Computers & Industrial Engineering 62 (2012) 245–255

Contents lists available at SciVerse ScienceDirect

Computers & Industrial Engineering

journal homepage: www.elsevier.com/locate/caie

A construction and improvement heuristic for a liquefied natural gas inventory routing problem

Magnus Stålhane [a,*], Jørgen Glomvik Rakke [a], Christian Rørholt Moe [a], Henrik Andersson [a], Marielle Christiansen [a], Kjetil Fagerholt [a,b]

[a] Norwegian University of Science and Technology, Department of Industrial Economics and Technology Management Technology Trondheim, Norway
[b] Norwegian Marine Technology Research Institute (MARINTEK), Trondheim, Norway

ARTICLE INFO

Article history:
Received 12 April 2011
Received in revised form 13 September 2011
Accepted 14 September 2011
Available online 24 September 2011

Keywords:
Maritime transportation
Inventory routing
Heuristics

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.

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.

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 worldwide to whom the producer has a contractual obligation to deliver LNG, see Fig. 1. 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.

* Corresponding author. Tel.: +47 73590811.
E-mail addresses: magnus.stalhane@iot.ntnu.no (M. Stålhane), jorgen.rakke@iot.ntnu.no (J.G. Rakke), chrism@stud.ntnu.no (C.R. Moe), henrik.andersson@iot.ntnu.no (H. Andersson), marielle.christiansen@iot.ntnu.no (M. Christiansen), kjetil.fagerholt@iot.ntnu.no (K. Fagerholt).

0360-8352/$ - see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.cie.2011.09.011

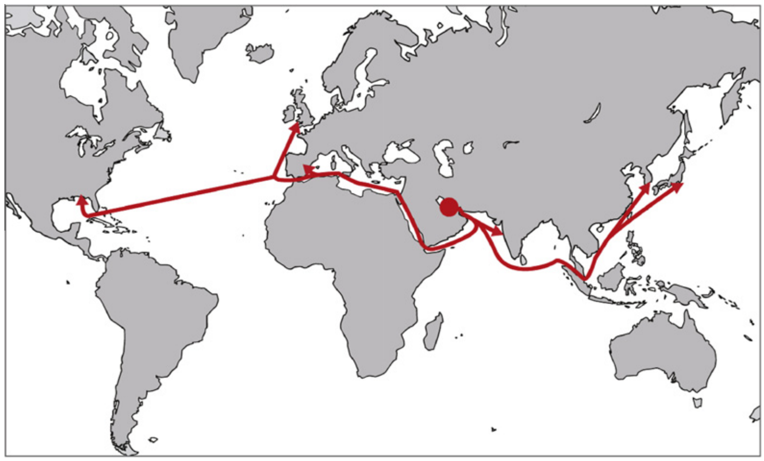

Fig. 1. Map showing some of the main sailing routes of the producer.

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 et al. (2011) for the same problem.

The paper is outlined as follows. First we present an overview of recent literature in Section 2. In Section 3 we present a more detailed description of the problem studied in this paper, before giving a mathematical formulation of the problem in Section 4. The construction and improvement heuristic is described in Section 5, followed by a computational study in Section 6. Finally, some concluding remarks are given in Section 7.

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. 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.

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 $\overline{T}_p$. 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 $\mathcal{T}_p = \{\underline{T}_p, \ldots, \overline{T}_p\}$. 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.

$$\min \sum_{c\in\mathcal{C}} \sum_{v\in\mathcal{V}_c} \sum_{t\in\mathcal{T}_v} C^T_{cv} x_{cvt} + \sum_{c\in\mathcal{C}^{LT}} \sum_{p\in\mathcal{P}} C^D_{cp} y_{cp} - \sum_{c\in\mathcal{C}^S} \sum_{v\in\mathcal{V}_c} \sum_{t\in\mathcal{T}_v} R^S_c L_v x_{cvt} - \sum_{g\in\mathcal{G}} R^I_g i_{g|\mathcal{T}|}, \tag{4.1}$$

$$\sum_{c\in\mathcal{C}_g\setminus\mathcal{C}^M} \sum_{v\in\mathcal{V}_c} x_{cvt} + \sum_{c\in\mathcal{C}^M\bigcap\mathcal{C}_g} \sum_{v\in\mathcal{V}^M} x_{cv(t-T_{cv}-1)} \leqslant B_g, \quad \forall g\in\mathcal{G},\ t\in\mathcal{T}, \tag{4.2}$$

