



Maritime inventory routing with transshipment: the case of Yamal LNG

Mingyu Li¹ · Kjetil Fagerholt¹ · Peter Schütz¹ 

Accepted: 19 October 2022
© The Author(s) 2022

Abstract

The LNG-ADP problem is a tactical planning problem for creating an annual delivery program (ADP) for a liquefied natural gas (LNG) producer. An ADP specifies the departure dates of the LNG carriers as well as the delivery dates at the different customers for a period of 12–18 months. The problem can be formulated as maritime inventory routing problem, as it is an important requirement to plan the deliveries such that inventory levels are kept within minimum and maximum limits at the customers as well as the LNG production facility. Inspired by the case of Yamal LNG, we propose a novel discrete-time formulation for the LNG-ADP problem with transshipment and intermediate storage. Our formulation also allows for waiting at the unloading ports. The problem is solved using a rolling horizon heuristic (RHH) for a case based on the Yamal LNG project. We study the impact of different RHH configurations on run time and solution quality. The results show that using a central period that is shorter than the forecast period provides the best objective function, whereas a central period that is longer than the forecast period improves run time. We also explore the effect of allowing waiting at the unloading ports. Waiting does not necessarily improve the objective function value, despite increasing the solution space. However, we observe a reduction in run time for instances where waiting is allowed.

Keywords Maritime inventory routing · Transshipment · LNG-ADP

Kjetil Fagerholt and Peter Schütz contributed equally to this work.

✉ Peter Schütz
peter.schuetz@ntnu.no

Mingyu Li
mingyu.li@ntnu.no

Kjetil Fagerholt
kjetil.fagerholt@ntnu.no

¹ Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology, 7491 Trondheim, Norway

1 Introduction

Traditionally, liquefied natural gas (LNG) has been shipped directly from the liquefaction plant at the producer port to the regasification plant at the customer port. However, in 2015 Yamal LNG signed a 20-year contract with Fluxys LNG for transshipment of LNG at the Zeebrugge LNG terminal in Belgium (Schach and Madlener 2018). Due to the liquefaction plant's location in Sabetta on the Yamal peninsula in the Russian Arctic, ice-going LNG carriers are needed to ensure year-round deliveries. These carriers can serve Asian and European customers directly, but if the eastbound route along the North East Passage becomes unavailable, they will sail to Zeebrugge where the LNG is transshipped to an open-water LNG carrier that is then serving the Asian customers, sailing through the Suez Canal.

Historically, LNG has been traded in long-term contracts, but the share of cargoes sold in spot or short-term transactions is increasing (GIIGNL 2021). Still, planning deliveries to customers remains an important task for LNG producers. As part of their tactical planning, LNG producers create an Annual Delivery Program (ADP). The purpose of the ADP is to ensure that the producer satisfies the customers' demand at lowest possible cost (Andersson et al. 2017). The plan typically specifies the departure dates of the LNG carriers as well as the delivery dates at the different customers for a period of 12–18 months. An important requirement for the ADP is to plan the deliveries such that minimum and maximum inventory limits at the customers are observed to prevent any disruption in their production. The volume committed to the long-term customers usually accounts for most, but not all produced LNG (Stålhane et al. 2012). The rest of the production is sold in the spot market, which may release the inventory pressure at the producer while also bringing extra revenues. The ADP therefore often includes sales in spot markets and the corresponding deliveries as well.

The LNG-ADP problem is categorized as a maritime inventory routing problem (MIRP) where ship routing and scheduling are combined with inventory management. Please see Papageorgiou et al. (2014), Christiansen and Fagerholt (2014) and Christiansen et al. (2013) for a more general discussion of MIRPs. The LNG-ADP problem is known to be very difficult to solve and most literature resorts to heuristics in order to solve the problem. One way is to divide the entire planning horizon into several sub-horizons, and solve the subproblems sequentially in a rolling horizon heuristics, see Rakke et al. (2011) and Al-Haidous et al. (2016). Alternatively, initial solutions can be constructed based on some greedy principle, e.g., by prioritizing cheaper and less flexible vessels when assigning unfulfilled contracts (Stålhane et al. 2012). Mutlu et al. (2016) also construct feasible solutions by applying greedy techniques for assigning vessels to contracts. Attempts to solve the LNG-ADP problem with exact methods are scarce in the literature. Two notable examples are Andersson et al. (2017) and Halvorsen-Weare and Fagerholt (2013). Andersson et al. (2017) introduce four different types of valid inequalities to strengthen the formulation of the LNG-ADP problem, while Halvorsen-Weare and Fagerholt (2013) decompose the problem into a scheduling/

inventory management problem and a subproblem for each vessel. All of the papers mentioned so far assume direct shipments from the production port to the customer port. Li and Schütz (2020) are the first to introduce the possibility of transshipment in the LNG-ADP problem. They use a continuous-time formulation and apply a rolling horizon heuristic to solve the problem.

The LNG-ADP problem usually considers the optimization of ship schedules together with inventory management at the production terminals only (and in our case also the transshipment terminal). Goel et al. (2012) study an LNG inventory routing problem where ship schedules are optimized together with inventory management on both the production and regasification terminals. They present an arc-flow formulation and propose a set of construction and improvement heuristics to solve it. The heuristics are tested on a set of realistic test instances with a planning horizon of up to a year (like for the ADP) and the results indicate that the proposed methods find high-quality solutions in reasonable time. Later, Shao et al. (2015) and Goel et al. (2015) consider the same problem and develop a hybrid heuristic strategy and a constraint programming method, respectively. Other notable contributions on MIRPs, though not necessarily with applications from the LNG business are Papageorgiou et al. (2015) and Papageorgiou et al. (2018). These papers study a more generic MIRP problem. However, as for the LNG-ADP, they consider long planning horizons with up to 360 time periods.

Transshipment has been considered in inventory routing problems (IRPs) for many years (e.g., Coelho et al. 2012; Peres et al. 2017). However, the LNG-ADP differs from other IRPs and MIRPs: In LNG-ADP, the network structure is often simpler with only one production port. It is also common to only consider full-ship loads, and hence in each voyage, only a single loading port and a single unloading port are involved. On the other hand, the problem size is often larger due to the fleet size and the length of the planning horizon (Andersson et al. 2017). Tailor-made models and solution methods are therefore often more efficient than methods developed for other IRPs.

We contribute to the literature in the following aspects: First, we extend the existing literature on the LNG-ADP problem by providing a novel discrete-time formulation for the LNG-ADP problem with transshipment and waiting at customer ports. Existing formulations either do not consider transshipment or use a continuous-time formulation. Second, we use our formulation to solve a case based on the Yamal LNG project. Yamal LNG is the northernmost producer of LNG in the world and can serve customers directly or through a transshipment port if the direct route is not available. Third, we show that a rolling-horizon based heuristic (RHH) can solve the problem efficiently. We also discuss how different waiting times and different configurations for the RHH impact solution quality and run time.

The remainder of this paper is organized as follows. We describe the LNG-ADP problem with transshipment in more detail in Sect. 2. Modeling approach and mathematical formulation are presented in Sect. 3. Our solution approach is described in Sect. 4, followed by the Yamal LNG case study in Sect. 5. Results from the computational study are discussed in Sect. 6. We conclude in Sect. 7.

2 Problem description

We consider the following LNG-ADP: The network consists of one production port, one transshipment port and multiple customers. We categorize the customers into three customer types according to two criteria. The first criterion is based on how customers can be reached from the production port, see also Fig. 1: Customers of type 1 can only be served directly from the production port in all time periods. These will typically be customers located in Europe for our Yamal case study. Customers of type 2, which represent customers located in Asia, can be served from the production port directly in some periods of the planning horizon when the ice conditions along the North East Passage allow it (i.e., during summer time), but not all of them (i.e., during winter time when there is too much ice). However, these customers can always be served through the transshipment port. Customers that can only be served through the transshipment port belong to type 3. The second criterion for categorizing customers is the type of contract the customer has: A customer is either a long-term customer with a contract specifying delivery requirements for different periods of the planning horizon, or a short-term customer, representing a spot market that can receive deliveries at any point in time. Note that all customers of types 1–3 are either a long-term or a short-term customer.

