. The RHH is better than CIH on six of the eight instances based on cases A and C, i.e. the cases based on real case data provided by the producer.

Table 2.6 shows the total running time in cpu(s) for the ESM, RHH and CIH. The ESM has been run for 24 h for all the instances, while the RHH reached the time limit for some, but not all instances. The instances solved by the CIH all have shorter CPU time than the instances solved by the RHH and ESM.

In addition, Table 2.6 gives information about the CPU time it takes before the ESM obtains a better solution than the one found by the RHH (ESM>RHH). For five of the instances the ESM does not find a better solution than the RHH, hence the value of the column for these instances is ">86400". When comparing objective values from the RHH and the ESM, considering running time, the RHH outperforms the ESM in 10 out of 16 instances. Only studying the instances with a one year planning horizon, the RHH performs better than the ESM in three

Table 2.6: Comparison of computational time

Instance	ESM (s)	RHH (s)	CIH (s)	ESM > RHH (s)
A1	86400	6626	1587	>86400
A2	86400	4564	642	>86400
A3	86400	3539	318	>86400
A4	86400	2521	157	23577
B1	86400	3043	152	325
B2	86400	2528	76	207
B3	86400	2022	42	2932
B4	86400	2015	23	128
C1	86400	5854	1788	>86400
C2	86400	3727	729	42149
C3	86400	3027	410	14348
C4	86400	2552	218	349
D1	86400	5174	530	>86400
D2	86400	3607	226	82392
D3	86400	3069	134	54
D4	86400	2307	66	326

out of four instances. The table also shows that the RHH is more stable with regard to the solution time than the ESM.

2.8 Concluding Remarks

In this paper a rolling horizon heuristic (RHH) for a large scale ship routing and inventory management problem is presented. It is a tactical problem with multiple products, where the goal of the producer is to create an annual delivery program (ADP) that minimizes the cost of fulfilling its long-term contracts, while maximizing revenue from selling LNG in the spot market. The producer operates one loading port with limited inventory and berth capacities for each type of LNG, and a heterogeneous fleet of ships. A long-term contract either outline monthly demands, or state that a certain volume of LNG is to be delivered fairly evenly spread throughout the planning horizon to a given regasification terminal. No inventory management is considered at the regasification terminals. A mixed integer programming model of the ADP planning problem is presented. Over- and under-deliveries within specified time intervals are penalized in the model.

To provide good solutions within reasonable time a RHH is developed. The RHH solves the problem by iteratively solving subproblems with shorter time horizons. Each such subproblem consists of a central period and a forecasting period. After solving a subproblem, the variables in the central period are frozen and the algorithm moves to the next subproblem. A method for solution space reduction (SSR) is also developed to further reduce the number of variables in each subproblem. Finally, an improvement heuristic is implemented on top of the RHH. Here, the ADP created by the RHH is used as input to a reduced version of the mathematical model.

For the computational study instances reflecting the current and estimated fu-

ture situations of the producer are generated. The success of the RHH depends on several factors. (1) The penalties used in the mathematical model and in the RHH are critical. An extensive parameter testing has been performed, and the penalties that performed best on average have been used in the rest of the computational study. (2) Initial testing indicated that the variables representing the scheduled voyages in the central period should remain binary, while they should be continuous in the forecasting period. (3) In addition, the length of the forecasting period is set to twice the length the central period (4) The computational results indicate that the RHH should apply SSR as both solution time and quality improved. (5) The improvement heuristic is important for the success of the RHH and improved the results by 2%–10% for the instances tested.

The computational study shows that the RHH produces high quality solutions to real-world problem instances in a relative short amount of time. None of the instances could be solved to optimality using an exact solution method (ESM), and for most of the instances the RHH produced better quality solutions than the ESM in the same amount of time. The computational study also showed that the RHH is more stable with regard to the solution time than the ESM. When using the RHH, the user can specify the computational time limit for each of the subproblems and the improvement heuristic, such that a maximum computational time is given for the RHH to provide a solution. This makes the RHH well suited as a heuristic in a decision support system (DSS), where the user may want quick feedback in some situations where alternative scenarios can be tested in a short amount of time while considerable time is available in others. The RHH is also compared to a construction and improvement heuristic (CIH) developed by Stålhane et al. (2010). This heuristic is fast compared to the RHH, but for half of the instances the RHH produced better solutions. Both methods are valuable in a DSS, where it is expected that the RHH will produce relatively better solutions when there is plenty of computational time available, while the CIH produces good solutions in a short amount of time.

In the future it would be interesting to study exact solution methods for the ADP planning problem. Here, different decomposition approaches combined with problem specific cuts may be promising in order to produce better bounds and hopefully provide optimal solutions to some of the instances. It would also be interesting to study other formulations of the problem in order to strengthen the bound and reduce the symmetry. By finding, and proving, the optimal solutions to some of the instances we would also be able to provide a better measure of the performance of the RHH.

Other planning problems faced by the producer also contain interesting research alternatives. At the strategic planning level, both contract portfolio management and investment planning in physical assets are important decision areas for this major producer of LNG. At the operational level, the ADP has to be

updated due to unforeseen events and other disruptions to the plan. Therefore it would be interesting to study how to re-optimize the solution, given that certain parts of the plan is fixed, while other parts of it have to be discarded.

The planning horizon for the ADP spans over an entire year, so it is a major issue to produce a robust plan for the producer. Therefore, an interesting topic for further research would be to study how to create an ADP that minimizes the effect of possible disruptions, while ensuring a low cost ADP.

Acknowledgments

This work was carried out with financial support from the Research Council of Norway through the OPTIMAR, DOMinant and Desimal projects.

Bibliography

- Al-Khayyal, F., Hwang, S.-J., 2007. Inventory constrained maritime routing and scheduling for multi-commodity liquid bulk, part I: Applications and model. *European Journal of Operational Research* 176, 106–130.
- Andersson, H., Christiansen, M., Fagerholt, K., 2010. Transportation planning and inventory management in the LNG supply chain. In E. Bjørndal and M. Rönnqvist (Eds.) *Energy, Natural Resources and Environmental Economics*. Springer, 429–441.
- Araujo, S. A., Arenales, M. N., Clark, A. R., 2007. Joint rolling-horizon scheduling of materials processing and lot-sizing with sequence-dependent setups. *Journal of Heuristics* 13, 337–358.
- Baker, K., 1977. An experimental study of the effectiveness of rolling schedules in production planning. *Decision Sciences* 8, 19–27.
- Baker, K. R., Peterson, D. W., 1979. An analytic framework for evaluating rolling schedules. *Management Science* 25, 341–351.
- Bredström, D., Rönnqvist, M., 2006. Supply chain optimization in pulp distribution using a rolling horizon solution approach. Discussion paper 2006:17, Norwegian School of Economics and Business Administration. Department of Finance and Management Science.
- Chidinma, O. U., 2004. What effect does the emerging spot market for LNG have on the financing of gas projects? Tech. rep., University of Dundee, Centre for Energy, Petroleum & Mineral Law & Policy.
- Christiansen, M. 1999. Decomposition of a combined inventory and time constrained ship routing problem. *Transportation Science* 33, 3–26.
- Christiansen, M., Fagerholt, K., 2009. Ch. Maritime inventory routing problems. In: C.A. Floudas and P. Pardalos (Eds.), *Encyclopedia of Optimization*, 2nd Edition. Springer, 1947–1955.
- Christiansen, M., Fagerholt, K., Nygreen, B., Ronen, D., 2007. Maritime transportation. In: Barnhart, C., Laporte, G. (Eds.), *Transportation*. Vol. 14 of

- Handbooks in Operations Research and Management Science. Elsevier Science, 189–284.
- Christiansen, M., Nygreen, B., 1998a. A method for solving ship routing problems with inventory constraints. *Annals of Operations Research* 81, 357–378.
- Christiansen, M., Nygreen, B., 1998b. Modelling path flows for a combined ship routing and inventory management problem. *Annals of Operations Research* 82, 391–413.
- Dimitriadis, A., Shah, N., Pantelides, C., 1997. RTN-based rolling horizon algorithms for medium term scheduling of multipurpose plants. *Computers & Chemical Engineering* 21, 1061 – 1066.
- EIA, 2007. International energy outlook 2007. Tech. rep., Energy Information Administration, U.S. Department of Energy, Energy Information Administration U.S. Department of Energy Washington, DC 20585.
- EIA, 2008. International energy outlook 2008. Tech. rep., Energy Information Administration, U.S. Department of Energy, Energy Information Administration U.S. Department of Energy Washington, DC 20585.
- Flatberg, T., Haavardtun, H., Kloster, O., Løkketangen, A., 2000. Combining exact and heuristic methods for solving a vessel routing problem. *Ricerca Operativa* 36, 94–118.
- Grønhaug, R., Christiansen, M., 2009. Supply chain optimization for the liquefied natural gas business. In: Bertazzi, L., Nunen, J. v., Speranza, M. (Eds.), *Innovations in Distribution Logistics*. Vol. 619 of *Lecture Notes in Economics and Mathematical Systems*. Springer, 195–218.
- Grønhaug, R., Christiansen, M., Desaulniers, G., Desrosiers, J., 2010. A Branch-and-price method for a liquefied natural gas inventory routing problem. *Transportation Science*, trsc.1100.0317v1-trsc.1100.0317.
- Halvorsen-Weare, E., Fagerholt, K., 2009. Routing and scheduling in a liquefied natural gas shipping problem with inventory and berth constraints, working paper. NTNU, Trondheim, Norway.
- IEA, 2007. Natural Gas Market Review 2007: Security in a globalising market to 2015. International Energy Agency.
- Mercé, C., Fontan, G., 2003. MIP-based heuristics for capacitated lotsizing problems. *International Journal of Production Economics* 85, 97 – 111.

- Ronen, D., 2002. Marine inventory routing: shipments planning. *Journal of the Operational Research Society* 53, 108–114.
- Stauffer, L., Liebling, T., 1997. Rolling horizon scheduling in a rolling-mill. *Annals of Operations Research* 69, 323–349.
- Stålhane, M., Rakke, J., Moe, C., Andersson, H., Christiansen, M., Fagerholt, K., 2009. A construction and improvement heuristic for a liquefied natural gas inventory routing problem, working paper. NTNU, Trondheim, Norway.
- Stremersch, G., Michalek, J., Hecq, S., 2008. Decision support software tools for LNG supply chain management. Tech. rep., Gastech.
- Zanakis, S. H., Evans, J. R., Vazacopoulos, A. A., 1989. Heuristic methods and applications: a categorized survey. *European Journal of Operational Research* 43, 88–110.

Paper II

M. Stålhane, J. G. Rakke, C. R. Moe, M. Christiansen, H.
Andersson, K. Fagerholt and I. Norstad:

A Construction and Improvement Heuristic for a Liquefied Natural Gas Inventory Routing Problem

Computers & Industrial Engineering 62 (1), 245–255.

A Construction and Improvement Heuristic for a Liquefied Natural Gas Inventory Routing Problem

Abstract:

We present a large scale ship routing and inventory management problem for a producer and distributor of liquefied natural gas (LNG). The problem contains multiple products, inventory and berth capacity at the loading port and a heterogeneous fleet of ships. The goal is to create an annual delivery program to fulfill the producer's long-term contracts at minimum cost, while maximizing the revenue from selling LNG in the spot market. To solve this problem we have developed a construction and improvement heuristic (CIH). The CIH is a multi-start local search heuristic that constructs a set of solutions using a greedy insertion procedure. The solutions are then improved using either a first-descent neighborhood search, branch-and-bound on a mathematical formulation, or both. Tests on real-life instances show that the CIH provides good solutions in a short amount of time.

3.1 Introduction

During the past decade the demand for energy has soared across the globe and is estimated to increase a further 50% from 2005 to 2030. Natural gas is becoming increasingly important in fulfilling this demand, with its share of the total energy consumption expected to increase from 20% in 2005 to 25% in 2030, EIA (2008). Until recently, natural gas has mainly been transported through pipelines. However, technological advances have made it possible to transport natural gas cost efficiently over greater distances by converting it into liquefied natural gas (LNG) and transporting it in specially designed ships.

More and more countries consider LNG as an alternative way of satisfying their energy demand. The largest suppliers of LNG have customers spread across far-east Asia, Europe and North-America. The increased production is also reflected in the number of LNG ships in service today. The world LNG fleet has grown from 105 ships in 1998 to 257 ships at the start of 2008, and is estimated to expand to about 400 ships by 2015 according to Fearnley (2008). Taking into account that an LNG ship with a capacity of 145 000 m^3 costs roughly 200–250 million USD to build, and the daily charter rate can be as much as 80 000 USD,

it is clear that there is a huge amount of money invested in the transportation of LNG. In addition, a spot market similar to that of crude oil is developing. The value of a single shipload of LNG is approximately 20–30 million USD, making this market very profitable. However, to have a stable and predictable income, most LNG producers have tied the majority of their production to long-term delivery contracts, while the remaining quantity is sold in the spot market. In such a complex and shifting environment decision support systems for routing and scheduling of LNG ships become ever more important.




Figure 3.4: Map showing some of the main sailing routes of the producer.

In this paper we present a ship routing and inventory management problem for a producer and distributor of LNG. The producer has a single liquefaction plant with storage tanks and a connected loading port with limited capacity. From the loading port the LNG is shipped by a heterogeneous fleet of ships to customers world-wide to whom the producer has a contractual obligation to deliver LNG, see Figure 3.4. All shipments are full shiploads. In addition, the producer wants to utilize any excess LNG by selling it in the spot market. The goal is to create an annual delivery program (ADP) to fulfill the producer’s long-term contract obligations at minimum cost, while maximizing the expected marginal contribution from selling excess LNG in the spot market. An ADP is a complete schedule of every ship’s sailing plan for the coming year.

The problem faced by the producer in this case study is similar to problems faced by other producers of LNG, but also producers of other unrefined products like crude oil. These problems often have one location where the product is extracted with limited storage and berth capacity, and many customers to which it is distributed. The fleet is usually fixed and heterogeneous, and each voyage is a round trip, where products are unloaded at one or more ports, before the ship returns empty to the loading port. The main difference between LNG and most other products is the constraint of full ship loads being delivered to only one destination. However, many unrefined products are being extracted in remote locations, with the primary markets located within fairly limited geographical areas. For instance crude oil tankers are fully loaded in the Middle East, and make deliveries to either the US, Western Europe or Japan. Thus the sailing trip may be fairly accurately modelled as one destination when making a schedule for the coming year.

