Noname manuscript No.
(will be inserted by the editor)

Planning interrelated voyages with separation requirements in roll-on roll-off shipping

Jone R. Hansen[1] · Kjetil Fagerholt[1] · Frank Meisel[2] · Jørgen G. Rakke[3] .

Received: date / Accepted: date

Abstract We consider a new problem of planning interrelated voyages with separation requirements along a single trade in roll-on roll-off shipping. Along a given trade with a sequence of port calls, there is a number of contracts for transportation of cargoes between the different port pairs, where each contract states its service frequency as well as that these services should be evenly separated in time. Instead of visiting all ports every time a trade is serviced, as closely resembles current planning practice, we aim at determining the sailing routes of each voyage along the trade, i.e. which ports to visit when, which contracts to serve, and the sailing speeds, so that all contract requirements regarding frequency and separation are satisfied at minimum cost. We propose and compare two novel mixed integer programming models for the problem, both including a new way of modeling the contracts' separation requirements. Then we show through a computational study on a set of realistic test instances that significant gains can be obtained compared to current planning practice.

Keywords Maritime transportation · Roll-on roll-off shipping · Separation requirements · Speed optimization · Transit time constraints

1 Introduction

In maritime transportation it is common to distinguish among three, not necessarily mutually exclusive, modes of operation: *industrial*, *tramp* and *liner*. In industrial shipping, the shipper controls the fleet of ships, trying to minimize the cost of transporting its cargoes, similar to a private fleet. In a tramp operation the ships follow the available cargoes, some of which may be optional from the spot market,

J. R. Hansen
E-mail: jone.hansen@ntnu.no

[1] Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology, Norway
[2] Faculty of Business, Economics and Social Sciences, Christian-Albrechts-University Kiel, Germany
[3] Wallenius Wilhelmsen Logistics, Lysaker, Norway

trying to maximize profit, similar to a taxi service. Bulk ships usually operate in one of these two modes. In liner shipping, the ships operate on more or less fixed services or trades according to a published schedule. See Christiansen et al (2013) for a general review on ship routing and scheduling for the various modes and Meng et al (2014) for a specific one on liner shipping.

In this paper we focus on a planning problem arising in the roll-on roll-off (RoRo) shipping segment, which is normally categorized in the liner shipping mode. In RoRo-shipping, a wide variety of rolling cargoes, such as cars, trucks, and heavy rolling machinery, is transported. Additionally, RoRo vessels are designed to carry complex cargoes that are placed on trolleys and rolled on and off the vessels, such as turbines, yachts, and windmill blades. RoRo-shipping is an important segment with a world fleet of around 5000 vessels with a total capacity of more than 24 million deadweight tons (ISL 2017).

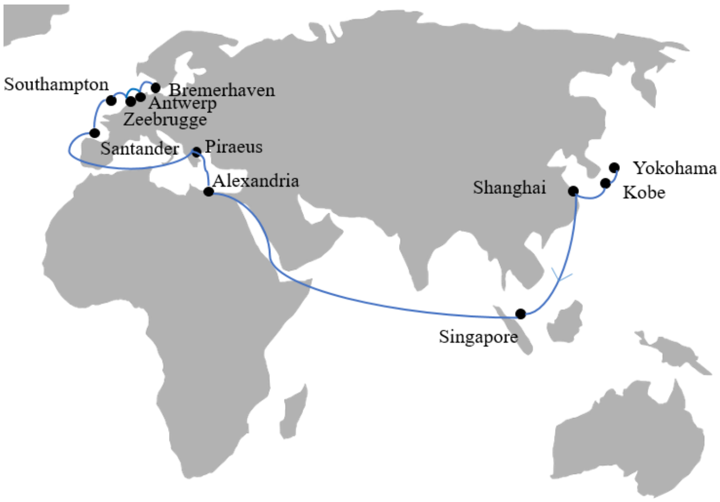

Fig. 1 Example of a trade from Asia to Europe.

We study the single trade ship routing and scheduling problem (STSRSP) in RoRo-shipping, which is the problem of planning the voyages to fulfill the demand for transportation along a single trade for a given planning period (e.g. a month). A *trade* connects two geographical regions, such as Asia and Europe as exemplified in Figure 1, where each region contains a set of ports that are called upon for loading and unloading the cargoes. A vessel that is deployed on the trade performs a *voyage*, i.e. it visits a sequence of port along the trade, see Figure 1. For each trade a number of voyages must be performed during the planning period.

The trades to be serviced in liner shipping are usually designed based on a large number of contracts for transportation of cargoes between the different port pairs along a trade, e.g. from Shanghai to Santander along the Asia - Europe trade shown in Figure 1. The trades are directed, such that Asia-Europe is the trade from Asia to Europe and the Europe-Asia trade covers the opposite direction. Each contract states a given total quantity to be transported during the planning horizon. This quantity should be distributed among a given number of services (i.e. with some given frequency), which must also be fairly evenly spread or separated in time. This means that, for a given contract, vessels must visit the corresponding loading port (e.g. Shanghai) to pick up cargoes under that contract with the regularity as stated by the contract and transport that cargo to the corresponding unloading port (e.g. Santander). In contrast to container liner shipping, transshipment between vessels and trades is rarely performed in RoRo-shipping due to the time consuming loading and unloading operations, and is therefore disregarded in this study.