We distinguish between two types of LNG carriers: there is one fleet of carriers owned and operated by the producer and a fleet of carriers that can be chartered on a voyage-by-voyage basis. The producer's own fleet serves customers of types 1 and 2 (if the direct route is available) as well as the transshipment port. For routes from the transshipment port to customers of type 2 and 3, chartered LNG carriers are used. Each of the fleets is homogeneous, but the fleet of all carriers is heterogeneous. We assume that the size of the producer's fleet is given, whereas chartered carriers are always available when needed.

Voyages carried out by the producer's own carriers start by loading at the production port, from which the carrier sails to either the transshipment port or a customer

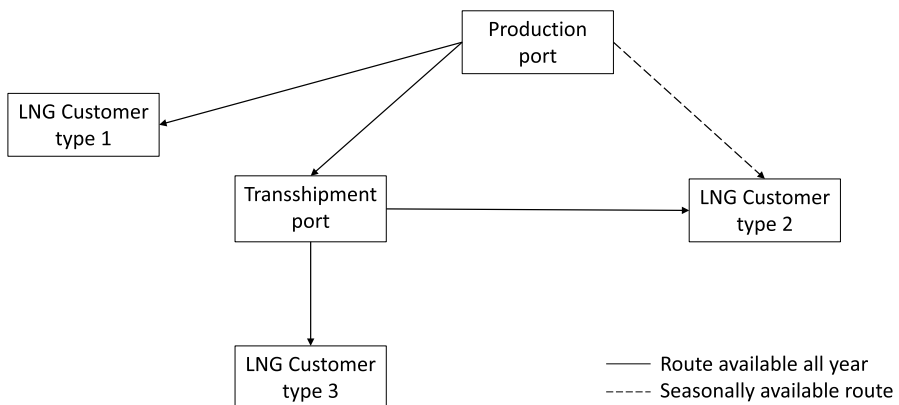


Fig. 1 Illustration of the transportation network for shipping LNG and the availability of the routes. Customer types are distinguished by how they can be reached from the production port

port. The return leg of the voyage starts with unloading cargo and ends with the arrival back at the production port. For chartered carriers, we only model the delivery leg of the voyage, so a voyage starts with loading cargo at the transshipment port and ends with unloading at the customer port. For both types of carriers, we allow waiting at the customer port before starting to unload cargo. This introduces some flexibility for the producer, e.g., in case a carrier has to leave early from the loading port due to congestion, but wants to satisfy demand in a later period. Waiting time at the customer port is limited, as an overly long waiting time before delivery imposes additional operating costs and reduces utilization of the fleet capacity.

We only consider full shiploads. However, due to the presence of boil-off of LNG, which means that a small amount of the LNG onboard the vessel evaporates during the voyage, the amount delivered to customers is less than the amount loaded at the production or transshipment port (see Hasan et al. 2009, for details).

The long-term customers' demand for LNG is given on a monthly basis for the entire planning horizon. It is common to allow some flexibility in deliveries, i.e., small deviations from the annual demand (Andersson et al. 2017). In this paper we consider deviations within one shipload as small deviations. We impose a piecewise linear penalty cost to incentivize meeting the annual demand, as illustrated in Fig. 2. Small deviations receive a much lower penalty per m^3 than large deviations, i.e., deviations above one shipload. Note that deviations may come in as under-deliveries and over-deliveries. We set the under-delivery penalty rates to be higher than the penalty rates for over-delivery, as the former is less desirable for the customers.

We also introduce penalties for deviations from monthly demand to motivate LNG cargoes to be delivered during the specified month. The monthly deviations are penalized at fixed rates per m^3 . The unit under-delivery penalty is, except for small deviations, set to be greater than the unit revenue from spot sales such that satisfying long-term customers is prioritized over the sales to spot market. Penalties and spot price are set in the following ascending order: penalty for small deviations, penalty for monthly over-delivery, price at spot markets, penalty for monthly under-delivery, penalty for large deviations. With this setting, the producer has some

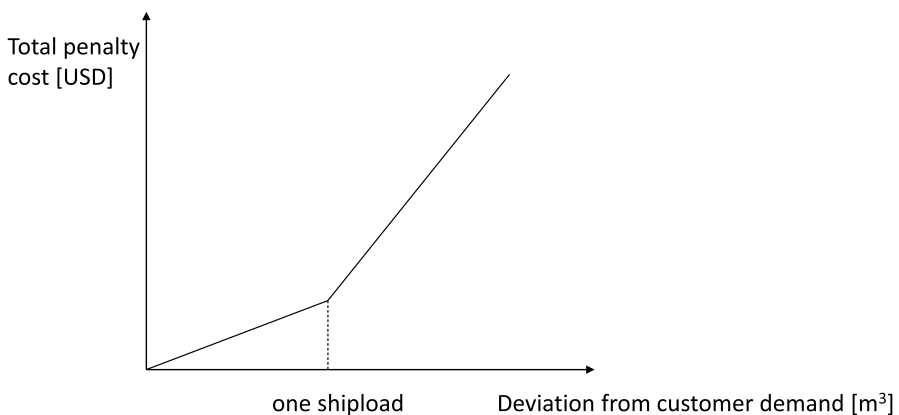


Fig. 2 Penalty cost function for deviating from a customer's long-term demand

flexibility to alter delivery time and volume, but also has clear incentive to meet the annual demand.

The producer is responsible for maintaining the inventory levels of the storage tanks within given limits, both at the production port and the transshipment port. This implies that loading at either production port or transshipment port cannot be scheduled before the inventory level exceeds one shipload and that unloading cannot happen unless there is sufficient storage capacity left at the transshipment port. Inventory management at the customers is handled through the deviation penalties.

The number of berths at each port for loading and unloading operations is limited. It is not possible to schedule more operations at any port during any period of the planning horizon at any port than berths are available.

The purpose of creating an ADP is to minimize the total costs of serving the long-term customers. Besides transportation and charter costs, the total costs include penalties for deviating from customer demand. We also include revenues from spot sales in the objective function to reflect that surplus LNG can be sold in a spot market. With the option of using a transshipment port, the LNG producer decides on the monthly quantities delivered to the long-term customers, which carrier and which route is used to serve the customers, the corresponding loading and unloading times, as well as when and how much to sell to the spot markets. The ADP has to be feasible with respect to the available routes, inventory limits and berth capacities at ports.

3 Model formulation

We formulate the LNG-ADP problem with transshipment as a discrete-time mixed integer linear programming (MILP) model. Our formulation is inspired by other discrete-time models for MIRPs, e.g., Agra et al. (2013) and Goel et al. (2012), though significantly extended in order to handle the transshipment part of our problem. We first discuss discrete-time and continuous-time formulations for MIRPS, before we introduce our approach to modeling voyages and explain how waiting is included in the model. We then present the notation used in the model, followed by the mathematical formulation. The duration of the underlying time period in the model is a day.

3.1 Discrete-time versus continuous-time

Continuous-time formulations are usually smaller than discrete-time formulations (Agra et al. 2017). However, the main characteristic of a discrete-time model is that decision variables and parameters can be indexed by time. This facilitates modeling and implementing time-dependencies in decisions and data. For example, varying production and/or consumption rates are often modeled using discrete-time formulations in MIRPs, while continuous-time formulations are most common when these rates are constant (Christiansen et al. 2013). Sailing speed is another typical time-dependent parameter (see e.g., Lianes et al. 2021). In continuous-time formulations, these time-dependencies have to be handled using additional variables and big-M

constraints. This might both give weaker linear relaxations and increase the size of the problem, possibly removing some of the advantage of the smaller formulation.

Implementing a RHH in a discrete-time framework is also simpler than for the corresponding continuous-time formulation as all variables that need to be moved to the locked period are known in advance for each iteration. In a continuous-time formulation, decision variables have to be locked based on value rather than time index. Fixing a continuous variable for subsequent iterations might even introduce rounding errors, potentially causing the new problem to become infeasible. We therefore use a discrete-time formulation for the LNG-ADP problem with transshipment.

3.2 Modeling voyages

We adopt different modeling approaches for voyages with owned and chartered carriers. For owned carriers, we model both loading and unloading voyages. Voyages by owned carriers start with loading LNG at the production port (delivery voyages) and end with the arrival at the customer port. The return voyage starts by unloading LNG at the transshipment or customer port and ends with the arrival of the carrier at the production port. The decision variables keep track of which carrier is assigned to the different voyages to ensure that each carrier is assigned to only one voyage at a time and that the carriers return to the production port before starting a new delivery voyage. Between the end of a delivery voyage and the start of the return voyage, the carrier is allowed to wait for a few days. Waiting is in this case modeled by means of a separate decision variable.