The purpose of this paper is to present a construction and improvement heuristic (CIH) that quickly solves large real world instances of the LNG planning problem described above. The heuristic includes a diversification strategy by constructing a set of initial solutions, and the search for good solutions is intensified by applying branch-and-bound on a mathematical formulation with some variables from the original formulation fixed. We further show that under certain realistic conditions, the heuristic performs much better than the one presented by Rakke, Stålhane, Moe, Andersson, Christiansen, Fagerholt, and Norstad (2011) for the same problem.

The paper is outlined as follows. First we present an overview of recent literature in Section 3.2. In Section 3.3 we present a more detailed description of the problem studied in this paper, before giving a mathematical formulation of the problem in Section 3.4. The construction and improvement heuristic is described in Section 3.5, followed by a computational study in Section 3.6. Finally, some concluding remarks are given in Section 3.7.

3.2 Related literature in maritime LNG transportation

The problem described in this paper is also studied by Rakke et al. (2011), who present a rolling horizon heuristic to create good ADPs. The heuristic solves the problem by iteratively solving mixed integer subproblems with shorter time horizons by branch-and-bound, taking into account decisions made in the previous subhorizon. The solution quality is comparable to the one obtained by the CIH described in this paper, but the solution time is much shorter for the CIH. Halvorsen-Weare and Fagerholt (2009) study a similar problem from the LNG

business. The main difference is that the cargoes for each long-term contract are pre-generated and given time windows. The problem is decomposed into a routing subproblem and a scheduling master problem where berth, inventory and scheduling decisions are handled. Halvorsen-Weare and Fagerholt (2009) do not consider spot market opportunities.

Two other studies of optimizing the LNG supply chain are done by Grønhaug and Christiansen (2008) and Grønhaug, Christiansen, Desaulniers, and Desrosiers (2010). The problem has a more complex structure than the one presented here, as it involves more than one production port, inventory constraints both at the production and consumption ports, variable production and sales, and each LNG ship can visit more than one regasification terminal on a given voyage. When it comes to the computational studies, the problem has fewer ships, and a much shorter planning horizon than the large scale problem presented in this paper. Grønhaug and Christiansen (2008) give both an arc flow and path flow formulation of the problem. There the paths are pregenerated, while Grønhaug et al. (2010) use a branch-and-price algorithm to solve the problem.

Andersson, Christiansen, and Fagerholt (2010) present a more general study of the LNG supply chain, and some of its main characteristics. They consider two problems, one for a producer of LNG and one for a vertically integrated company. Mathematical models are presented for each problem and solution approaches to both models are briefly discussed, but no computational results are given.

3.3 Problem description

In this section we give a detailed description of a real planning problem faced by a producer and distributor of LNG. In addition, we explicitly state the assumptions made in order to formulate a mathematical model of the problem.

At the producer's storage and liquefaction plant, two types of LNG are produced, rich LNG (RLNG) and lean LNG (LLNG). Connected to the plant is a single loading port with one storage tank for each type of LNG, as well as several berths for loading LLNG and one berth for loading RLNG. The inventory levels of the storage tanks should always be between time dependent upper and lower limits. Due to these limitations, ships may potentially wait outside the port for a significant amount of time before loading.

The production rates of LNG are assumed to be fixed, since the producer always aims to produce as much LNG as possible. However, rates may vary, for instance due to maintenance. Since the producer is not the sole supplier to many of its customers, storage on the consumption side is assumed to be unknown and is not taken into account. This means that we assume an infinite storage and berthing capacity at the customer side.

The producer distributes LNG from the loading port to its customers world-wide using a fleet of LNG ships. The fleet is heterogeneous and considered as fixed for the planning horizon. All ships may transport both types of LNG, but only one type at a time. The ships are always fully loaded at the loading port. During a voyage, it is company practice to visit only one regasification terminal before returning to the loading port. We assume that the sailing time of a voyage is known. The total time of a port-stay is approximated to one day for all ships and ports.

All ships have to be maintained at least once every fifth year, and the rule of thumb is that it should be performed within the last month before this deadline. Maintenance is performed at a dry-dock along the sailing lane of a voyage, but the ship occupies one berth at the loading port for one day after the maintenance to perform purge and cool-down procedures.

As of today, the producer controls 35 ships, but the size of the fleet is expected to increase significantly in the coming years. In peak periods, we assume that additional ships can be chartered at a daily rate.

There are also compatibility restrictions on which ships can visit which regasification terminals. This is due to ship acceptance policies at the ports, and the fact that some ships are owned by one, or a group of customers, limiting them to visit only their owner's regasification terminals.

The producer has committed itself to fulfill a set of long-term contracts, each with a duration of 10–25 years. A contract states the annual volume of LNG purchased and the regasification terminal where it is to be delivered. The volumes to be delivered are either specified per month, or simply stated as being delivered fairly evenly spread throughout each year. Since the contracts span many years, there is some flexibility in the volumes that have to be delivered in a given year, but the producer always aims to deliver as closely as possible to the contractual demand.

Since production often is greater than the contractual demands, we assume that the producer has the opportunity to sell excess LNG in the spot market. When planning, we assume that the spot prices are known.

To plan for the coming year, the producer creates an ADP which is a complete annual sailing plan for the entire fleet. The ADP consists of a set of scheduled voyages, where each represents one voyage made by a specific ship from the loading port to a contract's designated regasification terminal and back again. The voyage starts by loading in the producer's port at a particular time and finishes when it arrives back at the loading port. The day of delivery at the regasification terminal and the day the ship returns to the loading port are given implicitly by the start time of the scheduled voyage. Each ship typically undertakes anywhere between 5 and 30 such voyages during the planning horizon.

The aim is to create an ADP that fulfills the long-term contractual demands

at minimum cost, while maximizing the expected marginal contribution from selling excess LNG in the spot market. The ADP created by the producer will then be subject to negotiations with the customers, who may wish to move certain delivery dates. It is therefore in the interest of the producer to make an initial ADP that the customers are likely to accept and thus it is important that the deliveries are fairly evenly spread throughout the planning horizon.

Even though an ADP is a complete plan for a given year, the producer has expressed the need for a tool that enables it to quickly (within 30 min) produce good ADPs in order to perform “what-if” scenario analyses. These include, for instance, the evaluation of adding a new contract, the impact of changing the production schedule, the consequences of chartering out a ship, or the effect of customer demands during negotiations.

3.4 Mathematical model

We start this section by defining the mathematical notation. Then the model is presented and described.

Let \mathcal{G} be the set of liquefied natural gas types and \mathcal{C} be the set of all contracts. The contract set \mathcal{C} can be divided into disjoint sets \mathcal{C}_g containing only the contracts with demand for gas type g . The set of long-term contracts is denoted \mathcal{C}^{LT} , each element having one defined regasification terminal. Let \mathcal{C}^S be the set of artificial spot contracts, one element for each g . This set is created in order to model sale of LNG in the spot market as a scheduled voyage. Let \mathcal{C}^M be the set of artificial maintenance contracts, one element for each g . This set is created to model maintenance as a scheduled voyage. We introduce one maintenance contract for each gas type in order to decide the berth type to use for purge and cool-down operations after the maintenance is performed.

Let \mathcal{V} be the set of all ships available, while \mathcal{V}_c represents the ships allowed to serve contract c . The ships operated by the producer are given by \mathcal{V}^P , where $\mathcal{V}^P \subseteq \mathcal{V}$. In peak periods, additional spot ships may be chartered to serve the contracts or spot market. The ships that need maintenance during the planning horizon is given by the set $\mathcal{V}^M \subseteq \mathcal{V}^P$.

The set of time periods in the planning horizon is given by \mathcal{T} . We partition the time periods into time intervals p , where each p has an earliest time period \underline{T}_p and a latest time period \bar{T}_p . For instance if each time period t is one day, then p may be one week, month, or year. The set of all time intervals is given by \mathcal{P} and the set of time periods of time interval p is given by $\mathcal{T}_p = \{\underline{T}_p, \dots, \bar{T}_p\}$. The set \mathcal{P} may include overlapping time intervals, for instance both a month and a week within the same month. Let $\mathcal{T}_v \subseteq \mathcal{T}$ be the set of time periods where ship v is available to start a voyage, and let \mathcal{T}_v^M be the set of time periods where maintenance may start for all $v \in \mathcal{V}^M$.

Let the parameter T_{cv} be the total time of a scheduled voyage to long-term contract c for ship v . In addition, T_{cv} gives the duration of the maintenance for $c \in \mathcal{C}^M$. Since the destination of a scheduled spot voyage is not known in advance, it was considered most robust to set $T_{cv} = \max_{c' \in \mathcal{C}^{LT}} \{T_{c'v}\}$ for all spot contracts.

Parameter C_{cv}^T represents the transportation and port costs of sailing a scheduled voyage to the regasification terminal of contract c and back again, using ship v . For the spot ships the cost also includes the time charter cost of chartering the ship. The loading capacity of ship v is given by L_v , and the number of berths available at the loading port for loading gas type g is represented by B_g . Parameter P_{gt} is the production volume, while \underline{I}_{gt} and \bar{I}_{gt} are the lower and upper limits on the capacity of the storage tank at the loading port for gas type g at time period t . The estimated value of having LNG of gas type g left in the storage tanks at the end of the planning horizon is given by R_g^I , and R_c^S represents the revenue of selling one unit of gas type g to spot contract c .

One major advantage with the voyage structure in the formulation is that “boil-off” – the daily evaporation of LNG from the ship’s storage tanks – can be pre-calculated. Boil-off is considered a major challenge in the LNG inventory routing problems by Grønhaug et al. (2010), but does not need to be explicitly handled in our model. The demand D_{cp} for LNG of contract c in time interval p is adjusted for boil-off. As mentioned, the long-term contracts state a monthly demand or that the annual demand should be fairly evenly spread throughout the year. In the latter case, the annual demand is partitioned into periodic demand corresponding to the length of time interval p . A cost C_{cp}^D is introduced to penalize the under-delivery in time interval p to contract c . Any over-deliveries are implicitly penalized, since they will reduce the amount of LNG the producer is able to sell in the spot market. It is important that the penalty of under-delivery is higher than the revenue in the spot market, to ensure that the long-term contracts are served. It is also important that the penalty of under-delivery over the entire planning horizon is significantly higher than the penalty of each shorter time interval, to even out deviations over the planning horizon.

The binary variable x_{cvt} represents one scheduled voyage by ship v serving contract c starting in time period t , and exists only for the days the ship is available. The inventory level of gas type g in time period t is given by i_{gt} , while y_{cp} represents the under-delivery of LNG to contract c in time interval p .

$$\begin{aligned} \min \sum_{c \in \mathcal{C}} \sum_{v \in \mathcal{V}_c} \sum_{t \in \mathcal{T}_v} C_{cv}^T x_{cvt} + \sum_{c \in \mathcal{C}^{LT}} \sum_{p \in \mathcal{P}} C_{cp}^D y_{cp} \\ - \sum_{c \in \mathcal{C}^S} \sum_{v \in \mathcal{V}_c} \sum_{t \in \mathcal{T}_v} R_c^S L_v x_{cvt} - \sum_{g \in \mathcal{G}} R_g^I i_{g|\mathcal{T}|}, \end{aligned} \quad (3.13)$$

$$\begin{aligned} & \sum_{c \in \mathcal{C}_g \setminus \mathcal{C}^M} \sum_{v \in \mathcal{V}_c} x_{cvt} \\ & + \sum_{c \in \mathcal{C}^M \cap \mathcal{C}_g} \sum_{v \in \mathcal{V}^M} x_{cv(t-T_{cv}-1)} \leq B_g, \end{aligned} \quad \forall g \in \mathcal{G}, t \in \mathcal{T}, \quad (3.14)$$

$$i_{gt} - i_{g(t-1)} + \sum_{c \in \mathcal{C}_g \setminus \mathcal{C}^M} \sum_{v \in \mathcal{V}_c} L_v x_{cvt} = P_{gt}, \quad \forall g \in \mathcal{G}, t \in \mathcal{T}, \quad (3.15)$$

$$\underline{I}_{gt} \leq i_{gt} \leq \bar{I}_{gt}, \quad \forall g \in \mathcal{G}, t \in \mathcal{T}, \quad (3.16)$$

$$\sum_{v \in \mathcal{V}_c} \sum_{t \in \mathcal{T}_p} L_v x_{cvt} + y_{cp} \geq D_{cp}, \quad \forall c \in \mathcal{C}^{LT}, p \in \mathcal{P}, \quad (3.17)$$

$$\sum_{c \in \mathcal{C}} \sum_{t-T_{cv} < \tau \leq t} x_{cv\tau} \leq 1, \quad \forall v \in \mathcal{V}^P, t \in \mathcal{T}_v, \quad (3.18)$$

$$\sum_{c \in \mathcal{C}^M} \sum_{t \in \mathcal{T}_v^M} x_{cvt} = 1, \quad \forall v \in \mathcal{V}^M, \quad (3.19)$$

$$x_{cvt} \in \{0, 1\}, \quad \forall c \in \mathcal{C}, v \in \mathcal{V}_c, t \in \mathcal{T}_v, \quad (3.20)$$

$$y_{cp} \geq 0, \quad \forall c \in \mathcal{C}, p \in \mathcal{P}. \quad (3.21)$$

The objective function (3.13) minimizes the variable costs incurred by the producer. The first term calculates the transportation and port visit costs of sailing the scheduled voyages. The second term penalizes under-delivery to long-term contracts in different time intervals. The third term subtracts the revenue associated with selling LNG in the spot market. Finally, the fourth term subtracts the value of having LNG left in the storage tanks at the end of the planning horizon.

Constraints (3.14) make sure that the number of ships occupying a berth on any given day does not exceed the number of berths. In the second term the time index is adjusted because a scheduled voyage to maintenance contracts occupies a berth after the ship returns to the loading port. These constraints are only defined for the loading port.

Constraints (3.15) and (3.16) make sure that the inventory levels in the storage tanks are between their upper and lower limits. Constraints (3.17) are soft and ensure that the demand of each long-term contract is met in every time interval.

Constraints (3.18) restrict each ship to sail at most one scheduled voyage on any given day. Constraints (3.19) state that every ship that is planned for maintenance will undertake maintenance exactly once. Finally, constraints (3.20) limit the scheduled voyages to be binary on the days the ship is available, while (3.21) make sure the under-deliveries are non-negative.

We now define three expressions that will be used extensively in the following section of the paper. An ADP is considered to be *berth-feasible* if it does not violate (3.14), *inventory-feasible* if it does not violate (3.15) and (3.16), and *routing-feasible* if it does not violate (3.18).

For the remaining of this paper we assume that the length of a time period is one day, and that the set of time intervals \mathcal{T}_p consists of all the months in the planning horizon, as well as one time interval covering the entire planning horizon.