$$i_{gt} - i_{g(t-1)} + \sum_{c\in\mathcal{C}_g\setminus\mathcal{C}^M} \sum_{v\in\mathcal{V}_c} L_v x_{cvt} = P_{gt}, \quad \forall g\in\mathcal{G},\ t\in\mathcal{T}, \tag{4.3}$$

$$\underline{I}_{gt} \leqslant i_{gt} \leqslant \bar{I}_{gt}, \quad \forall g\in\mathcal{G},\ t\in\mathcal{T}, \tag{4.4}$$

$$\sum_{v\in\mathcal{V}_c} \sum_{t\in\mathcal{T}_p} L_v x_{cvt} + y_{cp} \geqslant D_{cp}, \quad \forall c\in\mathcal{C}^{LT},\ p\in\mathcal{P}, \tag{4.5}$$

$$\sum_{c\in\mathcal{C}} \sum_{t-T_{cv}<\tau\leqslant t} x_{cv\tau} \leqslant 1, \quad \forall v\in\mathcal{V}^P,\ t\in\mathcal{T}_v, \tag{4.6}$$

$$\sum_{c\in\mathcal{C}^M} \sum_{t\in\mathcal{T}^M_v} x_{cvt} = 1, \quad \forall v\in\mathcal{V}^M, \tag{4.7}$$

$$x_{cvt} \in \{0,1\}, \quad \forall c\in\mathcal{C},\ v\in\mathcal{V}_c,\ t\in\mathcal{T}_v, \tag{4.8}$$

$$y_{cp} \geqslant 0, \quad \forall c\in\mathcal{C},\ p\in\mathcal{P}. \tag{4.9}$$

The objective function (4.1) 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 (4.2) 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 (4.3) and (4.4) make sure that the inventory levels in the storage tanks are between their upper and lower limits. Constraints (4.5) are soft and ensure that the demand of each long-term contract is met in every time interval.

Constraints (4.6) restrict each ship to sail at most one scheduled voyage on any given day. Constraints (4.7) state that every ship that is planned for maintenance will undertake maintenance exactly once. Finally, constraints (4.8) limit the scheduled voyages to be binary on the days the ship is available, while (4.9) 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 (4.2), *inventory-feasible* if it does not violate (4.3) and (4.4), and *routing-feasible* if it does not violate (4.6).

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.

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 4.

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 1 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 $\overline{T}_m$ 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 $\mathcal{C}^Q_{\alpha\beta}$, ranked according to the contract ranking parameter α and the inventory control parameter β. $\mathcal{C}^Q_{\alpha\beta}$ contains all contracts from $\mathcal{C}^{LT}$ and $\mathcal{C}^S$.

Ships are then selected from the ordered set $\mathcal{V}^Q_c$, 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^{RF}_v be the first day it is routing-feasible to send out ship v, and T^{IF}_{cv} 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 $T^E = max\left\{T^{RF}_v, T^{IF}_{cv}\right\}$. 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^{RF}_v 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.

Algorithm 1. CreateADP (α, β, κ)

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

5.1.1. 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:

$$D^*_{c(m+1)} = D_{c(m+1)} + \left[D^*_{cm} - \sum_{t=\underline{T}_m}^{\overline{T}_m} \sum_{s \in \mathcal{S}_{ct}} L_{v(s)} \right]. \tag{5.1}$$

5.1.2. 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 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 2.

Algorithm 2. Contract ranking scheme

---

if $|\rho_{c_1m} - \rho_{c_2m}| \leqslant \alpha$ **then**
  return $argmax\left\{\rho_{c_1m} D^*_{c_1m}, \rho_{c_2m} D^*_{c_2m}\right\}$
else
  return $argmax\{\rho_{c_1m}, \rho_{c_2m}\}$
end if

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 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 $\mathcal{C}^Q_{\alpha\beta}$ 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.

5.1.3. Ship ranking

Each contract in $\mathcal{C}^Q_{\alpha\beta}$ has a queue of ships associated with it, $\mathcal{V}^Q_c$, 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.