For chartered carriers, we model only the voyage from the transshipment port to the customer port. This delivery voyage includes both loading and unloading operations. The reason for this that we do not have to keep track of the chartered carrier's return voyage as we assume that the carrier is only chartered for a single voyage. Waiting at customer ports before starting to unload the cargo is modeled by means of a double time index, where the first time index denotes the loading time and the second time index denotes unloading time. As a consequence, we have multiple decision variables for a delivery voyages starting on the same day, but delivering on different days.

3.3 Notation

We introduce the following notation:

3.3.1 Sets

\mathcal{B} Set of types of LNG carriers that are available for chartering. These carriers only load at the transshipment port.

\mathcal{G} Set of aggregated delivery periods.

- \mathcal{P} Set of ports.
- \mathcal{P}_v^A Set of ports that can be visited by carrier v , $v \in \mathcal{V}$.
- \mathcal{P}_{vi}^{AF} Set of ports that can be visited by carrier v from port i , $v \in \mathcal{V}, i \in \mathcal{P}$.
- \mathcal{P}^C Set of customer ports.
- \mathcal{P}_b^C Set of customer ports that can be visited by a chartered carrier of type b , $b \in \mathcal{B}$.
- \mathcal{P}^N Set of customers of type 2.
- \mathcal{P}^L Set of long-term customers.
- \mathcal{P}^P Set of production ports.
- \mathcal{P}^S Set of short-term customers/spot markets.
- \mathcal{P}^T Set of transshipment ports.
- \mathcal{T} Set of time periods in the entire planning horizon.
- \mathcal{T}^W Set of time periods when the direct route to customers of type 2 is closed.
- \mathcal{T}_g Set of time periods belonging to aggregated delivery period g , $g \in \mathcal{G}$.
- \mathcal{T}_{bjt}^C Set of unloading time periods for a chartered carrier of type b that loads at time t and unloads at port j , taking into account sailing and waiting times, $b \in \mathcal{B}, j \in \mathcal{P}_b^C, t \in \mathcal{T}$.
- \mathcal{V} Set of owned LNG carriers. These carriers only load at the production port.

3.3.2 Parameters

- A Maximum allowed waiting time before delivery.
- C_{vijt} Voyage cost of an owned carrier from port i to port j with carrier v when departing at time t , $v \in \mathcal{V}, i \in \mathcal{P}_v^A, j \in \mathcal{P}_{vi}^{AF}, t \in \mathcal{T}$.

- C_{bjte}^C Voyage cost of a chartered carrier of type b from the transshipment port to port j when departing at time t and delivering at time e , $b \in \mathcal{B}, j \in \mathcal{P}_b^C, t \in \mathcal{T}, e \in \mathcal{T}_{bjt}^C$.
- D_{jg} Demand (in m^3) at customer j in delivery period g , $j \in \mathcal{P}^L, g \in \mathcal{G}$.
- K Breakpoint in penalty cost function for deviating from annual LNG demand.
- K_v Loading capacity of owned carrier v , $v \in \mathcal{V}$.
- K_b^C Loading capacity of chartered carrier of type b , $b \in \mathcal{B}$.
- L_{vijt} Amount of LNG unloaded at port i by owned carrier v before sailing back to port j at time t , $v \in \mathcal{V}, i \in \mathcal{P}_v^A \setminus \mathcal{P}^P, j \in \mathcal{P}_{vi}^{AF}, t \in \mathcal{T}$.
- L_{bjte}^C Amount of LNG unloaded at port j by chartered carrier b at time e when voyage starts at time t , $b \in \mathcal{B}, j \in \mathcal{P}_b^C, t \in \mathcal{T}, e \in \mathcal{T}_{bjt}^C$.
- O_v Initial position of carrier v , $v \in \mathcal{V}, O_v \in \mathcal{P}_v^A$.
- P_{it} Production rate at port i in time period t , $i \in \mathcal{P}^P, t \in \mathcal{T}$.
- Q_j^+ Penalty per m^3 for annual over-delivery at long-term customer j above one shipload, $j \in \mathcal{P}^L$.
- Q_j^- Penalty per m^3 for annual over-delivery at long-term customer j below one shipload, $j \in \mathcal{P}^L$.
- Q_{jg} Penalty per m^3 for over-delivery in delivery period g at long-term customer j , $j \in \mathcal{P}^L, g \in \mathcal{G}$.
- R_{vijt} Revenue from selling one shipload LNG by owned carrier v to spot market i at time t before sailing back to port j , $v \in \mathcal{V}, i \in \mathcal{P}^S \cap \mathcal{P}_v^A, j \in \mathcal{P}_{vi}^{AF}, t \in \mathcal{T}$.
- R_{bjte}^C Revenue from selling one shipload LNG by chartered carrier of type b to spot market j at time e when voyage starts at time t , $b \in \mathcal{B}, j \in \mathcal{P}^S \cap \mathcal{P}_b^C, t \in \mathcal{T}, e \in \mathcal{T}_{bjt}^C$.
- S_i Initial storage at port i , $i \in \mathcal{P}^P \cup \mathcal{P}^T$.
- \bar{S}_i Upper bound on inventory level at port i , $i \in \mathcal{P}^P \cup \mathcal{P}^T$.
- \underline{S}_i Lower bound on inventory level at port i , $i \in \mathcal{P}^P \cup \mathcal{P}^T$.
- T_{eoy} End of year.

- T_{vijt} Time for owned carrier v to load/unload and sail from port i to j if loading/unloading starts at time t , $v \in \mathcal{V}$, $i \in \mathcal{P}_v^A$, $j \in \mathcal{P}_{vi}^{AF}$, $t \in \mathcal{T}$
- T_v Time when owned carrier v becomes available, $v \in \mathcal{V}$.
- U_j^+ Penalty per m^3 for annual under-delivery at long-term customer j above one shipload, $j \in \mathcal{P}^L$.
- U_j^- Penalty per m^3 for annual under-delivery at long-term customer j below one shipload, $j \in \mathcal{P}^L$.
- U_{jg} Penalty per m^3 for under-delivery in delivery period g at long-term customer j , $j \in \mathcal{P}^L$, $g \in \mathcal{G}$.
- Z_i Berth capacity at port i , $i \in \mathcal{P}$.

3.3.3 Decision variables

- q_j^+ Annual over-delivery (in m^3) at long-term customer j above one shipload, $j \in \mathcal{P}^L$.
- q_j^- Annual over-delivery (in m^3) at long-term customer j below one shipload, $j \in \mathcal{P}^L$.
- q_{jg} Over-delivery (in m^3) at long-term customer j in delivery period g , $j \in \mathcal{P}^L$, $g \in \mathcal{G}$.
- s_{it} Inventory level at port i at the beginning of time period t , $i \in \mathcal{P}^L$, $t \in \mathcal{T}$.
- u_j^+ under-delivery (in m^3) at long-term customer j above one shipload, $j \in \mathcal{P}^L$.
- u_j^- Annual under-delivery (in m^3) at long-term customer j below one shipload, $j \in \mathcal{P}^L$.
- u_{jg} Under-delivery (in m^3) at long-term customer j in delivery period g , $j \in \mathcal{P}^L$, $g \in \mathcal{G}$.
- w_{vit} 1 if own carrier v is waiting outside port i in time period t , 0 otherwise, $v \in \mathcal{V}$, $i \in \mathcal{P}_v^A$, $t \in \mathcal{T}$.

x_{vijt} 1 if owned carrier v starts loading/unloading at port i in time period t and sails to port j afterwards, 0 otherwise, $v \in \mathcal{V}, i \in \mathcal{P}_v^A, j \in \mathcal{P}_{vi}^{AF}, t \in \mathcal{T}$.

y_{bjte} 1 if chartered carrier of type b loads on day t at the transshipment port and unloads at customer port j on day e , 0 otherwise, $b \in \mathcal{B}, j \in \mathcal{P}_b^C, t \in \mathcal{T}, e \in \mathcal{T}_{bjt}^C$.

3.4 Mathematical model

With the notation introduced above, we present the discrete-time formulation for the LNG-ADP planning problem with transshipment.