3.5 Solution approach

This section presents the construction and improvement heuristic (CIH) that solves the ADP planning problem presented above. The heuristic uses a multi-start approach where a set of initial solutions are constructed, and then improved by intensifying the search in the neighborhood of each solution. The multi-start heuristic is chosen because it is difficult to design a search operator which will allow the search to investigate a large portion of the feasible region, due to the global inventory- and berth-constraints. Thus a more natural approach is to use a multi-start heuristic to diversify the search by producing a large number of initial solutions, each of which may then be improved using simple search operators.

We start by outlining the construction phase, addressing how ships, contracts and start days are put together to form a feasible set of scheduled voyages. We then describe two different ways in which this initial ADP may be improved. One is based on a local search, while the other is based on the mathematical programming formulation in Section 3.4.

3.5.1 Construction heuristic

To plan for the coming year, the producer creates an ADP consisting of a set of scheduled voyages \mathcal{S} . A scheduled voyage s represents one voyage made by a specific ship from the loading port to a contract's designated regasification terminal and back again. Let $s = (c, v, t)$, where v denotes the ship sailing the voyage, c the contract served and t the day the loading process starts.

The construction heuristic creates a feasible ADP by going through the planning horizon from beginning to end, adding new scheduled voyages to the solution. Algorithm 2 shows the pseudo code for the construction heuristic. Let \mathcal{S} be an empty set of scheduled voyages. Further, let \mathcal{M} be the set of months in the planning horizon and \underline{T}_m and \bar{T}_m the first and last day of month m . Each iteration starts with the inventory being updated, before the algorithm goes through each contract c in a greedy fashion according to the ordered set of contracts $\mathcal{C}_{\alpha\beta}^Q$,

ranked according to the contract ranking parameter α and the inventory control parameter β . $\mathcal{C}_{\alpha\beta}^Q$ contains all contracts from \mathcal{C}^{LT} and \mathcal{C}^S .

Ships are then selected from the ordered set \mathcal{V}_c^Q , that may serve the chosen contract c . The ship selected must be available at the loading port at most κ days after t , where κ is a look-ahead parameter. This look-ahead functionality is similar to that recommended by both Ronen (1986) and Atkinson (1994), and tells us how long it is acceptable to wait for a better ship to become available.

Let T_v^{RF} be the first day it is routing-feasible to send out ship v , and T_{cv}^{IF} the first day where it is both berth- and inventory-feasible to load v with the LNG demanded by contract c . Then, the earliest possible start day of a scheduled voyage serving c using v is $T^E = \max\{T_v^{RF}, T_{cv}^{IF}\}$. When a ship v is found that is able to start within the same month and within κ days, a scheduled voyage (c, v, T^E) is added to \mathcal{S} and the algorithm continues with the next contract. Using T^E , and data input regarding sailing times to the selected contract's regasification terminal, T_v^{RF} is updated to be the next day it is routing-feasible to use v .

The berth capacity and inventory levels must also be updated after a new scheduled voyage is added. It must be ensured that there is LNG available to sail the voyage on T^E , and at the same time make sure that the storage tank will not exceed its capacity between t and T^E .

Adjusting monthly demand

Due to the different sizes of the ships and the fact that all cargoes have to be full shiploads, contracts will almost never have their monthly demands met exactly. This may accumulate to large deviations between the amount of LNG delivered and the contractual demands over the entire planning horizon if we do not adjust the remaining monthly demands accordingly. Let $\mathcal{S}_{ct} \subset \mathcal{S}$ be the set of scheduled voyages starting on day t , serving contract c and let $v(s)$ be the ship sailing the scheduled voyage s . Further let D_{cm} and D_{cm}^* be the original and updated demand for contract c in month m . After the last day of each month, the next month's demand is calculated in the following way:

$$D_{c(m+1)}^* = D_{c(m+1)} + [D_{cm}^* - \sum_{t=T_m}^{\bar{T}_m} \sum_{s \in \mathcal{S}_{ct}} L_{v(s)}]. \quad (3.22)$$

Contract ranking

The main objective of our contract ranking is to get the deliveries to each contract fairly evenly spread throughout the planning horizon. We apply two approaches that both consider the remaining demand in the current month, the first uses the remaining volume and the second the remaining percentage of the monthly

Algorithm 2 CreateADP(α, β, κ)

```

 $\mathcal{S} = \emptyset$ 
 $m = 1$ 
for  $t \in \mathcal{T}$  do
  if  $t > \bar{T}_m$  then
     $m = m + 1$ 
  end if
  for  $g \in \mathcal{G}$  do
     $i_{gt} = i_{gt-1} + P_{gt}$ 
  end for
  for  $c \in \mathcal{C}_{\alpha\beta}^Q$  do
    for  $v \in \mathcal{V}_c^Q$  do
       $T^E = \max\{T_v^{RF}, T_{cv}^{IF}\}$ 
      if  $T^E \leq \min\{t + \kappa, \bar{T}_m\}$  then
         $\mathcal{S} = \mathcal{S} \cup (c, v, T^E)$ 
         $T_v^{RF} = T_v^E + T_{cv}$ 
        update inventory levels and berth capacities
        go to next contract
      end if
    end for
  end for
end for

```

demand. The reason for considering monthly demands instead of the demand for the entire planning horizon is that some of the contracts have seasonal variations in demand that would not have been adjusted for, if we only considered demand for the entire planning horizon.

Contract ranking parameter α

We combined the two approaches into one ranking scheme for the contracts. The scheme sorts the contracts on remaining percentage if the difference in remaining percentage is large, and on remaining volume if it is not. More formally, let ρ_{cm} be the percentage of the monthly demand for contract c that is left to be scheduled for delivery in month m , and D_{cm}^* the demand for the contract in the current month. The algorithm for ranking two contracts, c_1 and c_2 , in month m is given in Algorithm 3.

Inventory control parameter β

For most months during the planning horizon the total production will be greater than the total demand of the long-term contracts. This excess LNG may be sold

Algorithm 3 Contract ranking scheme

```

if  $|\rho_{c_1m} - \rho_{c_2m}| \leq \alpha$  then
    return  $\operatorname{argmax}\{\rho_{c_1m}D_{c_1m}^*, \rho_{c_2m}D_{c_2m}^*\}$ 
else
    return  $\operatorname{argmax}\{\rho_{c_1m}, \rho_{c_2m}\}$ 
end if

```

in the spot market. As with the long-term contracts the spot sales should be spread across the planning horizon. The main reason is that only a subset of the ships can take spot cargoes and we therefore want to spread out these voyages to avoid having to charter additional ships.

In order to spread the spot cargoes throughout the planning horizon, one artificial spot contract is added to $\mathcal{C}_{\alpha\beta}^Q$ for each gas type. The demand of these artificial contracts are set to be β times the difference between total production and total adjusted contractual demand for the given month, where $\beta \in [0, 1]$. The parameter β can thus be seen as a parameter controlling the inventory levels at the start of the next month, and the value of β will be varied in the multi-start heuristic.

Ship ranking

Each contract in $\mathcal{C}_{\alpha\beta}^Q$ has a queue of ships associated with it, \mathcal{V}_c^Q , which contains an ordered sequence of the ships that are allowed to serve the contract. The ships are ranked according to the total number of contracts the ship may serve, where the ships are sorted in increasing order. Ships which may serve the same number of contracts are sorted in increasing order by their cost to capacity ratio. This approach is similar to the one used in the heuristic described by Ronen (1986), and ensures that less flexible ships are used before more flexible ones, and that the cheapest ships among equally flexible ships are used first.

Maintenance

Some ships need to be scheduled for maintenance each year. The maintenance for ship v is required to begin during a given time interval $[\underline{T}_v^M, \bar{T}_v^M]$. In order to add a scheduled maintenance voyage in the construction heuristic we add a rule to the heuristic ensuring that a ship cannot be selected to sail a normal scheduled voyage that ends after the last day of the maintenance period \bar{T}_v^M , unless maintenance has already been scheduled. Once it is decided to schedule maintenance, it is scheduled as early as possible.

The objective function

The objective function presented here is equivalent to objective function (3.13) using the planning horizon and months as time intervals, but will be restated using the symbols defined in Section 3.5.

Let $D_c = \sum_{m \in \mathcal{M}} D_{cm}$ for all $c \in \mathcal{C}^{LT}$ be the total demand of contract c , $\mathcal{S}_c \subseteq \mathcal{S}$ the set of all scheduled voyages serving contract c , and $\mathcal{S}_{cm} \subseteq \mathcal{S}_c$ the set of all scheduled voyages serving c in month m . Further, let C_s^T be the transportation cost associated with voyage s , and let C_c^D and C_{cm}^D be the penalty cost per unit of under-delivery of LNG to contract c for the planning horizon and each month respectively. Finally, $v(s)$ denotes the ship sailing scheduled voyage s , and $c^S(g) \in \mathcal{C}^S \cap \mathcal{C}_g$ the spot contract for gas type g . The function used to evaluate the ADP is then given as:

$$\begin{aligned}
 z = & \sum_{s \in \mathcal{S}} C_s^T + \sum_{c \in \mathcal{C}^{LT}} \sum_{m \in \mathcal{M}} \max\{0, D_{cm} - \sum_{s \in \mathcal{S}_{cm}} L_{v(s)}\} \cdot C_{cm}^D \\
 & + \sum_{c \in \mathcal{C}^{LT}} \max\{0, D_c - \sum_{s \in \mathcal{S}_c} L_{v(s)}\} \cdot C_c^D \\
 & - \sum_{g \in \mathcal{G}} \sum_{s \in \mathcal{S}_{c^S(g)}} L_{v(s)} R_{c^S(g)}^S - \sum_{g \in \mathcal{G}} R_g^I i_{g|\mathcal{T}|}
 \end{aligned} \tag{3.23}$$

The first term summarizes the cost associated with each scheduled voyage, while the second penalizes monthly under-delivery to the long-term contracts. The third term penalizes under-delivery to the long-term contracts over the entire planning horizon, and the fourth term subtracts the revenue of selling LNG in the spot market. Finally, the fifth term subtracts the value of having LNG left in the storage tanks on the last day of the planning horizon.

3.5.2 Local search heuristic

Here we present an improvement heuristic which takes a feasible ADP and tries to decrease the value of the objective function described in Section 3.5.1 through a local search. Since the CIH is a multi-start heuristic the goal of the local search is to quickly find better solution by exploring the neighborhood around a feasible solution. We have implemented five local search operators that each defines a neighborhood $N(\mathcal{S})$ to a given feasible ADP \mathcal{S} . The search is done by going through these operators in a “first improvement” fashion until no more improvements are found. The five neighborhood operators are:

1. *Changing contract:* $(c, v, t) \rightarrow (c', v, t)$

The ship v starting a scheduled voyage to contract c on day t is assigned another contract c' .

2. *Changing ship*: $(c, v, t) \rightarrow (c, v', t)$
The contract c served on a scheduled voyage starting day t is assigned another ship v' .
3. *Swapping contracts*: $(c, v, t), (c', v', t') \rightarrow (c', v, t), (c, v', t')$
The contracts of two scheduled voyages are swapped.
4. *Swapping ships*: $(c, v, t), (c', v', t') \rightarrow (c, v', t), (c', v, t')$
The ships of two scheduled voyages are swapped.
5. *Remove a scheduled voyage*
Remove the scheduled voyage (v, c, t) from \mathcal{S} .

Since the local search uses the “first improvement” strategy, the sequence in which the scheduled voyages are ordered and tested for each operator will influence the quality of the solution. Each operator sorts the scheduled voyages according to c using the contract ranking $\mathcal{C}_{\alpha\beta}^Q$, and then according to t in chronological order. Note that the contract ranking at the end of the planning horizon will be ranked according to fulfillment of the total yearly demand, as deviation from demand in each of the previous months are passed onto the next. For operator 2, \mathcal{V}_c^Q is used to order the sequence in which ships are inserted into a scheduled voyage. All operators are checked with respect to berth-, inventory-, and routing-feasibility to evaluate if they are feasible.

By construction, there will only be LNG left in the storage tanks at the end of the planning horizon if all the long-term contracts with a demand for a particular gas type are satisfied. We may therefore add additional scheduled voyages at the end of the planning horizon for the artificial spot contracts. Adding these scheduled voyages is basically a multiple-knapsack problem, and we solve it through dynamic programming. The multiple-knapsack problem is NP-hard and thus only non-polynomial algorithms for solving it are known. However, since the number of variables is small, the solution time is negligible.

3.5.3 Mixed integer programming heuristic

In this section we describe a mixed integer programming (MIP) heuristic based on a modified version of the mathematical model presented in Section 3.4. The idea is to explore more thoroughly the neighborhood around local minima in order to further improve the value of an ADP. Archetti, Bertazzi, Speranza, and Hertz (2010) use a similar approach to improve the solutions of their tabu search heuristic to solve a vehicle routing problem with inventory constraints. The same MIP heuristic as the one presented here is also used by Rakke et al. (2011) as part of their rolling horizon heuristic.

Variable generation

In the mathematical model, every feasible (c, v, t) combination is represented by a binary variable x_{cvt} . Here we suggest creating only a subset of these binary variables based on the set of scheduled voyages \mathcal{S} given by the ADP used as input to the model. Recall that \mathcal{S}_{ct} is the set of scheduled voyages sailing to contract c on day t . For each scheduled voyage going to a long-term contract we keep the contract c and day t fixed, but allow all ships $v \in \mathcal{V}_c$ to sail this scheduled voyage. Thus we create the following variables:

$$x_{cvt} \in \{0, 1\}, \quad \forall c \in \mathcal{C}^{LT}, v \in \mathcal{V}_c, t \in \mathcal{T}_v, \mathcal{S}_{ct} \neq \emptyset. \quad (3.24)$$

For scheduled voyages going to the artificial spot contracts we allow the new formulation to re-route the voyages to a long-term contract if it is profitable to do so. Therefore we create the following variables:

$$x_{cvt} \in \{0, 1\}, \quad \forall g \in \mathcal{G}, c \in \mathcal{C}_g \setminus \mathcal{C}^M, v \in \mathcal{V}_c, t \in \mathcal{T}_v, \mathcal{S}_{c^S(g)t} \neq \emptyset. \quad (3.25)$$

Maintenance of ships is an area where there is potential for improvement because taking a ship out of commission may have great impact on the ADP. Since maintenance on average only concerns a fifth of the ships and the impact of changing maintenance by only a few days may be substantial, all the variables involving the maintenance contract will be created:

$$x_{cvt} \in \{0, 1\}, \quad \forall c \in \mathcal{C}^M, v \in \mathcal{V}_v^M, t \in \mathcal{T}_v^M. \quad (3.26)$$

By only creating the variables mentioned above, we limit the size of the problem to a very small fraction of its original size.