5.1.4. 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 $\left[\underline{T}^M_v, \overline{T}^M_v\right]$. 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 $\overline{T}^M_v$, unless maintenance has already been scheduled. Once it is decided to schedule maintenance, it is scheduled as early as possible.

5.1.5. The objective function

The objective function presented here is equivalent to objective function (4.1) using the planning horizon and months as time intervals, but will be restated using the symbols defined in Section 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, $\mathcal{S}_c \subset \mathcal{S}$ 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^T_s be the transportation cost associated with voyage s, and let C^D_c and C^D_{cm} 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 $c^S(g) \in \mathcal{C}^S \bigcap \mathcal{C}_g$ the spot contract for gas type g. The function used to evaluate the ADP is then given as:

$$z = \sum_{s\in\mathcal{S}} C_s^T + \sum_{c\in\mathcal{C}^{LT}} \sum_{m\in\mathcal{M}} max\{0, D_{cm} - \sum_{s\in\mathcal{S}_{cm}} L_{v(s)}\} \cdot C_{cm}^D + \sum_{c\in\mathcal{C}^{LT}} max\left\{0, D_c - \sum_{s\in\mathcal{S}_c} L_{v(s)}\right\} \cdot C_c^D - \sum_{g\in\mathcal{G}} \sum_{s\in\mathcal{S}_{c^{S}(g)}} L_{v(s)} R^S_{c^S(g)} - \sum_{g\in\mathcal{G}} R^I_g i_{g|\mathcal{T}|} \quad (5.2)$$

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.

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 5.1.5 through a local search. Since the CIH is a multi-start heuristic the goal of the local search is to quickly find better solutions 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*: $(c, v, t) \to (c', v, t)$
   The ship v starting a scheduled voyage to contract c on day t is assigned another contract c'.
2. *Changing ship*: $(c, v, t) \to (c, v', t)$
   The contract c served on a scheduled voyage starting day t is assigned another ship v'.
3. *Swapping contracts*: $(c, v, t), (c', v', t') \to (c', v, t), (c, v', t')$
   The contracts of two scheduled voyages are swapped.
4. *Swapping ships*: $(c, v, t), (c', v', t') \to (c, v', t), (c', v, t')$
   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 $\mathcal{C}^Q_{\alpha\beta}$, 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, $\mathcal{V}^Q_c$ 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.

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 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 (2011) 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.

5.3.1. 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:

$$x_{cvt} \in \{0, 1\}, \quad \forall c \in \mathcal{C}^{LT}, \ v \in \mathcal{V}_c, \ t \in \mathcal{T}_v, \ \mathcal{S}_{ct} \neq \emptyset. \quad (5.3)$$

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:

$$x_{cvt} \in \{0, 1\}, \quad \forall g \in \mathcal{G}, \ c \in \mathcal{C}_g \setminus \mathcal{C}^M, \ v \in \mathcal{V}_c, \ t \in \mathcal{T}_v, \ \mathcal{S}_{c^S(g)t} \neq \emptyset. \quad (5.4)$$

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:

$$x_{cvt} \in \{0, 1\}, \qquad \forall c \in \mathcal{C}^M, \ v \in \mathcal{V}^M, \ t \in \mathcal{T}^M_v. \quad (5.5)$$

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

5.3.2. 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:

$$\sum_{v\in V_c} x_{cvt} = |\mathcal{S}_{ct}|, \quad \forall c \in \mathcal{C}^{LT}, \ t \in \mathcal{T}, \ \mathcal{S}_{ct} \neq \emptyset, \quad (5.6)$$

$$\sum_{v\in V_c} x_{cvt} \leqslant |\mathcal{S}_{ct}|, \quad \forall c \in \mathcal{C}^S, \ t \in \mathcal{T}. \quad (5.7)$$

The reason why constraints (5.7) are $\leqslant$ constraints is that we have generated some variables in (5.4) which are not part of these constraints. If one of these variables are equal to one, all the variables in one constraint in (5.7) will most likely have to be equal to zero, or the ADP will not be inventory-feasible.