3.4.1 Objective function

The objective of the LNG-ADP planning problem with transshipment is to minimize the total costs of transporting the LNG and given as

$$\begin{aligned} \min & \sum_{v \in \mathcal{V}} \sum_{i \in \mathcal{P}_v^A} \sum_{j \in \mathcal{P}_{vi}^{AF}} \sum_{t \in \mathcal{T}} C_{vijt} x_{vijt} + \sum_{b \in \mathcal{B}} \sum_{j \in \mathcal{P}_b^C} \sum_{t \in \mathcal{T}} \sum_{e \in \mathcal{T}_{bjt}^C} C_{bjte}^C y_{bjte} \\ & + \sum_{j \in \mathcal{P}^L} \sum_{g \in \mathcal{G}} (Q_{jg} q_{jg} + U_{jg} u_{jg}) + \sum_{j \in \mathcal{P}^L} (Q_j^+ q_j^+ + Q_j^- q_j^- + U_j^+ u_j^+ + U_j^- u_j^-) \quad (1) \\ & - \sum_{v \in \mathcal{V}} \sum_{i \in \mathcal{P}_v^A \cap \mathcal{P}^S} \sum_{j \in \mathcal{P}_{vi}^{AF}} \sum_{t \in \mathcal{T}} R_{vijt} x_{vijt} - \sum_{b \in \mathcal{B}} \sum_{j \in \mathcal{P}_b^C \cap \mathcal{P}^S} \sum_{t \in \mathcal{T}} \sum_{e \in \mathcal{T}_{bjt}^C} R_{bjte}^C y_{bjte} \end{aligned}$$

The first line of objective function (1) is the total sailing cost from owned carriers and chartered carriers, respectively. The second line sums up the penalties for monthly and annual deviations of LNG delivered at long-term customers. The last line calculates the total revenues from the spot market.

3.4.2 Routing and waiting constraints

Constraints (2)–(9) control the routing and waiting of the LNG carriers.

$$\sum_{i \in \mathcal{P}_v^A} w_{vit} + \sum_{i \in \mathcal{P}_v^A} \sum_{j \in \mathcal{P}_{vi}^{AF}} x_{vijt} \leq 1 \quad v \in \mathcal{V}, t \in \mathcal{T}, \quad (2)$$

$$w_{v,O_v,T_v} + \sum_{j \in \mathcal{P}_{v,O_v}^{AF}} x_{v,O_v,j,T_v} = 1 \quad v \in \mathcal{V}, \quad (3)$$

$$\begin{aligned} w_{vi,t-1} + \sum_{\{j \in \mathcal{P}_v^A | i \in \mathcal{P}_{vj}^{AF}\}} \sum_{\{\tau \in \mathcal{T} | \tau + T_{vj\tau} = t\}} x_{vj\tau} &= w_{vit} + \sum_{j \in \mathcal{P}_{vi}^{AF}} x_{vijt} \quad (4) \\ v \in \mathcal{V}, i \in \mathcal{P}_v^A, t \in \mathcal{T}, \end{aligned}$$

$$\sum_{v \in \mathcal{V}} \sum_{j \in \mathcal{P}_{vi}^{AF}} x_{vijt} \leq Z_i \quad i \in \mathcal{P}^P, t \in \mathcal{T}, \quad (5)$$

$$\sum_{v \in \mathcal{V}} \sum_{j \in \mathcal{P}_{vi}^{AF}} x_{vijt} + \sum_{b \in \mathcal{B}} \sum_{j \in \mathcal{P}_b^C} \sum_{e \in \mathcal{T}_{bjt}^C} y_{bjte} \leq Z_i \quad (6)$$

$$i \in \mathcal{P}^T, t \in \mathcal{T},$$

$$\sum_{\{v \in \mathcal{V} | i \in \mathcal{P}_v^A\}} \sum_{j \in \mathcal{P}_{vi}^{AF}} x_{vijt} + \sum_{\{b \in \mathcal{B} | i \in \mathcal{P}_b^C\}} \sum_{\{e \in \mathcal{T} | t \in \mathcal{T}_{bjt}^C\}} y_{bjet} \leq Z_i \quad (7)$$

$$i \in \mathcal{P}^C, t \in \mathcal{T},$$

$$\sum_{v \in \mathcal{V}} \sum_{i \in \mathcal{P}^P} \sum_{j \in \mathcal{P}_{vi}^{AF} \cap \mathcal{P}^N} \sum_{t \in \mathcal{T}^W} x_{vijt} = 0 \quad (8)$$

$$x_{vijt} - \sum_{\{e \in \mathcal{T} | e + T_{vje} \leq t \leq e + T_{vje} + A\}} x_{vjie} \leq 0 \quad (9)$$

$$v \in \mathcal{V}, j \in \mathcal{P}^P, i \in \mathcal{P}_{vj}^{AF}, t \in \mathcal{T}$$

Constraints (2) ensure that each carrier at any time period is either waiting at a port ($w_{vit} = 1$), starting a new voyage ($x_{vijt} = 1$), or sailing (left-hand-side = 0). Constraints (3) state that when the owned carriers become available at their origin ports, they can either be waiting, or start loading there. Constraints (4) keep track of the status of a carrier from day to day. If a carrier is waiting in port or just arrives at a port, then in the next day, it may either continue to wait or start port operations. Constraints (5)–(7) handle the berth limitation for the production port, the transshipment port and customer ports, respectively. On the left-hand-side are the cargo flows at the port, which vary with the port types. For the production port, we only consider loading owned carriers. For the transshipment port, both unloading owned carriers and loading chartered carriers are considered while for the customer ports, cargoes may be unloaded from either owned carriers or chartered carriers. However, the number of carriers operating cannot exceed berth capacity. Constraints (8) ensure that no carrier starts a new voyage using the direct route from the production port to a customer of type 2, when the direct route is closed. Constraints (9) limit waiting at the customer port before starting unloading and returning to the production port. If unloading is scheduled at time t , the corresponding delivery voyages must have started no earlier than the sailing time from the production port to the customer plus the maximum waiting time and no later than the time it takes to sail between the two ports.

3.4.3 Inventory management

For each port, inventory constraints have to be considered. We define inventory levels at the beginning of a time period as this allows us to initialize inventory levels inside the planning horizon.

$$s_{i0} = S_i \quad i \in \mathcal{P}^P \cup \mathcal{P}^T, \tag{10}$$

$$s_{it} = s_{i,t-1} + P_{i,t-1} - \sum_{v \in \mathcal{V}} \sum_{j \in \mathcal{P}_{vi}^{AF}} K_v x_{vij,t-1} \quad i \in \mathcal{P}^P, t \in \{T | t \leq T_{\text{eoy}}\}, \tag{11}$$

$$s_{it} \leq S_i^U \quad i \in \mathcal{P}^P, t \in \{T | t \leq T_{\text{eoy}}\}, \tag{12}$$

$$s_{it} \geq S_i^L \quad i \in \mathcal{P}^P, t \in \{T | t \leq T_{\text{eoy}}\}, \tag{13}$$

$$s_{it} = s_{i,t-1} + \sum_{\{v \in \mathcal{V} | i \in \mathcal{P}_v^A\}} \sum_{j \in \mathcal{P}_{vi}^{AF}} L_{vij,t-1} x_{vij,t-1} - \sum_{b \in \mathcal{B}} \sum_{j \in \mathcal{P}_b^C} \sum_{e \in \mathcal{T}_{bj,t-1}^C} K_b^C y_{bj,t-1,e} \quad i \in \mathcal{P}^T, t \in \mathcal{T} \tag{14}$$

$$s_{it} \leq S_i^U \quad i \in \mathcal{P}^T, t \in \mathcal{T}, \tag{15}$$

$$s_{it} \geq S_i^L \quad i \in \mathcal{P}^T, t \in \mathcal{T}, \tag{16}$$

Constraints (10) initialize the initial inventory at the production port as well as the transshipment port. Constraints (11) update the inventory at the production port with the amount produced and loaded. Constraints (12) and (13) require the inventory level at the production port to be within the limits. Note that since we do not allow new departures from the production port after the end of year, inventory management at the production port is out-of-scope of our model for the final period of the planning horizon. Similarly, constraints (14) update the inventory at the transshipment port with the amounts unloaded from the own carriers and loaded to the chartered carriers. Constraints (15) and (16) maintain the inventory level at the transshipment port within the upper and lower limits.