Additional constraints

For some instances the problem may still be hard to solve using branch-and-bound. To ease the solution process further, additional sets of constraints are added to the formulation. For each scheduled voyage, only one x_{cvt} variable can be equal to one. The reasoning behind this is that we want to keep the number of scheduled voyages constant. We add the constraints:

$$\sum_{v \in \mathcal{V}_c} x_{cvt} = |\mathcal{S}_{ct}|, \quad \forall c \in \mathcal{C}^{LT}, t \in \mathcal{T}, \mathcal{S}_{ct} \neq \emptyset, \quad (3.27)$$

$$\sum_{v \in \mathcal{V}_c} x_{cvt} \leq |\mathcal{S}_{ct}|, \quad \forall c \in \mathcal{C}^S, t \in \mathcal{T}. \quad (3.28)$$

The reason why constraints (3.28) are \leq constraints is that we have generated

some variables in (3.25) which are not part of these constraints. If one of these variables are equal to one, all the variables in one constraint in (3.28) will most likely have to be equal to zero, or the ADP will not be inventory-feasible.

Overview of the mixed integer programming heuristic

The MIP heuristic solves a modified version of the formulation presented in Section 3.4, based on input from a feasible ADP. The objective function (3.13) and constraints (3.14) – (3.19) are all kept, and, in addition, (3.27) and (3.28) are added to the model. Instead of generating variables as stated in (3.20) and (3.21), only the variables given by (3.24) – (3.26) are created. By solving this reduced version of the problem we can find good feasible solutions quickly using branch-and-bound.

3.5.4 Summary of the construction and improvement heuristic

We have presented a CIH to solve the ADP planning problem with respect to the criteria mentioned above. It is a multi-start heuristic where a number of initial solutions are generated by a constructive heuristic. These solutions are then improved by either a local search, and/or by solving a restricted version of the mathematical model presented in Section 3.4. Figure 3.5 shows the flow of data through the different parts of the heuristic. The naming convention of the three different outputs will be used in Section 3.6, where a computational study of the heuristic is conducted.

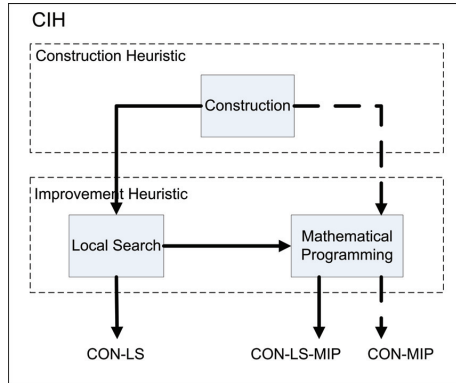


Figure 3.5: Flowchart showing the different modules of the CIH and how they interact.

3.6 Computational results

In this section we test the CIH to see how suitable it is to solve the ADP planning problem. First we give a short description of the instances used in the computational study. The testing of the parameters α , β and κ is then presented, in order to determine the best value range for each parameter. Finally we present the results given by different variants of the CIH on all instances.

The CIH has been implemented in Java version 6.0. The development is done using the Eclipse Ganymede workbench and compiler. Both the full mathematical model and the MIP heuristic were implemented in Xpress-IVE 1.19.00, with Xpress-Mosel 2.4.0, and solved by Xpress-Optimizer 19.00.00. The tests have been performed on a HP DL 160 G5 computer with an Intel Xeon QuadCore E5472 3.0 GHz processor, 16 GB of RAM and running on a Linux operating system. Even though the processors used to run these tests have multiple cores, only single thread versions of the programs have been run, to give running times comparable to using a single core computer.

3.6.1 Test case description

We have tested the CIH on four cases, A, B, C, and D, based on real case data provided by the producer. From these cases we have created 20 instances which are listed in Table 3.7. For each instance we give the number of ships, number of contracts for each gas type, the length of the planning horizon, the number of berths for each gas type, the inventory to production ratio (I/P), and a range for the number of scheduled voyages needed. The I/P-ratio is calculated as the size of the storage tanks in the loading port divided by the average daily production. This measure was used by Halvorsen-Weare and Fagerholt (2009) and gives an indication of how tightly constrained the problem is with respect to the inventory constraints. The number of scheduled voyages needed is calculated based on the amount of LNG to be shipped and the minimum and maximum ship sizes available in the instance.

The instances in Table 3.7 are organized such that instances A1, B1, C1, and D1 represent each of the four cases with a one year planning horizon, while the instances numbered 2 – 5 are equivalent, with the exception of the length of the planning horizon.

Case A is tightly constrained with respect to the RLNG inventory, while case C is tightly constrained with regard to inventory for both gas types. On the other hand, case B has a lot of maneuvering room in how to schedule its voyages carrying LLNG, since it will take more than a month to fill up the storage tank.

The cases can be divided into two main scenario classes where cases A and B reflect the current situation, and C and D reflect a future situation when planned production-trains, ships and loading docks all have been completed. The main

A Construction and Improvement Heuristic for a LNG Inventory Routing Problem

Instance	Number of ships	Number of contracts		Planning horizon days	Number of berths		I/P		Number of voyages
		LLNG	RLNG		LLNG	RLNG	LLNG	RLNG	
A1	34	5	3	366	4	1	9.21	6.17	310 - 350
A2	34	5	3	244	4	1	9.02	6.08	200 - 230
A3	34	5	3	182	4	1	9.05	5.98	140 - 170
A4	34	5	3	121	4	1	9.01	5.50	100 - 120
A5	34	5	3	91	4	1	9.23	5.38	80 - 100
B1	16	1	3	366	2	1	44.60	11.35	110 - 140
B2	16	1	3	244	2	1	44.21	11.29	75 - 90
B3	16	1	3	182	2	1	44.37	11.11	50 - 70
B4	16	1	3	121	2	1	44.14	10.21	40 - 50
B5	16	1	3	91	2	1	45.23	10.00	30 - 40
C1	46	6	11	365	4	1	6.05	6.07	310 - 350
C2	46	6	11	243	4	1	6.08	6.19	200 - 240
C3	46	6	11	181	4	1	6.04	6.39	160 - 200
C4	46	6	11	120	4	1	6.07	6.22	100 - 130
C5	46	6	11	90	4	1	6.18	5.82	70 - 100
D1	30	4	4	365	4	1	9.08	9.37	200 - 240
D2	30	4	4	243	4	1	9.14	9.73	130 - 160
D3	30	4	4	181	4	1	9.60	9.97	100 - 130
D4	30	4	4	120	4	1	9.11	9.55	70 - 100
D5	30	4	4	90	4	1	9.35	8.94	50 - 70

Table 3.7: Instances

differences between the cases are: cases A and B have fewer and smaller ships than cases C and D, the production rates are lower, and there is more excess LNG available to sell in the spot market. They also include seasonal variations, the contract demands are of similar volume sizes, and the contract destinations are, on average, situated closer to the producer's loading port. Cases B and D are down-scaled versions of cases A and C respectively, where some of the ships, contracts, and production volumes are removed. All cases are also used by Rakke et al. (2011), where a comparison between the heuristics is presented.

3.6.2 Exact solution method

In this section the computational results of the mathematical model given in Section 3.4 are presented. Each instance was run for a maximum of 24 h (86400 s) and Table 3.8 shows the computational time (CPU(s)), the optimality gap in percentage (opt. gap) and the difference between the upper and lower bound (UB - LB) given in 1000 USD. Optimality could not be proven for any of the instances within the time limit, though the 3 and 4 month instances are very close. The 6 months instances provide fairly good results, however for the 8 and 12 months instances of cases A, C and D, the best solutions found after 24 h are far from the lower bound. Especially the 12 month instances of these cases have big gaps.

Focusing only on the percentage gap indicates a poor performance on the first 10 instances compared to the last 10 instances. However, the picture is quite dif-

Instance	CPU(s)	opt. gap	UB - LB
A1	86400	396.64 %	54641
A2	86400	20.62 %	20141
A3	86400	22.28 %	8700
A4	86400	3.44 %	2492
A5	86400	3.82 %	3107
B1	86400	12.24 %	2985
B2	86400	6.23 %	3279
B3	86400	6.06 %	1387
B4	86400	10.63 %	2371
B5	86400	8.41 %	1464
C1	86400	7.77 %	112224
C2	86400	3.66 %	37209
C3	86400	0.68 %	5514
C4	86400	0.45 %	2476
C5	86400	0.22 %	877
D1	86400	8.66 %	74957
D2	86400	3.39 %	21499
D3	86400	1.53 %	7581
D4	86400	0.51 %	1677
D5	86400	0.47 %	1132

Table 3.8: Test of the exact solution method

ferent if looking at the difference in value between the upper and lower bounds. It is reasonable that the difference is bigger for the instances with longer planning horizons as more voyages are made, and there are more time intervals for which under-delivery is penalized, resulting in a larger objective value. However, instance A1, which has the biggest percentage gap, is much closer to the optimal solution than instances C1 and C2 which have small percentage gaps. The reason why instances from cases A and B have larger percentage gaps is that these instances have a lot more excess LNG. For cases A and B the income from selling LNG in the spot market is approximately the same as the total costs, and therefore the objective values become close to zero. This makes the percentage gap large, even though the total costs related to the instances are similar to those in cases C and D.

3.6.3 Parameter testing

The CIH is a multi-start heuristic, where the starting points are determined by the values of contract ranking parameter α , the inventory control parameter β , and the look-ahead parameter κ . In this section we test instances A3, B3, C3, and D3 for several combinations of the three parameters to see which are more likely to provide good solutions. We test the following setting for the parameters: $\alpha \in [0, 1]$ with increments of 0.025, $\beta \in [0, 1]$ with increments of 0.025, and $\kappa \in \{1, \dots, 31\}$. Each tested parameter value for a given parameter is tested over all the possible combinations of the other parameter values.

Figure 3.6 shows the deviation between the best obtained ADP and the best




Figure 3.6: Graph showing the variation in the minimal objective value for each κ value, compared to the best found objective value, for each of the four instances.

ADP for that particular κ value, for each of the four instances. Which κ value performs better differs a lot between the instances, and especially instance C3 is a very sensitive to the κ value. However, there is a correspondence between the κ values that provide the best ADPs, and the I/P-ratio given in Section 3.6.1. For instance A3 which has I/P-ratios of 9.05 and 5.98 the best ADPs are created with κ values between 4 and 10. These observations indicate that there is a connection between the I/P-ratio and the best κ values and we therefore use the I/P-ratio as a basis for our κ values. This is logical since the I/P-ratio gives the number of days before an empty storage tank fills up, and thus gives an indication of the frequency with which voyages need to be scheduled. Since some instances have different I/P-ratios for each gas type, we replace κ with κ_g in Algorithm 2.

To test what range of κ_g provides the best ADPs, we tested all four instances with $\kappa_g = \lfloor (\bar{I}_g - \underline{I}_g) / P_g \rfloor + \sigma_g$, where P_g is the daily average production rate of gas type g , and $\sigma_g \in \{-5, \dots, 2\}$. These tests did not give any clear indication on what combinations of κ values performed better, and it was therefore decided to do the remaining testing with the same σ value for both gas types.

For parameter β low values seem to be inferior. We therefore limit β to the range $[0.25, 1]$ in the remaining tests. This result is reasonable since for instances with much excess LNG, one needs to ship some LNG out to keep the solution inventory-feasible. For instances with little excess LNG, the construction heuris-

tic is less sensitive to the value of the β multiplier, since the demand of the artificial spot contract is so small that it will not be served. The parameter tests did not give any indication of which values of the α parameter are superior, and thus the α parameter will be tested over the range $[0, 1]$ in the remaining computational study.

3.6.4 Computational results

In order to evaluate the effect of using the local search and the MIP heuristic to improve the ADP created by the construction heuristic, we decided to test six different variants of the CIH. For each variant the heuristic is run over all the possible combinations of the parameters, and returns the best ADP.

- CON-LS-1: The construction heuristic with the local search, using a granularity of 0.1 for α and β , and $\sigma \in \{-5, \dots, 2\}$.
- CON-LS-2: The construction heuristic with the local search, using a granularity of 0.05 for α and β , and $\sigma \in \{-5, \dots, 2\}$.
- CON-LS-3: The construction heuristic with the local search, using a granularity of 0.025 for α and β , and $\sigma \in \{-5, \dots, 2\}$.
- CON-MIP: The best ADP produced by the construction heuristic using a granularity of 0.05 for α and β , and $\sigma \in \{-5, \dots, 2\}$, and then run the MIP heuristic.
- CON-LS-1-MIP: The best ADP produced by CON-LS-1, and then run the MIP heuristic.
- CON-LS-2-MIP: The best ADP produced by CON-LS-2, and then run the MIP heuristic.

Table 3.9 shows the computational results for CON-LS-1, CON-LS-2, and CON-LS-3. The computational time (CPU(s)) is given in seconds. The gap between the upper bound found by the CIH and the lower bound obtained by running the mathematical model is given both as a percentage (opt. gap), and as an absolute difference in value (UB – LB) given in 1000 USD. None of the instances were solved to optimality, and thus the quality of the lower bound is not known. We see that the optimality gap is much lower for the CON-LS variants than it is for the mathematical model, see Table 3.8, even though the running time is much shorter.

CON-LS-3 produces better ADPs than CON-LS-2 in 11 out of 20 instances, and is, on average, 1.3 million USD closer to the lower bound. Especially for instances A2, C2, and C3 the improvements are substantial (more than 4 million USD).

A Construction and Improvement Heuristic for a LNG Inventory Routing Problem

Instance	CON-LS-1			CON-LS-2			CON-LS-3		
	CPU(s)	opt. gap	UB-LB	CPU(s)	opt. gap	UB - LB	CPU(s)	opt. gap	UB-LB
A1	108	61.00 %	25923	418	51.91 %	23381	1587	48.79 %	22435
A2	45	19.41 %	19149	171	19.41 %	19149	642	14.32 %	14757
A3	22	46.88 %	15238	83	44.83 %	14779	318	40.36 %	13729
A4	11	16.91 %	10837	41	15.54 %	10074	157	13.67 %	9011
A5	6	18.83 %	13375	23	18.03 %	12895	87	15.67 %	11434
B1	10	20.10 %	4582	40	18.92 %	4356	152	18.92 %	4356
B2	5	9.31 %	4761	20	8.84 %	4541	76	8.84 %	4541
B3	3	18.26 %	3745	11	18.25 %	3744	42	18.25 %	3744
B4	1	21.32 %	4339	6	21.32 %	4339	23	14.54 %	3134
B5	1	21.00 %	3275	4	11.94 %	2013	15	11.94 %	2013
C1	124	2.29 %	33028	474	2.21 %	31858	1788	2.21 %	31858
C2	50	3.28 %	33338	192	3.28 %	33338	729	2.57 %	26138
C3	28	3.14 %	25529	109	3.14 %	25526	410	2.20 %	17924
C4	15	2.62 %	14350	58	2.62 %	14350	218	2.39 %	13095
C5	15	1.79 %	7207	55	1.79 %	7207	206	1.79 %	7207
D1	37	2.56 %	22155	141	2.47 %	21345	530	2.47 %	21345
D2	15	3.04 %	19310	60	2.99 %	18985	226	2.99 %	18965
D3	9	5.20 %	25767	35	5.20 %	25767	134	5.06 %	25107
D4	4	4.09 %	13429	17	3.51 %	11526	66	3.51 %	11526
D5	3	1.11 %	2655	11	1.11 %	2655	44	1.11 %	2655
Average	25.6	14.11 %	15100	98.45	12.87 %	14591	372.5	11.58 %	13249

Table 3.9: Local search heuristic tests

Taking into account the fact that CON-LS-3's running time is still relatively short for all instances, it seems to be worth the additional computational time used, compared to CON-LS-1 and CON-LS-2.