5.3.3. Overview of the mixed integer programming heuristic

The MIP heuristic solves a modified version of the formulation presented in Section 4, based on input from a feasible ADP. The objective function (4.1) and constraints (4.2)–(4.7) are all kept, and, in addition, (5.6) and (5.7) are added to the model. Instead

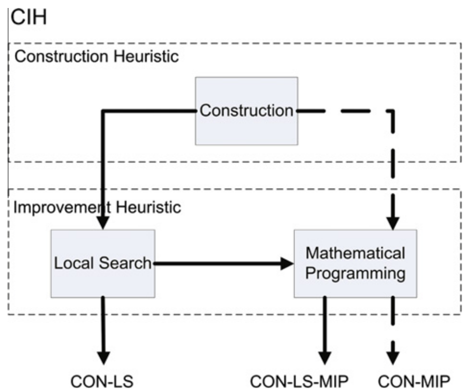

Fig. 2. Flowchart showing the different modules of the CIH and how they interact.

of generating variables as stated in (4.8) and (4.9), only the variables given by (5.3)–(5.5) are created. By solving this reduced version of the problem we can find good feasible solutions quickly using branch-and-bound.

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 4. Fig. 2 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 6, where a computational study of the heuristic is conducted.

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.

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 1. 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 1 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

Table 1
Instances.

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

docks all have been completed. The main 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.

6.2. Exact solution method

In this section the computational results of the mathematical model given in Section 4 are presented. Each instance was run for a maximum of 24 h (86 400 s) and Table 2 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 different 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.

Table 2
Test of the exact solution method.

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

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.

Fig. 3 shows the deviation between the best obtained ADP and the best 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 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 1.

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 heuristic 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.

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.

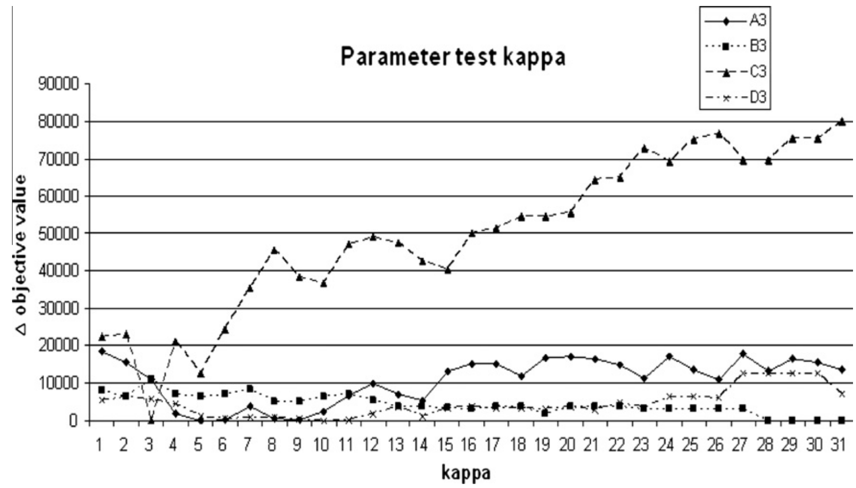

Fig. 3. Graph showing the variation in the minimal objective value for each κ value, compared to the best found objective value, for each of the four instances.

Table 3 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 2, 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). 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 4 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 search.

Table 3
Local search heuristic tests.

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 4
MIP heuristic tests.

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

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 5 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 4. The XpressMP implementation of the full mathematical model was run for a maximum of 86 400 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.

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 h 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).

Table 5
Comparison of computational times.

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

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.

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 inprovement 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.

References

Andersson, H., Christiansen, M., & Fagerholt, K. (2010). Transportation planning and inventory management in the LNG supply chain. In P. M. Pardalos, E. Bjørndal, M. Bjørndal, & M. Rönnqvist (Eds.), *Energy, natural resources and environmental economics. Energy systems* (pp. 427–439). Berlin, Heidelberg: Springer.

Archetti, C., Bertazzi, L., Speranza, M. G., & Hertz, A. (2011). A hybrid heuristic for an inventory routing problem. *INFORMS Journal on Computing.* doi:10.1287/ijoc.1100.0439.

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 (2008). International energy outlook 2008. Tech. rep., US Department of Energy, Energy Information Administration US 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 L. Bertazzi, J. v. Nunen, & M. Speranza (Eds.), *Innovations in distribution logistics. Lecture notes in economics and mathematical systems* (pp. 195–218). Springer.

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., Stslhåne, M., Moe, C., Andersson, H., Christiansen, M., Fagerholt, K., et al. (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.