3.4.4 Demand satisfaction and deviations from contracts

Satisfying demand and calculating the deviations from both monthly and annual demand is handled through constraints (17)–(20).

$$\sum_{\{v \in \mathcal{V} \mid i \in \mathcal{P}_v^\lambda\}} \sum_{j \in \mathcal{P}_{vi}^{AF}} \sum_{t \in \mathcal{T}_g} L_{vijt} x_{vijt} + \sum_{b \in \mathcal{B}} \sum_{\{e \in \mathcal{T} \mid t \in \mathcal{T}_{bje}^c\}} \sum_{t \in \mathcal{T}_g} L_{biet} y_{biet} \quad (17)$$

$$+ u_{ig} - q_{ig} = D_{ig} \quad i \in \mathcal{P}^L, g \in \mathcal{G},$$

$$\sum_{g \in \mathcal{G}} \sum_{t \in \mathcal{T}_g} \left(\sum_{\{v \in \mathcal{V} \mid i \in \mathcal{P}_v^\lambda\}} \sum_{j \in \mathcal{P}_{vi}^{AF}} L_{vijt} x_{vijt} + \sum_{b \in \mathcal{B}} \sum_{\{e \in \mathcal{T} \mid t \in \mathcal{T}_{bje}^c\}} L_{biet} y_{biet} \right) \quad (18)$$

$$+ u_i^+ + u_i^- - q_i^+ - q_i^- = \sum_{g \in \mathcal{G}} D_{ig} \quad i \in \mathcal{P}^L,$$

$$u_i^- \leq K \quad i \in \mathcal{P}^L, \quad (19)$$

$$q_i^- \leq K \quad i \in \mathcal{P}^L, \quad (20)$$

Constraints (17) and (18) sum the total amount of LNG delivered to long-term customers and calculate the deviations from the contract in each month and throughout the planning horizon, respectively. Restrictions (19) and (20) limit the small deviations to be within one shipload.

3.4.5 Non-negative and binary requirements

$$q, q^+, q^-, s, u, u^+, u^- \geq 0, \quad (21)$$

$$w, x, y \in \{0, 1\}. \quad (22)$$

Constraints (21) and (22) are the non-negativity constraints for the continuous variables and the binary requirements for the binary variables, respectively. Note that indices have been omitted.

3.4.6 Valid inequalities

To strengthen the model formulation further, we add delivery inequalities as presented in Andersson et al. (2017) to the model. The purpose of these valid inequalities is to remove solutions from the feasible space that reduce the penalties for deviating from the demand by allowing fractional solutions.

4 Solution approach

In this section we present our solution approach. First, we describe the overall rolling horizon heuristic (RHH) framework. We then introduce heuristic rules to reduce symmetry in solutions. We end this section with presenting the branching strategies.

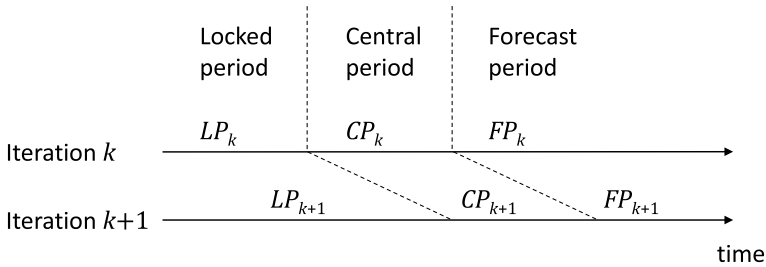


Fig. 3 Illustration of the different time periods in a rolling horizon heuristic

4.1 Rolling horizon heuristic

The complete LNG-ADP problem is usually too large to be solved by commercial solvers. A common approach to reduce the problem size for LNG-ADP problems is to reduce the planning horizon by means of a rolling horizon heuristic (RHH), see e.g., Rakke et al. (2011). The general concept of a RHH is illustrated in Fig. 3: In each iteration, the planning horizon is split into a locked period, a central period and a forecast period. Note that these three periods usually cover a period that is considerably shorter than the original planning horizon. In each iteration k , we inherit the already implemented decisions in the locked period, LP_k . We then solve the problem to find the optimal decisions in the central period, CP_k , and forecast period, FP_k . After the problem is solved for iteration k , we move the periods forward by the duration of CP_k and fix the decisions in CP_k as part of the new locked period LP_{k+1} for the new iteration $k + 1$. We repeat this process until the entire planning horizon belongs to the locked period.

With the discrete-time formulation, the decisions can be directly linked to each period through the time index t . All voyage variables, i.e., x_{vijt} and y_{bjte} , will be fixed in iteration $k + 1$ as long as the departure time t belongs to the central period in iteration k . Note that the delivery inequalities have to be updated in every iteration due to decisions taken in previous iterations.

A possible drawback of this method is that the generated solutions become too myopic. The configuration of RHH (i.e., the length of central and forecast periods) thus has an impact on the solution quality. We test and compare different configurations in Sect. 6.

4.2 Symmetry reduction

Due to the homogeneous fleet and the discrete-time formulation, the model has a large number of symmetric solutions that may slow down the solution process. We therefore slightly adjust voyage costs by a small amount to reduce symmetry and improve solution efficiency. Note that these adjustments are magnitudes smaller than cost parameters and therefore do not have an impact on the optimal decisions.

The voyage costs for owned carriers are adjusted based on two factors: first, we increase the voyage costs of owned LNG carriers, C_{vijt} based on their carrier

index v . Carriers with a larger index v are assigned a slightly higher sailing cost, i.e. $C_{vijt} < C_{v+1ijt}$, to reduce symmetry between carriers. We also slightly increase the voyage costs for a given carrier v based on departure time t during a month, i.e. $C_{vijt} < C_{vijt+1}$, to further reduce symmetry. This also encourages early loading/unloading during a month and makes vessels available for their next voyage as soon as possible. Note that the cost increase for the departure time is reset at the first of every month.

Similarly, we also increase the voyage costs for chartered carriers C_{bjte}^C by a small amount based on loading time t and unloading time e during a month.

4.3 Branching strategies

Often, branching decisions affecting larger sets of variables have a larger impact on the bound than decisions that only affect small sets of variables. We therefore aggregate certain voyages into sets of voyages to exploit this property. This idea has also been employed in other papers studying MIRPs, e.g., Song and Furman (2013) and Papageorgiou et al. (2014). We consider the following sets of voyages:

1. Total loads from the production port.
2. Total loads from the production port to the transshipment port.
3. Total loads from the transshipment port.

We introduce auxiliary integer variables n_i and add additional constraints of the following type:

$$\sum_{v \in \mathcal{V}} \sum_{j \in \mathcal{P}_{vi}^{AF}} \sum_{t \in \mathcal{T}} x_{vijt} + \sum_{b \in \mathcal{B}} \sum_{j \in \mathcal{P}_b^C} \sum_{t \in \mathcal{T}} \sum_{e \in \mathcal{T}_{bjt}^C} y_{bjte} = n_i \quad i \in \mathcal{P}^P \cup \mathcal{P}^T$$

We then force the solver to prioritize branching on variables n_i .

5 Case study

This section introduces the case study based on the Yamal LNG project. The input data is to a large degree based on the data set given in Li and Schütz (2020), but small changes may occur. The main features of the existing data set are included below for the sake of completeness.

5.1 Port information

We consider the distribution of LNG from a production facility in Sabetta on the Yamal peninsula in the Russian Arctic. Zeebrugge in Belgium is chosen as the transshipment port. The long-term customers and spot markets are located in both Asia and Europe. Figure 4 illustrates the locations of the ports, the availability of the possible routes and typical sailing times.