Table 3.10 shows how the MIP heuristic performs when used on the best solution found by the construction heuristic, CON-LS-1, and CON-LS-2, respectively. The total computational time for each variant is equal to the computational time of CON-LS-3 for the same instance to make the results comparable. The results that give better objective values than CON-LS-3 are marked in bold in the "opt. gap" column. The tests show that in 12 out of 20 instances the ADPs obtained by using the CON-LS-2-MIP have a better objective value than those obtained using only the local search in the same amount of time. For CON-LS-1-MIP the result, 13 out of 20, is almost the same but the instances which are improved are different.

It is interesting to note that using an ADP with a better objective value as a starting point does not necessarily lead to a better ADP after using the MIP heuristic. This can be explained by the fact that since the computational time of CON-LS-1 is much shorter than CON-LS-2, CON-LS-1-MIP spends more time doing branch-and-bound and thus XpressMP is able to find better results. Another explanation is that even though the starting solution for the MIP heuristic is worse for CON-LS-1, it may have a bigger potential for improvement than the solution found by CON-LS-2. This suggests that one may get better results by using the MIP heuristic on several local minima solutions produced by the local

3.6 Computational results

Instance	CON-MIP		CON-LS-1-MIP		CON-LS-2-MIP	
	opt. gap	UB - LB	opt. gap	UB - LB	opt. gap	UB - LB
A1	192.44 %	45022	61.00 %	25923	51.91 %	23381
A2	16.07 %	16310	11.63 %	12270	19.41 %	19149
A3	166.14 %	29804	25.13 %	9590	37.70 %	13071
A4	42.74 %	22430	13.22 %	8746	9.35 %	6404
A5	19.27 %	13638	10.64 %	8114	10.52 %	8032
B1	47.20 %	8776	16.46 %	3869	18.92 %	4356
B2	11.24 %	5645	7.94 %	4109	7.82 %	4052
B3	19.28 %	3921	16.50 %	3435	16.64 %	3461
B4	33.23 %	6158	17.37 %	3652	17.73 %	3718
B5	27.13 %	4026	8.58 %	1491	11.65 %	1969
C1	2.03 %	29317	1.87 %	27005	1.84 %	26643
C2	2.98 %	30333	2.56 %	25993	2.56 %	25993
C3	3.74 %	30397	2.85 %	23209	2.82 %	22918
C4	4.09 %	22367	2.40 %	13115	2.40 %	13159
C5	2.18 %	8793	0.39 %	1578	0.39 %	1578
D1	3.67 %	31756	2.56 %	22155	1.59 %	13732
D2	1.82 %	11544	3.04 %	19310	2.99 %	18985
D3	5.40 %	26760	5.20 %	25767	5.20 %	25767
D4	1.06 %	3468	3.26 %	10717	0.81 %	2677
D5	1.10 %	2630	1.10 %	2630	1.10 %	2630
Average	30.14 %	17655	10.68 %	12634	11.17 %	12084

Table 3.10: MIP heuristic tests

search. However, such an approach is too time consuming for our purpose. Looking at the average gaps, starting from an ADP generated from the construction heuristic is inferior to starting from an ADP improved by local search.

Table 3.11 repeats the running times of the different versions of the CIH. The columns named “Xpress < CIH” gives the time it takes for XpressMP to obtain a better solution than the one found by the CIH by solving the mathematical model presented in Section 3.4. The XpressMP implementation of the full mathematical model was run for a maximum of 86400 s (24 h).

With the exception of instance D3 solved using CON-LS-3, all three CIH variants find better solutions than XpressMP in the same amount of time on all instances. Notice that we here compare the time spent by XpressMP to obtain a specific objective value to the total running time of the CIH algorithm, though the best solution by the heuristics might have been found in a much shorter time.

Even though the CIH outperforms XpressMP given equal running times, the time spent by XpressMP until a better solution is found is still relatively short for most of the 3 and 4 month instances. Here we should keep in mind that XpressMP will continue its tree search, and may find better solutions if given more time. On the other hand, the CIH terminates after the given number of seconds, and is unable to improve on its current best value. Therefore it might be better to use commercial optimization software on small instances if computational time is not an issue.

Instance	CON-LS-1		CON-LS-2		CON-LS-3	
	CPU(s)	Xpress<CIH	CPU(s)	Xpress<CIH	CPU(s)	Xpress<CIH
A1	108	>86400	418	>86400	1587	>86400
A2	45	>86400	171	>86400	642	>86400
A3	22	1930	83	1930	318	1930
A4	11	936	41	1806	157	1813
A5	6	168	23	168	87	180
B1	10	9605	40	9605	152	9605
B2	5	1418	20	1418	76	1418
B3	3	159	11	159	42	159
B4	1	33	6	33	23	76
B5	1	15	4	130	15	130
C1	124	>86400	474	>86400	1788	>86400
C2	50	>86400	192	>86400	729	>86400
C3	28	1944	109	1944	410	14348
C4	15	337	58	337	218	337
C5	15	215	55	215	206	215
D1	37	>86400	141	>86400	530	>86400
D2	15	>86400	60	>86400	226	>86400
D3	9	54	35	54	134	54
D4	4	65	17	223	66	223
D5	3	31309	11	31309	44	31309

Table 3.11: Comparison of computational times

For case B, XpressMP finds better solutions than the CIH for all instances, though for the B1 and B2 instances it spends a lot of time to do so. For the other 8 and 12 months instances, XpressMP cannot find better solutions than those provided by the CIH, even after 24 hours of computational time. This suggests that the CIH is particularly effective when the instances become large and where the inventories are tightly constrained (low I/P ratio).

Rakke et al. (2011) compare the results obtained in this paper with their rolling horizon heuristic (RHH). The comparison show that the CIH is much faster than the RHH, while the results comparing solution quality is mixed. For test cases A and B the RHH performs better (on average), while for test instances C and D the reverse is true. On average the solution provided by the CIH is more than 12 million USD better than the solution provided by the RHH for cases C and D, while it is only on average 3 million USD worse on cases A and B. Especially for the largest instances C1 and D1 the CIH solution is more than 60 and 20 million USD better, respectively. This further strenghtens our claim above that the CIH provides very good solutions when the instances are large and inventories are tightly constrained. It should also be noted that cases C and D are based on the future outlook of the producer, which may suggest that the CIH is better suited for use in a future DSS for the producer.

3.7 Concluding remarks

In this paper we have presented a solution method for a large scale ship routing and inventory management problem for a producer and distributor of LNG. The goal is to create an annual delivery program to fulfill the producer's long-term contracts at minimum cost, while maximizing the revenue from selling LNG in the spot market.

To solve this problem we have proposed a multi-start construction and improvement heuristic (CIH). The computational results show that the CIH creates high quality solutions to real-world instances in a short amount of time. Especially for large instances with tightly constrained inventories the CIH performs much better than other known approaches to the problem.

The short running time of the CIH makes it attractive as a component in a decision support system (DSS). Here users may want to test out different alternatives - e.g. analyze the impact of signing a new contract, chartering out a ship in the fleet, changing the production schedule, and so on. Another component that makes the CIH attractive in such a DSS is the fact that it contains no random components. According to the authors' experience, most users who are not familiar with operations research will be sceptical of using any DSS which give different results in two runs with the same input data.

If the time allowed for the CIH to produce an ADP was far greater, and the above mentioned concerns were ignored, one could potentially get better solutions by using the MIP-based improvement heuristic on several instances, and also added random components to the construction phase in a GRASP fashion.

An interesting topic for future research is disruption management. For instance if a ship breaks down, or if production has to be stopped. In such circumstances there is a need to create a new feasible ADP that minimizes both the inconvenience of the customers and the additional cost experienced by the producer.

Robustness is another major issue when the planning horizon spans over an entire year and looking at how an ADP could be created that minimizes the effect of possible disruptions is important. Having an ADP that is robust to unforeseen events may be just as important to the producer as having one that minimizes cost.

Bibliography

- Andersson, H., Christiansen, M., Fagerholt, K., 2010. Transportation planning and inventory management in the LNG supply chain. In: Pardalos, P. M., Bjørndal, E., Bjørndal, M., Rönnqvist, M. (Eds.), *Energy, Natural Resources and Environmental Economics. Energy Systems*. Springer Berlin Heidelberg, 427–439.
- Archetti, C., Bertazzi, L., Speranza, M. G., Hertz, A., 2010. A hybrid heuristic for an inventory-routing problem, to appear in *INFORMS Journal on Computing*.
- Atkinson, J., 1994. A greedy look-ahead heuristic for combinatorial optimization: An application to vehicle scheduling with time windows. *Journal of the Operational Research Society* 45 (6), 673–684.
- EIA, July 2008. *International energy outlook 2008*. Tech. rep., U.S. Department of Energy, Energy Information Administration U.S. Department of Energy Washington, DC 20585.
- Fearnley, 2008. *Fearnleys review 2007*. Tech. rep., Fearnresearch, Oslo, Norway.
- Grønhaug, R., Christiansen, M., 2008. Supply chain optimization for the liquefied natural gas business. In: Bertazzi, L., Nunen, J. v., Speranza, M. (Eds.), *Innovations in Distribution Logistics. Lecture Notes in Economics and Mathematical Systems*. Springer, 195–218.
- Grønhaug, R., Christiansen, M., Desaulniers, G., Desrosiers, J., 2010. A branch-and-price method for a liquefied natural gas inventory routing problem. *Transportation Science* 44 (3), 400–415.
- Halvorsen-Weare, E., Fagerholt, K., 2009. Routing and scheduling in a liquefied natural gas shipping problem with inventory and berth constraints. *Annals of Operations Research*, doi:10.1007/s10479-010-0794-y.
- Rakke, J., Stålhane, M., Moe, C., Andersson, H., Christiansen, M., Fagerholt, K., Norstad, I., 2011. A rolling horizon heuristic for creating a liquefied natural gas annual delivery program. *Transportation Research Part C*, 19(5), 896–911.
- Ronen, D., 1986. Short-term scheduling of vessels for shipping bulk or semi-bulk commodities originating in a single area. *Operations Research* 34 (1), 164–173.

Paper III

H. Andersson, M. Christiansen, G. Desaulniers and J. G. Rakke:

Creating Annual Delivery Programs of Liquefied Natural Gas

*submitted as an invited chapter to Optimization and Analytics in
the Oil and Gas Industry*

Creating Annual Delivery Programs of Liquefied Natural Gas

Abstract:

The annual delivery program (ADP) setup is an important problem in liquefied natural gas (LNG) supply chain planning. The ADP for an actor in the chain is the complete sailing schedule of the ships in the fleet for the coming year. In this chapter we focus on the ADP planning problem for one of the world's largest producers of LNG. The producer is responsible for the LNG inventories at the liquefaction plant, the loading port with a limited number of berths, and the routing and scheduling of a heterogeneous fleet of LNG ships. In addition, the producer has to fulfill a set of long-term contracts to customers all around the world. The objective is to design an ADP to fulfill the long-term contracts at minimum cost, while maximizing revenue from selling LNG in the spot market. A mixed integer programming (MIP) formulation of the ADP planning problem is presented and solved with a branch-and-cut algorithm. Several types of valid inequalities are developed that allow us to reduce the linear programming gap of the MIP formulation. The computational study shows that the problem is very complex, but that the valid inequalities are effective.

4.1 Introduction

There are large reserves of natural gas worldwide. Several existing gas producers are increasing their production capacity and new reservoirs are explored. The major projected increase in natural gas production is expected to occur in non-OECD regions, with the largest increments coming from the Middle East, Africa and non-OECD Europe and Euroasia, including Russia and other former Soviet Republics, (EIA, 2011). The world liquefaction capacity is expected to double from 2008 to 2035. In this period, 20 % of the increase in world natural gas production is expected to come from Iran and Qatar.

Natural gas has traditionally been transported in pipelines to markets close to the production areas. However, in several of these areas there are no significant markets and a pipeline solution is not an economically viable alternative for longer distances, especially across oceans. Technological advances in the maritime industry have made ships a good alternative for long distance transportation of natural gas. The gas can be cooled down at atmospheric pressure to a temperature of -162°C (-260°F). At this temperature the natural gas reaches its liquid

state and turns into liquefied natural gas (LNG). This reduces the volume of the gas by a factor of 610, (EIA, 2011), which makes transportation and storage more efficient. In addition, LNG offers greater trade flexibility than pipeline transportation.

The huge expansion of capacity in liquefaction plants has resulted in a steady growth of the LNG shipping fleet capacity as well. In 2008, the record breaking 53 new built ships came out of the ship yards. This represented a 25 % increase in total capacity in a year of no trade growth due to the financial crisis. Together with the current order stock, forecasts predict that the LNG fleet is experiencing a growth of about 20 new built ships every year in the coming years. The capacity of the world fleet of LNG ships has reached about 52 million m³ and more than 350 LNG ships are in operation, (GIIGNL, 2010). The size of individual LNG carrier ships is also increasing. The new built Q-MAX carriers has a capacity of 260 000–266 000 m³, (EIA, 2011).

It is expected that the natural gas consumption will increase by 52 % from 2008 to 2035, (EIA, 2011). Although the global recession resulted in a decline in natural gas use in 2009, robust demand returned already in 2010. Natural gas continues to be the fuel of choice for many regions of the world in several sectors, partly because its relatively low carbon intensity compared with oil and coal. In the power sector, low capacity costs and fuel efficiency also favor natural gas.

Managing the LNG supply chain is an important task. The LNG supply chain, illustrated in Figure 4.7, begins with the natural gas being extracted and sent through pipelines to a nearby liquefaction plant. Impurities are removed from the gas, which is cooled down to its liquid state and stored in tanks built with full-containment walls and systems to keep the gas liquefied. The LNG is then transported by ships to its destination terminal. There, it is again stored until it is regasified and finally sent through pipelines to the end users, which can be power plants, industrial customers and households. Figure 4.7 highlights the parts of the supply chain we focus on in this chapter. It is common to distinguish




Figure 4.7: The LNG supply chain

between three planning levels with different time horizons when planning the supply chain, (Stremersch et al., 2008). Long-term planning typically includes decisions about investments and long-term contracts that will have an impact many years ahead. The annual delivery program (ADP) setup belongs to the

next planning level and represents a tactical planning problem with a typical planning horizon of 12–18 months. When creating an ADP, the aim is to determine an optimal fleet schedule, including the delivery dates (or time windows) at the different customers' terminals. This fleet schedule must also satisfy inventory and berth constraints, as well as contract constraints. Finally, the operational planning deals with updating fleet schedules because of various logistics, economic or contractual reasons. Examples of logistics reasons can be rescheduling due to unplanned events, such as equipment breakdown or ship delays. An example of an economic reason is when spot market prices change; new sales or purchase opportunities may create needs for rescheduling. The typical length of the operational planning horizon is up to three months.

In this chapter, we focus on the creation of an ADP for a large-scale producer and distributor of LNG. The LNG supply chain has become more complex in recent years, partly due to the increased volumes and higher number of supply sources and demand regions. Traditionally, most LNG ships have been tied up to specific long-term contracts (defined as contracts with duration of more than 4 years), where they would be dedicated to sail between given liquefaction plants and regasification terminals. Lately, this situation has changed. There is a growing spot market for LNG, which creates more business opportunities serving different markets and introduces more flexibility in utilizing the LNG ships. For instance, spot and short-term imports recorded a very strong increase (40 % and 727 cargoes) in 2010 compared with 16 % in 2009. In comparison, LNG traded under long-term contracts recorded a 17 % increase in 2010, (GIIGNL, 2010). Moreover, this also contributes to making the supply chain planning much more challenging.

The ADP planning problem considered in this chapter is a combined LNG ship routing and inventory management problem. The producer has a single liquefaction plant with storage tanks. Two types of LNG are produced at the liquefaction plant and the quantity of LNG in the storage tanks has to stay within its limits. Connected to the liquefaction plant is a loading port. The number of ships that can load simultaneously is restricted by the number of berths. From the loading port, the LNG is shipped to customers world-wide, see Figure 4.8. The producer has to fulfill a set of long-term contracts that either outline a monthly demand for a given regasification terminal, or state that a certain quantity of LNG is to be delivered fairly evenly spread throughout the year. Over- and under-deliveries are accepted, but penalized. In addition to this, the producer has the opportunity to sell LNG in the spot market using short-term contracts. The producer operates a large fixed fleet of heterogeneous LNG ships and has the option to charter in additional ships in peak periods. All shipments are full shiploads. To plan for the coming years, the producer creates an ADP which is a list of scheduled voyages. Each scheduled voyage includes information




Figure 4.8: A map showing the loading port and possible customer locations

about the ship chosen for the voyage, the day of loading at the loading port and the contract served. The objective is to create an ADP that abides by the long-term contractual agreements at lowest possible cost, while maximizing the expected revenue from spot contracts. Henceforth, this is referred to as the ADP planning problem.

The purpose of this chapter is twofold: 1) Present the LNG supply chain and a case study from a large-scale producer and distributor of LNG, where the objective is to create an ADP, and 2) Describe a mathematical formulation and valid inequalities for the ADP planning problem that is solved by a branch-and-cut algorithm. This problem is also studied in Rakke et al. (2011) and Stålhane et al. (2012) where it is solved by a rolling horizon heuristic and a construction and improvement heuristic, respectively. The research in this chapter complements the previous research for this important problem. The algorithm can act as a basis for other solution approaches like a rolling horizon heuristic and branch-cut-and-price methods for solving real large-scale instances of the problem.

The rest of the chapter is organized as follows: A brief overview of the recent literature within LNG supply chain optimization is presented in Section 3. Section 4 is devoted to a detailed description of the ADP planning problem, while the mixed integer programming formulation is given in Section 5. Valid inequalities are described in Section 6. Section 7 presents the real-world cases tested and the associated computational results. Finally, concluding remarks follow in Section 8.

4.2 Literature

In this section, we limit ourselves to present literature related to the optimization of the LNG supply chain.

Combined routing and inventory management within maritime transportation has only been considered in the literature during the last one and a half decades. A survey on maritime inventory routing problems is presented in Christiansen and Fagerholt (2009), while Christiansen et al. (2007) give a comprehensive review within maritime transportation in general.

The liquefied natural gas inventory routing problem (LNG-IRP) was introduced by Grønhaug and Christiansen (2009) and is an important problem within LNG supply chain optimization. The problem includes decisions about production and sales quantities, inventory management at both liquefaction plants and regasification terminals, and routing and scheduling of a fleet of heterogeneous ships. The hold of the LNG ships is separated into several cargo tanks. The ships are always fully loaded, but partial unloading is allowed. The need to consider boil-off adds extra complexity to the problem. The objective of the LNG-IRP is to maximize the profit by designing routes and schedules for the fleet, including determining the production and sales quantities at all plants and terminals, without exceeding the ship capacities or the inventory limits of the storages. The LNG-IRP is studied by Grønhaug and Christiansen (2009) and Grønhaug et al. (2010). Grønhaug and Christiansen (2009) propose two formulations for the problem; one arc-based formulation and one path-based. A priori generation of all feasible paths is used to solve the path-based formulation. The problem is further analyzed in Grønhaug et al. (2010) where a branch-price-and-cut algorithm is developed for the path-based formulation.

Fodstad et al. (2010) and Uggen et al. (2011) study a richer version of the LNG-IRP which also addresses contract management and trading in a spot market among other things. An arc-based formulation is developed and solved by Fodstad et al. (2010), while Uggen et al. (2011) propose a heuristic method based on fix-and-relax time decomposition for the same formulation.

There are some differences between the LNG-IRP as defined by Grønhaug and Christiansen (2009) and the ADP planning problem. The LNG-IRP has a more complex structure than the ADP planning problem; it involves more than one production port, inventory constraints both at the production and consumption ports, and partial unloading. The ADP planning problem on the other hand includes multiple types of LNG and is considerably larger with respect to the size of the fleet and the length of the planning horizon.

The ADP planning problem presented here is also studied in several other papers. Rakke et al. (2011) solve the problem using a rolling horizon heuristic (RHH). The planning horizon is partitioned into shorter time intervals and the

RHH iteratively solves subproblems over these time intervals. A construction and improvement heuristic for the problem is developed by Stålhane et al. (2012). Initial ADPs are built in a construction phase and then improved using two different heuristics, one local search and one based on a reduced MIP formulation, in the improvement phase. The construction and improvement phases are then repeated many times for different values of the parameters controlling the construction. A similar MIP-based improvement heuristic is also used to improve the ADPs produced by the RHH. A new formulation using delivery patterns is proposed by Rakke et al. (2012). A delivery pattern states the number of deliveries by each ship to a given contract within a time interval. A branch-price-and-cut algorithm is used to solve the problem.

Halvorsen-Weare and Fagerholt (2009) study a similar problem where cargoes with defined time windows are generated a priori for each long-term contract, and the fleet of ships can be divided into disjoint groups. This study does not include a spot market, as opposed to Rakke et al. (2011), Stålhane et al. (2012), Rakke et al. (2012), and this chapter. Halvorsen-Weare and Fagerholt (2009) decompose the problem using Dantzig-Wolfe decomposition and handle inventory, berth, and scheduling decisions in the master problem, while routing decisions are dealt with in the subproblem. The master problem is a feasibility problem minimizing penalties for time window violations and is solved using branch-and-bound, while the subproblems are minimizing the port and transportation costs and are solved using either a local search heuristic or branch-and-bound.

The most general study of the LNG supply chain including some of its main characteristics is presented in Andersson et al. (2009). They consider two problems, reflecting different actors in the LNG supply chain. Mathematical models for each problem are presented and solution approaches for both problems are discussed. One of the problems is very similar to the problem considered in this chapter, but the problem there is simplified with respect to the number of LNG types and the possibility of selling LNG in the spot market. The other problem is for a vertically integrated company. Here, inventory management and sales of LNG at the regasification terminals are also considered.

4.3 Problem description

The ADP planning problem includes inventory and port management at the producer, routing and scheduling of a fleet of ships, and contract management between the producer and its customers.

4.3.1 Inventory and port management at the liquefaction plant

We consider a large-scale producer of LNG with liquefaction facilities and storages in one production area. Two types of LNG are being produced; rich LNG (RLNG) and lean LNG (LLNG). The production rates of LNG are always such that the equipment available for gas refinement is 100 % utilized. Thus, fluctuations in production rates occur only because of planned maintenance or unforeseen events such as breakdowns or strikes. This means that the production rates are given as parameters, but may vary during a year. Here, stochasticity is not considered and the production rates are assumed to be known at the time of planning. The different types of LNG are stored in separate tanks with given upper and lower limits. From the tanks the LNG is transported in pipelines to the docking facilities, where a limited number of berths for loading LLNG and RLNG are available. We estimate the total time between arrival and departure at the loading port to 24 hours. It includes the port arrival, docking and loading operation. The inventory and port management consists of ensuring that the quantity of LNG in the storage tanks is within its limits and that the berth capacity is not violated.

4.3.2 Routing and scheduling of LNG ships

The producer operates a heterogeneous fleet of ships. These ships are either owned by the producer or by one, or a group of customers. However, the producer is responsible for the routing and scheduling of all ships.

Several factors influence the availability of the ships in the fleet. Since there is no depot, some of the ships may be en route at the start of the planning horizon and become available for operation when they return to the loading port. The ships may also be unavailable because of certain pre-allocated activities, e.g. each ship is required to dry-dock for maintenance for a certain number of days every fifth year. The maintenance is not fixed to certain dates, but is required to be carried out within a given time interval. Maintenance is always scheduled to a dry-dock along the sailing lane, thus minimizing the time the ship is out of commission. After the maintenance has been performed, the ship has to go through a purge and cool-down procedure at a berth in the loading port before starting its next voyage. This procedure takes on average 24 hours, but depends on the ship type and the initial temperature of the empty tank.

A shipload must contain either LLNG or RLNG, but a ship can carry different types of LNG on consecutive voyages without any intermediate preparations. Though it is technically possible for a ship to sail between the loading port and a regasification terminal with only some of its tanks filled, this is not considered economically feasible by the producer and is never done in practice. Due to this,

and in order to avoid sloshing during transportation, the ship tanks are always filled to their capacities at the loading port. The duration of the voyage influences the amount of LNG delivered to the customer because of boil-off, which causes a fixed ship-dependent percentage of the total capacity of a tank to evaporate each day. To keep the LNG tanks on board a ship cool during the return leg, a certain amount of LNG must be kept in the tanks to keep the boil-off process going until the ship returns to the loading port and the loading operation can begin. Consequently, the amount of LNG loaded at the loading port and the amount delivered to the customer will differ, depending on the time spent in transit. If boil-off causes the ship tanks to be completely empty, a 24-hour cool-down period is required at the loading port before the next loading operation can take place. However, the producer assumes that a ship never will be completely empty, which means that ships never have to go through a cool-down process, except before the first voyage following a maintenance period.

The travel times depend on the ship, because every ship has an individual cruising speed. The duration of a voyage may depend on the time of year, as sailing conditions vary between the summer and the winter months. Due to ocean currents the travel time may also vary between the outbound and return leg of a voyage. Not all ships can visit all regasification terminals. This is due to vessel acceptance policies at the ports, and the fact that some ships are owned by one, or a group of customers, limiting them to only visit their owners' regasification terminals.

The transportation costs consist of several components. The fixed costs are the time charter rates, while the variable costs are port and canal fees that are determined by the ship type and the contract served. A third variable component is the bunker oil cost, which is dependent on the ship size, the load on board and the duration of the voyage. Since the fixed costs cannot be changed during the time horizon, the cost of sailing a scheduled voyage is assumed to be dependent on the capacity class of the ship, the duration of the voyage and the regasification terminal visited.

If the producer does not have enough ships available at a given time, additional ships may be chartered in for one-off deliveries. A daily charter rate defines the cost of these ships.

4.3.3 Contract management: Long-term contracts and sale of spot cargoes

The producer has a set of long-term contracts by which it is obliged to deliver a certain amount of LNG to specified regasification terminals each year. These contracts have time frames of 20–30 years, so the total amount of LNG the producer will deliver in a given year is known well in advance. A contract also specifies

the regasification terminal to receive the LNG, but a regasification terminal may have more than one contract associated with it. A long-term contract either outlines the monthly demand that is to be delivered to the customer's regasification terminals, or simply states that the LNG is to be delivered fairly evenly spread throughout the year. Due to the contracts being long-term, there is some flexibility in the volumes that have to be delivered in a given year. The current "rule of thumb" is that it is acceptable for the delivered amount of LNG to deviate with up to 10 % of the contracted volume, as long as this is evened out the following year. However, the producer always aims to deliver as close as possible to the annual demands.

In addition to serving the long-term contracts, the producer has the opportunity to sell LNG in the spot market. The spot sales are short-term contracts for one-off deliveries from the producer to buyers on agreed terms, and represent possibilities for increased revenue. To maximize the expected contribution margin from this option while minimizing the total cost of fulfilling long-term contracts, spot sales are incorporated in the model.

4.4 Model

The ADP planning problem described in Section 4.3 can be formulated as a mixed integer linear program. We define G , C , V , T , I to be the sets of LNG types, contracts, ships, time periods, and time intervals respectively. Contract is here a generic term representing not only the deliveries of LNG but also a period of maintenance. The set of contracts can be divided into disjoint sets C_g , representing contracts for LNG type g . In each C_g , one contract, c^M , is a maintenance contract and one contract, c^S , is a spot contract; the sets of maintenance and spot contracts are denoted C^M and C^S respectively. The maintenance contracts are needed to model the purge and cool-down procedure which takes place at one of the berths in the loading port. All other contracts are associated with long-term contracts and the set of these contracts is denoted C^{LT} . The set of ships can be divided into two disjoint subsets; V^P denotes the set of ships operated by the producer and V^S denotes the set of spot ships. The set of ships that must undergo maintenance during the planning horizon is denoted V^M . We also introduce V_c as the set of ships that can serve contract c .

There are two types of LNG produced and stored at the liquefaction plant. The production rate on day t of LNG type g is denoted P_{gt} and the time-dependent upper and lower limits of the storage tanks are denoted \overline{Q}_{gt} and \underline{Q}_{gt} . The berth capacity B_g restricts the number of ships that can load LNG type g or go through a purge and cool-down procedure after maintenance using LNG type g in a given time period.