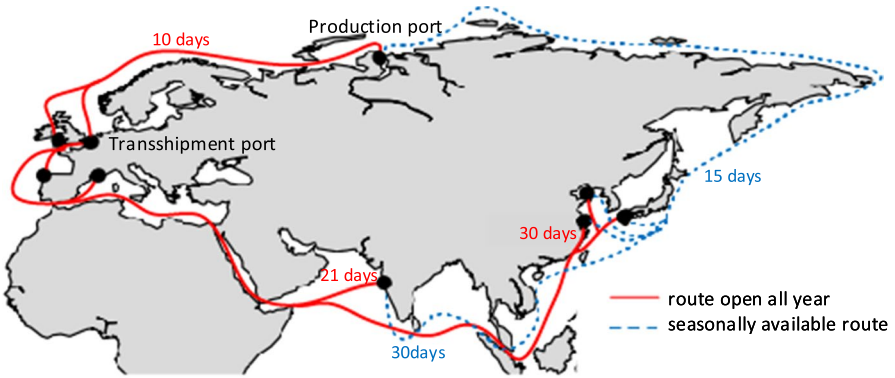


Fig. 4 Overview over the routes with typical sailing times to selected ports

Table 1 Information about the ports in the network

Port number	Type	Region	Connection		Distance [nm]	
			Direct	Transshipment	Direct	Transshipment
1	Production	–	–	–	–	–
2	Long-term	Asia	Seasonal	Yes	5959.7	11,067.2
3	Long-term	Asia	Seasonal	Yes	6054	11,160.6
4	Spot	Asia	Seasonal	Yes	5348.9	11,058.2
5	Long-term	Asia	Seasonal	Yes	10,268.2	10,955.9
6	Long-term	Asia	Seasonal	Yes	6010.9	10,862.2
7	Long-term	Europe	Yes	No	2784.9	–
8	Long-term	Europe	No	Yes	–	2080.2
9	Spot	Europe	No	Yes	–	741.7
10	Transshipment	–	Yes	–	2661.8	–

Table 1 summarizes the information about the ports included in the network. The type of each port is given in the second column. The third column lists the geographical region of the customers’ location. The fourth and fifth column indicate if a port can be reached from the production port directly or via the transshipment port. For Asian customers, the transshipment route is always available, whereas the direct route is assumed to only be available during the summer season. Dependent on the location of the European customer, she will be served either directly or via the transshipment port. Sailing distances are provided in columns 6 and 7 (taken from www.marinetraffic.com). At each port, there is only one berth. Loading and unloading operations take one day each.

5.2 Carrier information

The producer operates its own homogeneous fleet of 15 ice-going LNG carriers with a loading capacity of 172,600 m³. These carriers only load at the production port and serve either the transshipment port or a subset of the customers. Sailing speed, and thus sailing time, depends on the season. During the summer season, we assume that the ice-going carriers operate with an average speed of 18 knots on all voyages. During the winter season, we assume an overall average speed of 10.6 knots for westbound voyages from the production port, whereas we assume a speed of 9.5 knots when sailing in ice (along the Northeast Passage) and 16 knots when sailing in open water for vessels returning to Sabetta from the Asian customers. Note that east-bound voyages from Sabetta are not possible during the winter season. However, we do allow carriers to return to Sabetta during the winter season if they departed for an Asian customer during the summer season.

We assume that the ice-going vessels incur daily crew costs of 9,000 USD/day. Natural boil-off (see e.g., Hasan et al. 2009) is set to 0.11 %/day. We include a forced boil-off of 520 m³/day as fuel supply.

In addition to the ice-going carriers, there is a homogeneous fleet of open-water LNG carriers that are available for chartering whenever needed. These chartered LNG carriers have a loading capacity of 161,000 m³ and only load at the transshipment port. Their sailing speed also depends on the season and we use an average sailing speed of 19 knots during the summer and 18 knots during the winter season.

For chartered vessels we assume a daily charter cost of 60,000 USD/day and a natural boil-off of 0.12 %/day. Voyages from the transshipment port to Asian customers pass through the Suez Canal and hence, canal fees of 175,000 USD apply.

Please note that sailing times are determined using the information above, but the times used in the calculations have been rounded to full days due to the discrete-time model formulation.

5.3 Planning horizon

We consider a total planning horizon of 14 months. The planning horizon begins with six months of winter, followed by six summer months and two more winter months. We are interested in setting up an Annual Delivery Program for a period for 12 months, but we include an additional month both before and after the period of interest. The first month is used to initialize the problem, also reflecting the commitments from the previous Annual Delivery Program. We add the last month to allow voyages that start before the end of the year, but cannot be completed before, see Fig. 5.

5.4 Demand profiles and inventory management

We assume that demand to the long-term customers accounts for approximately 95% of the annual LNG production, with 54% being assigned to Asian customers while

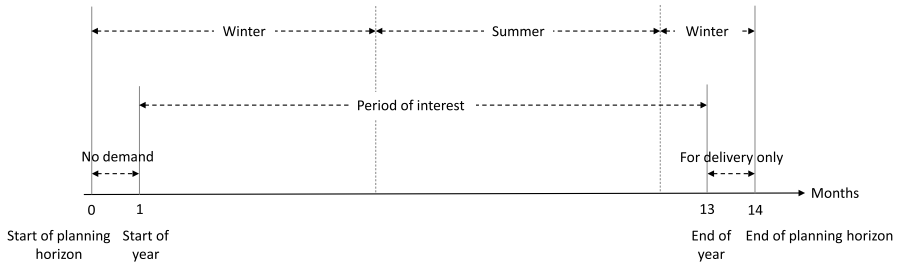


Fig. 5 Illustration of the planning horizon

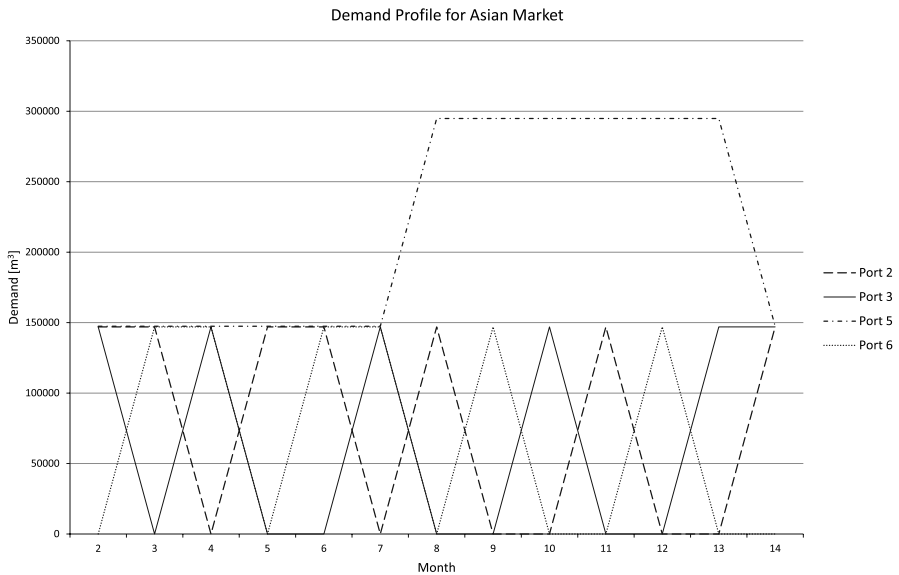


Fig. 6 Monthly demand for long-term customers in Asia

the European customers share the remaining 46%. We include seasonal changes in demand level for one Asian customer and both of the European customers. The different long-term demand profiles are illustrated in Figs. 6 and 7.

According to the principles described in Sect. 2, the monthly penalty for over-delivery is set at 45 USD/ m^3 and at 330 USD/ m^3 for under-delivery. We set the annual penalty for over-deliveries below one shipload to be 10 USD/ m^3 and 500 USD/ m^3 for over-deliveries above one shipload. For under-deliveries below one shipload, we set the penalty at 14 USD/ m^3 and at 700 USD/ m^3 for under-deliveries exceeding one shipload. Spot prices in Asia and Europe are given as 226 USD per m^3 and 160 USD per m^3 , respectively. We define “one shipload” based on the capacity of the ice-going LNG carriers, i.e. 172,600 m^3 , for all customers.