The capacity of ship v is denoted L_v , and the cost and duration of a sailing on contract c is denoted C_{cv}^T and T_{cv} , respectively. The set of time periods where ship v is available for service is denoted \mathbb{T}_v and the set of time periods where it can start maintenance is denoted \mathbb{T}_v^M .

The time horizon is divided into time intervals. These intervals cover the horizon and may be overlapping. The set of time periods in time interval i is denoted \mathbb{T}_i^I and the demand of contract c in time interval i is denoted D_{ci} . The demand does not need to be met exactly, but there are penalty costs associated with over-delivery, C_{ci}^{D+} , and under-delivery, C_{ci}^{D-} . The revenue from sending one unit of LNG to spot contract c is denoted R_c^S and the revenue from having one unit of LNG type g left in the storage tanks at the end of the planning horizon is denoted R_g^I . The end of the planning horizon is denoted \bar{T} .

The binary variable x_{cvt} is 1 if ship v starts sailing on contract c in time period t , and 0 otherwise. The outgoing inventory level of LNG type g at the liquefaction plant in time period t is denoted q_{gt} . Over- and under-delivery to contract c in time interval i is given by y_{ci}^+ and y_{ci}^- , respectively. Using this notation, the ADP planning problem can be formulated as:

$$\begin{aligned} \min z = & \sum_{c \in \mathbb{C}} \sum_{v \in \mathbb{V}_c} \sum_{t \in \mathbb{T}_v} C_{cv}^T x_{cvt} + \sum_{c \in \mathbb{C}^{LT}} \sum_{i \in \mathbb{I}} (C_{ci}^{D+} y_{ci}^+ + C_{ci}^{D-} y_{ci}^-) \\ & - \sum_{c \in \mathbb{C}^S} \sum_{v \in \mathbb{V}_c} \sum_{t \in \mathbb{T}_v} R_c^S L_v x_{cvt} - \sum_{g \in \mathbb{G}} R_g^I q_{g\bar{T}} \end{aligned} \quad (4.29)$$

$$\sum_{c \in \mathbb{C}_g \setminus \mathbb{C}^M} \sum_{\substack{v \in \mathbb{V}_c \\ t \in \mathbb{T}_v}} x_{cvt} + \sum_{c \in \mathbb{C}_g \cap \mathbb{C}^M} \sum_{\substack{v \in \mathbb{V}_c^M \\ t - (T_{cv} - 1) \in \mathbb{T}_v}} x_{cv, t - (T_{cv} - 1)} \leq B_g \quad g \in \mathbb{G}, t \in \mathbb{T}, \quad (4.30)$$

$$q_{g, t-1} + P_{gt} - q_{gt} - \sum_{c \in \mathbb{C}_g \setminus \mathbb{C}^M} \sum_{v \in \mathbb{V}_c} L_v x_{cvt} = 0 \quad g \in \mathbb{G}, t \in \mathbb{T}, \quad (4.31)$$

$$\sum_{c \in \mathbb{C}} \sum_{\substack{\tau \in \mathbb{T}_v \\ t - T_{cv} < \tau \leq t}} x_{cv\tau} \leq 1 \quad v \in \mathbb{V}^P, t \in \mathbb{T}_v, \quad (4.32)$$

$$\sum_{c \in \mathbb{C}^M} \sum_{t \in \mathbb{T}_v^M} x_{cvt} = 1 \quad v \in \mathbb{V}^M, \quad (4.33)$$

$$\sum_{v \in V_c} \sum_{t \in T_i^l \cap T_v} L_v x_{cvt} + y_{ci}^- - y_{ci}^+ = D_{ci} \quad c \in C^{LT}, i \in I, \quad (4.34)$$

$$\underline{Q}_{gt} \leq q_{gt} \leq \overline{Q}_{gt} \quad g \in G, t \in T, \quad (4.35)$$

$$x_{cvt} \in \{0, 1\} \quad c \in C, v \in V_c, t \in T_v, \quad (4.36)$$

$$y_{ci}^+, y_{ci}^- \geq 0 \quad c \in C^{LT}, i \in I. \quad (4.37)$$

The objective function (4.29) minimizes the transportation and penalty costs minus the revenue from LNG sold in the spot market and the LNG left in the tanks at the end of the planning horizon. Constraints (4.30) state the berth capacities. After maintenance, a ship goes through a purge and cool-down procedure at one of the berths, and because of this, the variables representing a sailing on a maintenance contract are added to the constraints on the last day of the sailing. Inventory balance at the liquefaction plant is given by constraints (4.31). Constraints (4.32) ensure that each ship serves at most one contract each day, and constraints (4.33) state that the ships that are scheduled for maintenance undergo maintenance. The demand is handled in constraints (4.34) where the delivered quantity is balanced with the demand using the variables for over- and under-delivery. The inventory limits are given by constraints (4.35) and the variable restrictions by constraints (4.36) and (4.37).

4.5 Strengthening the formulation

The LP-relaxation of the formulation presented in Section 4.4 is weak and does not provide a good bound. In this section we present valid inequalities that can be added to strengthen the formulation.

4.5.1 Delivery inequalities

The penalty costs for over- and under-deliveries are a substantial part of the total objective value. In a fractional solution, these costs can be avoided since it is always possible to deliver exactly the demand by using fractional ship loads. In Figure 4.9a, the penalty function of a small example with one contract and two ships is shown. The demand of the contract is 18 and the ships have capacities of 10 and 15 respectively. The penalty for over- and under-delivery is 15 per unit. The dots show the feasible integer solutions; no delivery, one delivery with the

small ship, one delivery with the large ship, two deliveries with the small ship and so on.




Figure 4.9: Illustration of the penalty cost as a function of delivered volume

It is clear that the penalty cost can be zero in a fractional solution; using the small ship 1.8 times is one possible solution. If this possibility could be cut away, we would have a strengthened formulation. Figure 4.9b illustrates the idea behind the delivery inequalities; forcing the penalty cost to be above the red line cuts away the zero cost solution. The red line can be represented as a valid inequality connecting the delivered quantity and the penalty cost. If we define (Q_{ci}^-, C_{ci}^-) as the maximum quantity that can be delivered to contract c in time interval i given that we under-deliver and the corresponding penalty cost, and (Q_{ci}^+, C_{ci}^+) as the minimum quantity that can be delivered given that we over-deliver and the corresponding penalty cost, the following inequalities are valid:

$$c_{ci}^D \geq \frac{(C_{ci}^+ - C_{ci}^-)}{(Q_{ci}^+ - Q_{ci}^-)} \cdot \left(\sum_{t \in T_i^I} L_v x_{cvt} - Q_{ci}^- \right) + C_{ci}^- \quad c \in \mathcal{C}^{LT}, i \in \mathcal{I}, \quad (4.38)$$

where c_{ci}^D is the penalty cost for contract c in time interval i . To incorporate this in the model, the following constraints must also be defined:

$$c_{ci}^D \geq C_{ci}^{D-} y_{ci}^- \quad c \in \mathcal{C}^{LT}, i \in \mathcal{I}, \quad (4.39)$$

$$c_{ci}^D \geq C_{ci}^{D+} y_{ci}^+ \quad c \in \mathcal{C}^{LT}, i \in \mathcal{I}. \quad (4.40)$$

The penalty costs in the objective function, $(C_{ci}^{D+} y_{ci}^+ + C_{ci}^{D-} y_{ci}^-)$, are then replaced with c_{ci}^D .

4.5.2 Symmetry breaking inequalities

Identical ships give rise to solutions that are mathematically different but similar from a practical point of view; with five identical ships and two sailings, there are 20 mathematically different solutions that are the same from a practical point of view. This causes symmetry and to reduce this symmetry, the following inequalities are derived.

Let V^σ be a set of identical ships and let \bar{V}^σ be the last of these ships. Two ships are identical if they have the same capacity, cruising speed and cost structure, and if they can serve the same contracts. Order the ships in V^σ in a non-decreasing order with respect to the day when the ships become available for operation. The symmetry between the ships can be reduced by adding

$$\sum_{c \in C} \sum_{t \in T} (T_{cj} x_{cjt} - T_{c,j+1} x_{c,j+1,t}) \geq 0 \quad j \in V^\sigma \setminus \bar{V}^\sigma, \quad (4.41)$$

where j is the j^{th} ship in V^σ . The inequalities force ship j to be used at least as many days as ship $j + 1$.

4.5.3 Loading inequalities

For any given time interval $T' = \{\underline{T}', \dots, \bar{T}'\}$, the maximum and minimum number of loading operations of LNG type g can be calculated from the incoming inventory level $q_{g,\underline{T}'-1}$ and the total production $P' = \sum_{t \in T'} P_{gt}$. Start by ordering all ships that can load LNG of type g in T' in a non-decreasing order with respect to capacity. From this order, two sequences of loading quantities can be created by taken into account the number of times each ship can load within T' . Assuming that there are three ships, v_1 , v_2 , and v_3 , with capacity 8, 11, and 12, respectively. Within T' , v_1 can load LNG type g at most 2 times, v_2 can load at most 1 time and v_3 at most 3 times. The minimum quantity that can be loaded given one loading operation is thus 8, i.e. using v_1 . Given two loading operations, the minimum loading quantity is 16, i.e. using v_1 twice. Given three loading operations, the minimum loading quantity is 27, i.e. using v_1 twice and v_2 once and so on. These minimum loading quantities form a sequence of minimum loading quantities $\underline{S} = (8, 16, 27, 39, 51, 63)$. Correspondingly, we can define a sequence of maximum loading quantities $\bar{S} = (12, 24, 36, 47, 55, 63)$ by reversing the order of the ships.

Disregarding the bounds on the incoming inventory we see that the incoming inventory level must be at least $q_{g,\underline{T}'-1} = Q_{g,\bar{T}'} - P'$ to be feasible with respect to the lower bound on the outgoing inventory. Thus, to allow i loading operations, $q_{g,\underline{T}'-1}$ must be at least $q_{g,\underline{T}'-1} + s_i$, where s_i is the i^{th} element in \underline{S} . Hence, the number of loading operations is bounded from above by $q_{g,\underline{T}'-1}$ and \underline{S} . Similar,

the lowest quantity at which we have to start loading is when $\bar{q}_{g,\underline{T}'-1} = \bar{Q}_{g,\bar{T}'} - P'$, which means that at least j loading operations are needed when $q_{g,\underline{T}'-1} > \bar{q}_{g,\underline{T}'-1} + \bar{s}_{j-1}$, where \bar{s}_j is the j^{th} element in \bar{S} and $\bar{s}_0 = 0$. The number of loading operations is thus bounded from below by $\bar{q}_{g,\underline{T}'-1}$ and \bar{S} .

Figure 4.10a shows an example of how the number of loadings can be bounded using the sequences \underline{S} and \bar{S} where $P' = 30$, $\underline{Q}_{g,\bar{T}'} = 5$ and $\bar{Q}_{g,\bar{T}'} = 35$. The red dotted line shows the maximum number of loadings that can be performed, while the black dotted line shows the minimum number of loadings that must be performed. The thick black lines define the feasible integer solutions.

We can strengthen this by including the upper and lower bounds on the incoming inventory. In Figure 4.10b, these bounds are included; $\underline{Q}_{g,\underline{T}'} = 0$ and $\bar{Q}_{g,\underline{T}'} = 25$. The feasible region of the LP relaxation can be further reduced by cutting away non-integer extreme points, this is illustrated in Figure 4.10c. Finally, the remaining feasible region of the LP relaxation and the feasible integer solutions are shown in Figure 4.10d. Each facet in Figure 4.10d can be expressed as a linear constraint containing the number of loadings and the incoming inventory level.

A slightly different version of this cut is proposed by Engineer et al. (2009) and later strengthened by Rakke et al. (2012). We have used the same ideas as in Rakke et al. (2012) but present it differently. See Rakke et al. (2012) for a more detailed description of the cuts.

4.5.4 Timing inequalities

A common phenomenon in fractional solutions is premature loadings, which means that fractions of ships load LNG before enough gas for a full ship load is available. Engineer et al. (2009) developed a cut to remedy this problem. Their cut relates the production rate and incoming inventory level to a time-weighted sum of loading days. The cut was later used in Rakke et al. (2012), and we have adopted this idea.

Start with a given time interval $T' = \{\underline{T}', \dots, \bar{T}'\}$ and the sequence of maximum loading quantities, \bar{S} , defined in Section 4.5.3. Let $l_{T'}$ be the total quantity loaded during T' and let K define the least number of ships needed to load $l_{T'}$, i.e. $\bar{s}_{K-1} < l_{T'} \leq \bar{s}_K$ where \bar{s}_i is the i^{th} element in \bar{S} . The earliest time for the k^{th} loading, \underline{t}_k , is such that

$$q_{\underline{T}'-1} + \sum_{t=\underline{T}'}^{\underline{t}_k-1} P_t < \underline{s}_k \leq q_{\underline{T}'-1} + \sum_{t=\underline{T}'}^{\underline{t}_k} P_t \quad (4.42)$$

where \underline{s}_k is the k^{th} element in \underline{S} defined in Section 4.5.3. For the sake of simplicity