The inventory capacity at the production port is 640,000 m^3 with an initial inventory of 320,000 m^3 . We set the production rate to be 34,279 m^3 /day, corresponding

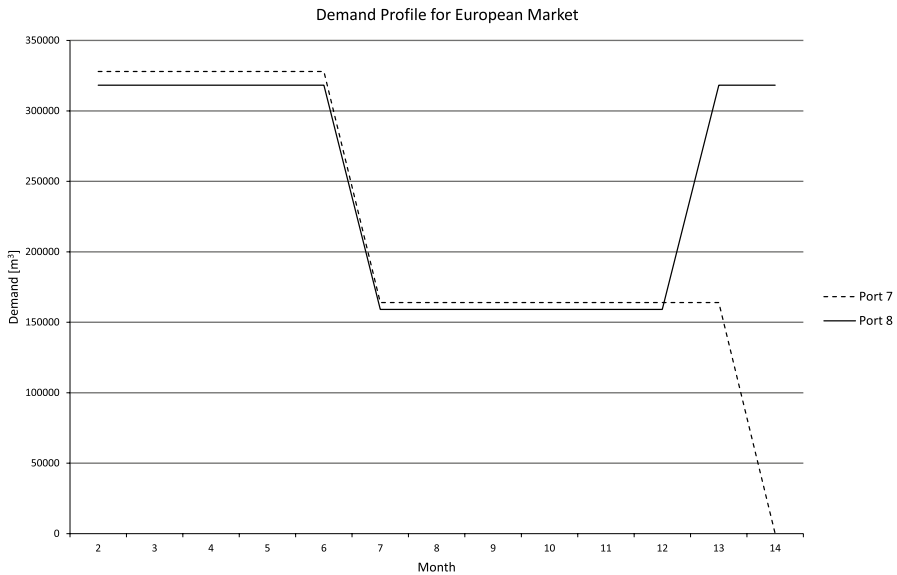


Fig. 7 Monthly demand for long-term customers in Europe

to an annual LNG production capacity of 5.8 MTPA for a single train. The inventory capacity at the transshipment port is set to 331,160 m³ with no initial inventory.

6 Computational results

This section presents the computational results for the case study described in Sect. 5. We first introduce our problem instances before presenting and discussing the results in more detail. The optimization model is implemented in Mosel and solved using FICO Xpress 8.10. All computations are carried out on Lenovo NextScale nx360 M5 computers with two 3.4 GHz Intel 6-core E5-2643 CPUs and 512 GB RAM running a Linux operating system.

Please note that Li and Schütz (2020) provide a continuous-time formulation for the LNG-ADP with transshipment. Even though this paper is inspired by the same case from the real world, voyages are modeled differently, different stopping criteria have been adopted and the cost data have been updated as compared to Li and Schütz (2020). The objective function values are therefore not directly comparable.

6.1 Problem instances

We create problem instances with different maximum waiting times, using upper limits of 0, 3, 5 and 7 days at the long-term customers to explore the effect of allowing waiting on the solution quality. Each of the waiting time instances is solved using different configurations for the RHH to study the impact of looking ahead on

Table 2 Combination of problem instances (W#) and RHH configuration (R#+#)

Combination	Waiting time	Central period	Forecast period
W0R1 + 1	0 days	1 month	1 month
W0R2 + 1	0 days	2 months	1 month
W0R1 + 2	0 days	1 month	2 months
W3R1 + 1	3 days	1 month	1 month
W3R2 + 1	3 days	2 months	1 month
W3R1 + 2	3 days	1 month	2 months
W5R1 + 1	5 days	1 month	1 month
W5R2 + 1	5 days	2 months	1 month
W5R1 + 2	5 days	1 month	2 months
W7R1 + 1	7 days	1 month	1 month
W7R2 + 1	7 days	2 months	1 month
W7R1 + 2	7 days	1 month	2 months

Table 3 Results for the different combinations

Combination	Obj. value (mill. USD)	Avg. opt. gap (%)	Run time (h)
W0R1 + 1	167.10	5.32	24.02
W0R2 + 1	165.29	3.30	8.10
W0R1 + 2	162.61	5.64	21.77
W3R1 + 1	165.76	5.11	20.04
W3R2 + 1	166.37	0.34	0.07
W3R1 + 2	163.89	2.06	4.15
W5R1 + 1	178.78	4.63	24.05
W5R2 + 1	166.31	0.38	0.11
W5R1 + 2	162.90	1.95	4.25
W7R1 + 1	178.78	4.78	24.19
W7R2 + 1	166.31	0.41	0.07
W7R1 + 2	163.95	0.45	4.17

the solution quality. We combine a duration of 1 and 2 months for the central period with durations of 1 and 2 months for the forecast period. All combinations of problem instance and RHH configuration are summarized in Table 2.

We allow at most four hours run time for each iteration of the RHH, but stop as soon as the optimality gap falls below 0.5%.

6.2 Solution quality

We provide the objective function value, the average optimality gap over all iterations and the total run time for the different combinations in Table 3. Most of the iterations are solved within 0.5% of optimality, but two can only be solved with an optimality gap as high as 21%.

We observe that the objective function value is on average lower for the configurations with a planning horizon of three months in the RHH (i.e., configurations R2 + 1 and R1 + 2) than for the R1 + 1 configurations with a two months planning horizon. This shows that a longer planning horizon in the RHH provides value in terms of lower costs. Note that the R1 + 2 configurations consistently provide the best solutions for the different waiting times. In addition, most of the R1 + 2 configurations have considerably shorter run times than the R1 + 1 configurations. Our results corroborate the findings by Papageorgiou et al. (2018), who report that a RHH with a central period of one month and a forecast period of two months outperforms other configurations for solving a group of MIRPs without transshipment.

The R1 + 1 and R1 + 2 configurations require the same number of iterations during the RHH to provide a complete ADP. Still, solving the R1 + 1 configurations takes almost six times longer than solving the R1 + 2 configurations. At the same, the R1 + 2 models are considerably larger than the R1 + 1 models as they include an additional month in the forecast period. This additional month seems to provide useful planning information that allows the solver to close the optimality gap much faster than for the smaller R1 + 1 problems. The shortest run times however come from the R2 + 1 configurations. These have approximately the same size as the R1 + 2 problems, but require only half the iterations to set up a complete ADP. The increase in speed however, comes at a slightly worse objective function value.

It is however interesting to see that introducing waiting to increase the shipping flexibility of the LNG producer does not have a consistent impact on the solution quality. In fact, the instances with the longest possible waiting time (W5 and W7) perform worst for the R1 + 1 configurations. Despite including all previous waiting times, they fail to provide a single best solution for any of the combinations. The increase in problem size due to additional waiting variables seems to affect the solution quality more negatively than the additional benefit from the increased solution space.

We therefore consider avoiding myopic planning horizons as important when configuring a RHH. In addition there might be a trade-off between run time and solution quality, even when different RHH configurations cover the same overall planning horizon, but split it differently between central period and forecast period.

6.3 Solution structure

We further examine the structure of the best solutions for the different waiting times, i.e., for the R1 + 2 configurations. In Table 4, we provide a breakdown of the objective function into sailing cost, deviation penalties (monthly and annual) and revenues from spot sales. We see that the different costs elements of the objective function are very similar across combinations, with relative difference being smaller than 1%.

The difference in the objective function for the solutions to these combinations is caused by variations in deviation penalties, mainly the monthly deviations. Looking further into the deliveries to the individual customers, we see that these deliveries only differ for customer port 5, see Table 5. While customer port 5 receives 18 shipments in all solutions, the routing changes, i.e., the number of shipments delivered directly or

Table 4 Details of the best known objective function values (mill.USD)

	WOR1 + 2	W3R1 + 2	W5R1 + 2	W7R1 + 2
Sailing cost	247.54	247.70	247.45	247.76
Monthly penalty	123.32	124.95	123.86	124.95
Annual penalty	3.98	3.47	3.81	3.47
Revenues	212.22	212.22	212.22	212.22
Obj.val	162.61	163.89	162.90	163.95

Table 5 Total deliveries per customer port

Combination	Route	Customer port								Sum
		2	3	4	5	6	7	8	9	
WOR1 + 2	Direct	2	3	6	5	3	17	-	-	36
	Transshipment	4	3	0	13	3	-	20	0	43
	Sum	6	6	6	18	6	17	20	0	79
W3R1 + 2	Direct	2	3	6	2	3	17	-	-	33
	Transshipment	4	3	0	16	3	-	20	0	46
	Sum	6	6	6	18	6	17	20	0	79
W5R1 + 2	Direct	2	3	6	4	3	17	-	-	35
	Transshipment	4	3	0	14	3	-	20	0	44
	Sum	6	6	6	18	6	17	20	0	79
W7R1 + 2	Direct	2	3	6	2	3	17	-	-	33
	Transshipment	4	3	0	16	3	-	20	0	46
	Sum	6	6	6	18	6	17	20	0	79

through the transshipment port. We also note that the solution with the highest number of direct shipments has the lowest overall costs. Customer port 5 is a port that is located such that it is almost as far to sail from the production port to the customer as it is from the transshipment port. That a clearly dominant route cannot be established could be an explanation for the combinations producing quite different solutions for this particular port.