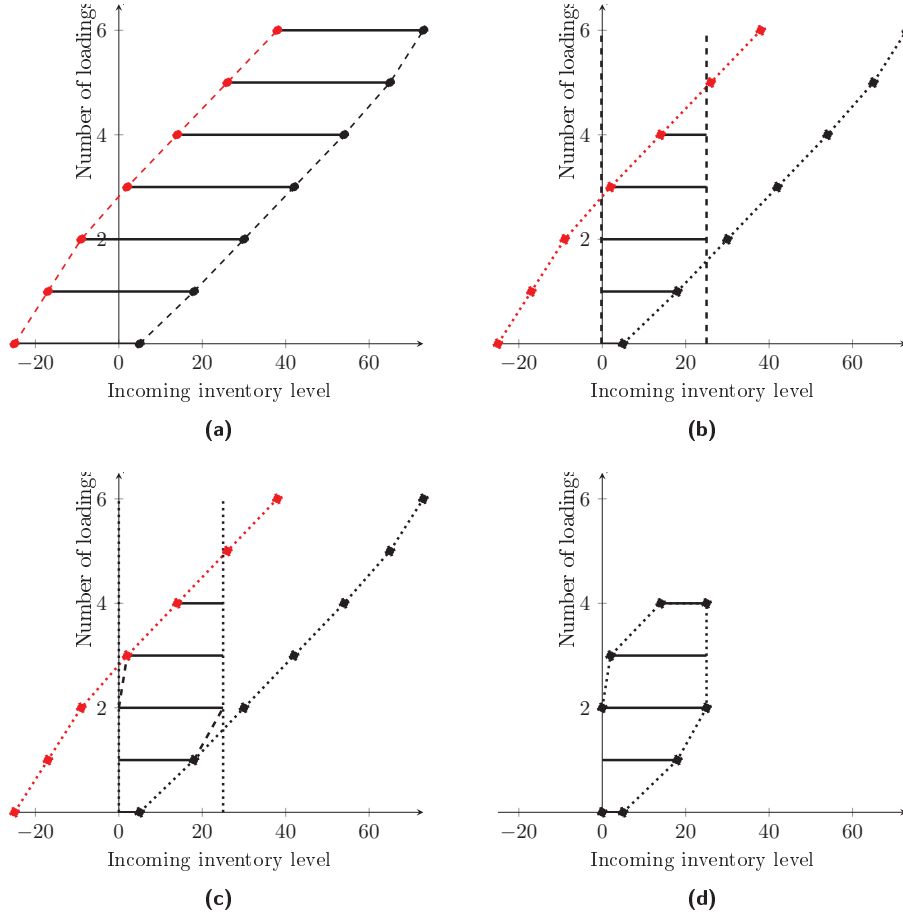


Figure 4.10: Illustration of the relation between the incoming inventory and the number of loadings

and without loss of generality, we have removed the index representing LNG type in this section. The least number of days, counting from $\underline{T}' - 1$, for the K loadings is thus $\sum_{k=1}^K (t_k - \underline{T}' + 1)$ and this can be used to bound the time-weighted sum of loadings. To estimate t_k , we introduce a production constant P instead of P_t . With P we can reformulate the second inequality of (4.42) into

$$\underline{s}_k \leq q_{\underline{T}'-1} + (t_k - \underline{T}' + 1)P \Rightarrow t_k \geq \frac{\underline{s}_k - q_{\underline{T}'-1}}{P} + \underline{T}' - 1 \quad k \in \{1, \dots, K\} \quad (4.43)$$

A natural way to estimate P is $P = \max_{t \in T'} P_t$, but a better estimate can be found by using the average production rate over a subinterval of T' starting from \underline{T}' . With this, we get the following estimate of P

$$P = \max_{\tau \in T'} \frac{\sum_{t=\underline{T}'}^{\tau} P_t}{\tau - \underline{T}' + 1} \quad (4.44)$$

The bounding inequality

$$\sum_{c \in C \setminus C^M} \sum_{v \in V_c} \sum_{t \in T'} (t - \underline{T}' + 1) x_{cvt} \geq \sum_{k=1}^K (t_k - \underline{T}' + 1) \quad (4.45)$$

relating the time-weighted sum of the loadings and the least number of loading days can now be stated as a valid inequality by substituting \underline{t}_k using (4.43). This gives

$$\sum_{c \in C \setminus C^M} \sum_{v \in V_c} \sum_{t \in T'} (t - \underline{T}' + 1) x_{cvt} \geq \sum_{k=1}^K \left(\frac{s_k - q_{\underline{T}'} - 1}{P} \right) \quad (4.46)$$

where P is given by (4.44). For a different and more detailed description of the cut, the reader is referred to Rakke et al. (2012).

4.6 Computational study

The branch-and-cut algorithm for the ADP planning problem is implemented using Xpress-MP 7.2 and Mosel 3.2.2. The cut separation was written in C++ and called as a module from Mosel. All tests were done on a HP dl165 G5 with a 2x 2.4GHz AMD Opteron 2431 and 24Gb of RAM.

The delivery inequalities, presented in Section 4.5.1, and the symmetry inequalities, presented in Section 4.5.2, are added to the formulation a priori. On the other hand, the loading and timing inequalities, presented in Sections 4.5.3 and 4.5.4 respectively, are implemented as cuts. The separation routines are run in every node of the branch-and-bound tree and if a violated cut is found, it is included and the node is reoptimized. Note that the loading and timing inequalities are implemented somewhat differently from the descriptions and that some algorithmic considerations have been made about which cuts to search for in the separation routines. This has a minor effect on the results.

4.6.1 Instances

We have tested the branch-and-cut algorithm on the instances presented in Rakke et al. (2011) and Stålhane et al. (2012). The test set consists of four groups of

instances. Groups A and C represent the situations in 2008 and 2012, respectively. Groups B and D are reduced versions of groups A and C where some contracts and ships have been removed. Within each group, four different instance are generated, having a planning horizon of 12, 8, 6, and 4 months respectively. The instances are denoted with the letter of the group and a number between 1 and 4 where 1 represents a 12-month planning horizon and 4 a 4-month planning horizon, i.e. B2 is the instance from group B with an 8-month planning horizon. This naming is the same as in Rakke et al. (2011) and Stålhane et al. (2012).

4.6.2 Results

The computational study focuses on two aspects, the improvement of the initial LP bounds and the improvement of the best bounds after 4 hours. As was noted in Rakke et al. (2011) and Stålhane et al. (2012), not only the relative gap, but also the absolute gap is interesting to study. The 2008 and 2012 cases are rather different; in the 2008 case there is a lot of excess LNG to sell, causing the revenue from selling LNG in the spot market to be approximately the same as the transportation and penalty costs giving a lower bound close to zero. This makes the relative differences large in these cases. The 2012 case is much more tightly constrained when it comes to excess LNG. Hence, the objective value is much closer to the sum of transportation and penalty costs, giving relatively smaller differences.

Table 4.12 shows the improvements of the initial LP bound for a small set of combinations of inequalities and cuts. Each column is marked with the section where the included inequalities are presented, the column denoted 'All' shows the results when all inequalities and cuts are included. The first set of results show the relative improvement of the objective value, in percentage, of including the inequalities compared with not include any inequalities. The second set shows the absolute difference between the objective values. It is clear that the delivery inequalities, presented in Section 4.5.1, are efficient and improve the objective value of the LP relaxation by 7.1 % on average. They are especially efficient on the 2008 reduced instances; fewer ships makes it harder to deliver the right quantity to the contracts and may thus create a larger deviation between the fractional and integer delivery quantities. The other inequalities are not as efficient, but still manage to improve the objective value of the LP relaxation on some instances.

The patterns from the LP values presented in Table 4.12 are repeated when we study the improvements of the best bound after four hours in Table 4.13. Here, not all combinations of inequalities and cuts improve the bound. Sometimes the separation routines take a long time, which means that the number of examined nodes in the branch-and-bound tree is reduced, something that leads to weaker

bounds.

4.7 Concluding remarks

This chapter presents a branch-and-cut algorithm and valid inequalities for the annual delivery program (ADP) planning problem. The ADP planning problem includes the inventory and port management at the producer, the routing and scheduling of a fleet of ships, and the contract management between the producer and its customers.

Four families of inequalities to strengthen the formulation are presented. The first family of inequalities cuts away the possibility to do fractional deliveries without incurring a penalty cost, while the second is symmetry breaking constraints. The third family of inequalities bounds the number of loadings in a given time interval by relating it to the incoming inventory. The fourth family handles the timing of loadings and prevents that fractions of ships load LNG before enough gas for a full ship load is available. The computational study shows that all valid inequalities strengthen the formulation.

Instance	Relative improvement					Absolute improvement				
	5.1	5.2	5.3	5.4	All	5.1	5.2	5.3	5.4	All
A1	7.1	0.0	0.0	0.1	7.1	7152	0	34	128	7186
A2	3.4	0.0	0.0	0.0	3.4	4841	0	0	47	4842
A3	5.2	0.0	0.0	0.1	5.4	3505	0	18	46	3605
A4	2.5	0.0	0.0	0.0	2.5	2221	0	21	3	2245
B1	26.6	0.0	0.0	0.3	26.9	13798	0	0	131	13928
B2	13.9	0.0	0.0	0.4	14.1	10350	0	0	275	10490
B3	16.1	0.0	0.0	0.6	16.5	6302	0	1	248	6434
B4	11.4	0.0	0.0	0.0	13.4	3848	0	0	0	4516
C1	4.6	0.0	0.0	0.0	4.6	60544	0	0	0	60544
C2	4.2	0.0	0.0	0.0	4.2	38140	0	0	0	38140
C3	4.1	0.0	0.0	0.5	4.1	29506	0	0	3783	29506
C4	4.1	0.0	0.0	1.1	5.2	19564	0	0	5322	24834
D1	2.9	0.0	0.0	0.0	2.9	23241	0	0	0	23241
D2	2.7	0.0	0.0	0.0	2.7	15367	0	0	0	15367
D3	2.3	0.0	0.0	0.0	2.5	10636	0	0	0	11258
D4	3.0	0.0	0.0	1.8	4.5	8590	0	0	5184	12910
Avg.	7.1	0.0	0.0	0.3	7.5	16100	0	5	948	16815

Table 4.12: Improvements of the LP optimal value

Instance	Relative improvement					Absolute improvement				
	5.1	5.2	5.3	5.4	All	5.1	5.2	5.3	5.4	All
A1	7.0	0.0	0.1	0.0	7.0	7015	1	66	17	7044
A2	3.3	0.0	0.0	0.0	3.3	4642	1	44	61	4677
A3	5.3	0.0	0.1	0.0	5.4	3567	0	50	5	3587
A4	2.3	0.0	0.0	0.0	2.3	2054	0	17	-17	2080
B1	27.5	0.0	-1.5	2.3	28.9	13872	0	-762	1178	14604
B2	9.5	0.0	0.3	-1.3	10.2	6145	0	171	-819	6605
B3	12.1	0.0	1.8	-3.6	11.9	3757	0	563	-1130	3682
B4	12.1	0.0	11.6	1.0	12.7	3494	0	3370	302	3672
C1	4.6	0.0	0.0	0.0	4.6	60567	0	0	10	60562
C2	4.2	0.0	0.0	0.0	4.2	38148	-16	-14	-1	38156
C3	3.5	0.0	0.0	0.0	4.3	25725	0	1	13	31222
C4	4.0	0.0	0.1	0.3	4.0	19533	1	420	1437	19568
D1	2.9	0.0	0.0	0.0	3.3	22927	303	-44	-91	26226
D2	2.7	0.0	0.0	0.0	3.7	15367	0	25	0	21478
D3	-0.1	-0.1	-0.6	-1.5	0.8	-480	-431	-2656	-7235	3620
D4	0.5	0.2	-0.4	-0.3	0.9	1550	560	-1119	-919	2722
Avg.	6.3	0.0	0.7	-0.2	6.7	14243	26	8	-449	15594

Table 4.13: Improvements of the best bound after four hours

The chapter also provides a detailed description of the ADP planning problem and an updated review of literature related to the optimization of the LNG supply chain.

Bibliography

- Andersson, H., Christiansen, M., Fagerholt, K., 2009. Transportation planning and inventory management in the LNG supply chain. In: Bjørndal, E., Bjørndal, M., Pardalos, P.M., Rönnqvist, M. (Eds.), *Energy, Natural Resources and Environmental Economics*. Springer, pp. 429–441.
- Christiansen, M., Fagerholt, K., 2009. Maritime inventory routing problems. In: Floudas, C. A., Pardalos, P. M. (Eds.), *Encyclopedia of Optimization*, 2nd Edition Springer, pp. 1947–1955.
- Christiansen, M., Fagerholt, K., Nygreen, B., Ronen, D., 2007. Maritime transportation. In: Barnhart, C., Laporte, G. (Eds.), *Transportation*. Vol. 14 of Handbooks in Operations Research and Management Science. Elsevier Science, pp. 189–284.
- EIA, 2011. International Energy Outlook 2011. Tech. rep., Energy Information Administration, U.S. Department of Energy, Washington, DC 20585.
- Engineer, F., Furman, K., Nemhauser, G., Savelsbergh, M., Song, J.-W., 2009. A branch-price-and-cut algorithm for a maritime inventory routing problem. DOMINANT workshop, Molde, Norway, 09.22.2009.
- Fodstad, M., Uggen, K. T., Rømo, F., Lium, A.-G., Stremersch, G., Hecq, S., 2011. LNGScheduler: a rich model for coordinating vessel routing, inventories and trade in the liquefied natural gas supply chain. *The Journal of Energy Markets*, **3**, 1–34.
- GIIGNL, 2010. The LNG Industry in 2010. The International Group of Liquefied Natural Gas Importers, Levallois-Perret, France. www.giignl.org/fr/home-page/lng-industry.
- Grønhaug, R., Christiansen, M., (2009). Supply chain optimization for the liquefied natural gas business. In: Bertazzi, L., Nunen, J. V., Speranza, M. (Eds.), *Innovations in Distribution Logistics*. Vol. 619 of Lecture Notes in Economics and Mathematical Systems. Springer, pp. 195–218.
- Grønhaug, R., Christiansen, M., Desaulniers, G., Desrosiers, J. (2010). A branch-and-price method for a liquefied natural gas inventory routing problem. *Transportation Science*, **44** (3), 400–415.

- Halvorsen-Weare, E., Fagerholt, K., (2012). Routing and scheduling in a liquefied natural gas shipping problem with inventory and berth constraints. *Annals of Operations Research*, DOI : 10.1007/s10479-010-0794-y.
- Rakke, J. G., Stålhane, M., Moe, C. R., Andersson, H., Christiansen, M., Fagerholt K., Norstad I. A rolling horizon heuristic for creating a liquefied natural gas annual delivery program. *Transportation Research C*, **19** (5), 896–911.
- Rakke, J. G., Christiansen, M., Andersson, H., Desaulniers, G. Branch-price-and-cut for creating an annual delivery program of multi-product liquefied natural gas. *Working paper*, Norwegian University of Science and Technology, 37 pages.
- Stremersch, G., Michalek, J., Hecq, S., (2008). Decision support software tools for LNG supply chain management. Gastech 2008.
- Stålhane, M., Rakke, J. G., Moe, C. R., Andersson, H., Christiansen, M., Fagerholt, K. A construction and improvement heuristic for a liquefied natural gas inventory routing problem. *Computers & Industrial Engineering*, **62** (1), 245–255, (2012).
- Uggen, K.T., Fodstad, M., Nørstebø, V. S. Using and extending fix-and-relax to solve maritime inventory routing problems. *TOP*, DOI : 10.1007/s11750-011-0174-z.

Paper IV

J. G. Rakke, H. Andersson, M. Christiansen and G. Desaulniers:

Branch-Price-and-Cut for Creating an Annual Delivery Program of Multi-Product Liquefied Natural Gas

Submitted to an international journal

Is not included due to copyright

Paper V

J. G. Rakke, M. Christiansen, K. Fagerholt and G. Laporte:

The Traveling Salesman Problem with Draft Limits

Computers & Operations Research, 39 (9), 2161–2167.

Is not included due to copyright