The detailed breakdown of the objective function also confirms once more that introducing the possibility of waiting does not improve the quality of the solutions for our case, despite that the solution space increases. Here, it should be noted that the reason for this might be that we are using a heuristic and that the solutions shown are not the optimal ones. On the other hand, and quite surprisingly, introducing the waiting results in an overall reduction in run time compared to the instances where waiting is not allowed.

7 Concluding remarks

We have in this paper considered a special version of the LNG-ADP problem, which is inspired by the case of Yamal LNG. Since Yamal's liquefaction plan is located in the Russian Arctic, ice-going LNG carriers are needed to ensure year-round distribution of LNG. These carriers can serve Asian and European customers directly. However, if the eastbound route along the North East Passage becomes unavailable, the carriers must sail to the European Continent where the LNG for the Asian customers is transshipped. In that case, the LNG is unloaded and stored at the transshipment port before being loaded onto chartered open-water LNG carriers that are serving the Asian customers, sailing through the Suez Canal.

We propose a novel discrete-time mixed integer programming formulation for the LNG-ADP problem with transshipment. Since the model is too large and complex to be solved using a commercial solver, we implement a rolling horizon heuristic (RHH) to solve the LNG-ADP problem for a case based on the Yamal LNG project. We test different configurations for the RHH and find that a myopic configuration for the RHH can provide good results, but also has the longest run time. Increasing the duration of the forecast period produces better solutions and considerably decreases run time, even with an increased problem size. Introducing waiting into the formulation does not necessarily improve the objective function value, despite increasing the solution space. Still, we observe a reduction in run time for the instances where waiting is allowed.

LNG-ADP problems are known to be difficult to solve and heuristics are therefore frequently used. But most of the existing heuristics developed and tested for the LNG-ADP problem do not consider transshipping LNG. Studying the need for new and/or improved heuristics to deal with transshipment is subject to future research. Another area for future research stems from the fact that LNG-ADP problems are traditionally formulated and solved as deterministic problems. However, maritime transportation problems are often characterized by uncertainty in sailing times. Incorporating uncertainty in the LNG-ADP will increase the complexity of the problem, but may also better reflect some of the challenges encountered in the real world. Considering uncertainty can therefore give rise to new model formulations and new solution methods for the LNG-ADP problem.

Funding Open access funding provided by NTNU Norwegian University of Science and Technology (incl St. Olavs Hospital - Trondheim University Hospital). The authors did not receive support from any organization for the submitted work.

Declarations

Conflict of interest The authors have no relevant financial or non-financial interests to disclose.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the

material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Agra A, Andersson H, Christiansen M, Wolsey L (2013) A maritime inventory routing problem: discrete time formulations and valid inequalities. *Networks* 62(4):297–314
- Agra A, Christiansen M, Delgado A (2017) Discrete time and continuous time formulations for a short sea inventory routing problem. *Optim Eng* 18(1):269–297
- Al-Haidous S, Msakni MK, Haouari M (2016) Optimal planning of liquefied natural gas deliveries. *Transp Res Part C Emerg Technol* 69:79–90
- Andersson H, Christiansen M, Desaulniers G, Rakke JG (2017) Creating annual delivery programs of liquefied natural gas. *Optim Eng* 18(1):299–316
- Christiansen M, Fagerholt K (2014) Ship routing and scheduling in industrial and tramp shipping. In: Toth P, Vigo D (eds) *Vehicle routing: problems, methods, and applications*, 2nd edn. Society for Industrial and Applied Mathematics, Philadelphia, pp 381–408
- Christiansen M, Fagerholt K, Nygreen B, Ronen D (2013) Ship routing and scheduling in the new millennium. *Eur J Oper Res* 228(3):467–483
- Coelho LC, Cordeau JF, Laporte G (2012) The inventory-routing problem with transshipment. *Comput Oper Res* 39(11):2537–2548
- GIIGNL (2021) GIIGNL Annual Report 2020. The International Group of Liquefied Natural Gas Importers, Neuilly-sur-Seine, France
- Goel V, Furman KC, Song JH, El-Bakry AS (2012) Large neighborhood search for LNG inventory routing. *J Heuristics* 18(6):821–848
- Goel V, Slusky M, van Hoeve WJ, Furman KC, Shao Y (2015) Constraint programming for LNG ship scheduling and inventory management. *Eur J Oper Res* 241(3):662–673
- Halvorsen-Weare EE, Fagerholt K (2013) Routing and scheduling in a liquefied natural gas shipping problem with inventory and berth constraints. *Ann Oper Res* 203(1):167–186
- Hasan MMF, Zheng AM, Karimi IA (2009) Minimizing boil-off losses in liquefied natural gas transportation. *Ind Eng Chem Res* 48(21):9571–9580
- Li M, Schütz P (2020) Planning annual LNG deliveries with transshipment. *Energies* 13(6):1490
- Lianes IM, Noreng MT, Fagerholt K, Slette HT, Meisel F (2021) The aquaculture service vessel routing problem with time dependent travel times and synchronization constraints. *Comput Oper Res* 134:105316
- Mutlu F, Msakni MK, Yildiz H, Sönmez E, Pokharel S (2016) A comprehensive annual delivery program for upstream liquefied natural gas supply chain. *Eur J Oper Res* 250(1):120–130
- Papageorgiou DJ, Cheon MS, Harwood S, Trespalacios F, and Nemhauser GL (2018). Recent progress using mathheuristics for strategic maritime inventory routing. In: *intelligent systems reference library*, vol 131. Springer, Cham, pp. 59–94
- Papageorgiou DJ, Cheon MS, Nemhauser G, Sokol J (2015) Approximate dynamic programming for a class of long-horizon maritime inventory routing problems. *Transp Sci* 49(4):870–885
- Papageorgiou DJ, Keha AB, Nemhauser GL, Sokol J (2014) Two-stage decomposition algorithms for single product maritime inventory routing. *INFORMS J Comput* 26(4):825–847
- Papageorgiou DJ, Nemhauser GL, Sokol J, Cheon MS, Keha AB (2014) MIRPLib - a library of maritime inventory routing problem instances: survey, core model, and benchmark results. *Eur J Oper Res* 235(2):350–366
- Peres IT, Repolho HM, Martinelli R, Monteiro NJ (2017) Optimization in inventory-routing problem with planned transshipment: A case study in the retail industry. *Int J Prod Econ* 193:748–756
- Rakke JG, Stålhane M, Moe CR, Christiansen M, Andersson H, Fagerholt K, Norstad I (2011) A rolling horizon heuristic for creating a liquefied natural gas annual delivery program. *Transp Res Part C Emerg Technol* 19(5):896–911
- Schach M, Madlener R (2018) Impacts of an ice-free northeast passage on LNG markets and geopolitics. *Energy Policy* 122:438–448
- Shao Y, Furman KC, Goel V, Hoda S (2015) A hybrid heuristic strategy for liquefied natural gas inventory routing. *Transp Res Part C Emerg Technol* 53:151–171

-
- Song JH, Furman KC (2013) A maritime inventory routing problem: practical approach. *Comput Oper Res* 40(3):657–665
- Stålhane M, Rakke JG, Moe CR, Andersson H, Christiansen M, Fagerholt K (2012) A construction and improvement heuristic for a liquefied natural gas inventory routing problem. *Comput Ind Eng* 62(1):245–255

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Mingyu Li received his PhD degree in Industrial Economics and Technology Management from the Norwegian University of Science and Technology. His research focuses on the impact of the Northern Sea Route on maritime transportation. He is currently working on decision support systems for planning LNG shipments for Quorum Software.

Kjetil Fagerholt is a Professor in Operations Research at the Norwegian University of Science and Technology, Department of Industrial Economics and Technology Management. He also holds a position as an Adjunct Professor at the Department of Marine Technology at the same university. His main research interest is on developing and applying Operations Research models and methods on planning problems in maritime transportation.

Peter Schütz is an Associate Professor in Managerial Economics and Optimization at the Norwegian University of Science and Technology, Department of Industrial Economics and Technology Management. His main research interest is decision-making under uncertainty with emphasis on applications from supply chain management, maritime transportation and quantitative logistics in general.