In container liner shipping, the port visit regularity is easily achieved as each trade is usually serviced on a strict weekly basis and each voyage along the trade includes all ports in the same order, see for example Brouer et al (2013), Ng (2015) and Wang and Meng (2017). However, RoRo-shipping entails more planning flexibility as both to when to start each voyage as well as when and how often to visit each port along the trade. For example, even though a trade, such as the one in Figure 1, may have a weekly frequency, it does not have to be exactly seven days between the voyages, but just seven days on average. In previous studies on fleet deployment in RoRo-shipping, one has therefore used time windows for when each voyage along each trade should start, see for example Fagerholt et al (2009), Andersson et al (2015) and Fischer et al (2016). One problem with using time windows, and especially if they are wide, is that one might obtain solutions where a given voyage starts at the beginning while the subsequent one on the same trade starts at the end of their time windows. This is not desirable from the customers' perspective, as they want their cargoes to be shipped out at relatively regular intervals, i.e. *fairly evenly spread* in time. Bakkehaug et al (2016), Norstad et al (2015) and Vilhelmsen et al (2017) overcome this, though for a problem from another shipping segment than RoRo-shipping, by imposing voyage separation constraints to make sure that the voyages on the same trade are evenly spread in time. However, this only makes sure that the starting times for each voyage along a given trade are evenly spread. This can be sufficient in the case, as for the previously mentioned studies, where the voyages are sailed in the exact same way every time, i.e. visiting the same ports in the same order and sailing with the same speed.

In the STSRSP we aim at utilizing the inherent planning flexibility in RoRo-shipping, also when it comes to when and how often to visit each port along a given trade. Depending on the frequency and spread requirements of each cargo contract, one might not need to visit each port on every voyage. Using the example from Figure 1; if the combined requirement for all contracts with cargoes to be picked up and delivered in Shanghai specifies that a visit to that port must be done only every second week, one might not need to visit that port along every voyage, even though the trade itself needs to be serviced with a weekly frequency. Therefore, instead of aggregating requirements from contracts to frequency requirements for the whole trade, which has been common to simplify planning and as done in the previous studies, we look at the frequency requirements for each contract and port along the trade. This is further utilized in the STSRSP to determine the sailing routes of each voyage along the trade, i.e. which ports to be visited, what contracts to serve, as well as the sailing speeds along the voyages, so that all contract requirements regarding frequency and spread are satisfied at minimum cost. This can also be seen as taking some of the flexibility one finds in tramp shipping regarding deciding which ports to visit and which cargoes to transport and combine this with operating on given trades as in liner shipping.

The STSRSP is a problem that, to the best authors' knowledge, is new to the research literature. However, in addition to the literature discussed above, the STSRSP has some similarities with the periodic vehicle routing problem (e.g. Campbell and Wilson (2014)) and the supply vessel planning problem (e.g. Kisialiou et al (2018) and Borthen et al (2017)), as well as the special liner shipping network design problem considered by Sigurd et al (2005), in the sense that one needs to separate the services for each customer in time. In the STSRSP, we also consider that some contracts may impose transit time limits like we also see in some studies on container liner shipping, e.g. Karsten et al (2015), Reinhardt et al (2016) and Wang et al (2013).

Our main contributions are twofold. Firstly, we introduce two novel mixed integer programming (MIP) models for the STSRSP. A central part of both MIP models includes a new way of modeling the fairly evenly spread (or separation) requirements for the contracts. We show that the less intuitive model outperforms the other one. Secondly, we demonstrate through a computational study based on data mostly from the case company the potential gains that can be achieved from utilizing the inherent planning flexibility. Furthermore, we discuss the trade-offs between service level with regards to the contract separation requirements and the costs of a solution.

A detailed description of the problem is given in Section 2, while the two MIP models we propose are presented in Section 3. Section 4 provides the computational study, while concluding remarks are given in Section 5.

2 Problem definition and some modeling assumptions

The STSRSP is the problem of planning the interrelated voyages along a single trade, i.e. it consists of determining the routes and schedules for the vessels to fulfill the demand for transportation in the next planning period. Let $\mathcal{N}$ be the set of all nodes in the network, and $\mathcal{N}^P$ be the set of ports along the trade route. Let $\mathcal{K}$ be the set of available vessels. Each vessel k has a starting position $o(k)$, an artificial ending position $d(k)$, and a graph $G_k = (\mathcal{N}_k, \mathcal{A}_k)$ associated with it. The set of nodes $\mathcal{N}_k$ consists of all ports that can be visited by vessel k, i.e. $\mathcal{N}_k^P \subseteq \mathcal{N}^P$, in addition to its starting position and artificial ending position, $\mathcal{N}_k = \mathcal{N}_k^P \cup \{o(k), d(k)\}$. The set of arcs $\mathcal{A}_k \subset \mathcal{N}_k \times \mathcal{N}_k$ defines the feasible movements for vessel k, while the set $\mathcal{A} \subset \mathcal{N} \times \mathcal{N}$ defines all feasible movements. In RoRo-shipping, the ports are most often serviced in a directed order, such that if ports i_1 and i_2 are both visited on a given voyage, port i_1 will always be visited before port i_2 (as illustrated in Figure 1). Due to the inherent geographical structure of the trades, the graphs G_k therefore become directed and acyclic.

Let $\mathcal{C}$ be the set of given contracts for cargoes to be transported along the trade during the planning horizon. Most contracts are inter-regional, which means that a given contract's cargoes are to be loaded at a specific port in one geographical region (e.g. Asia) and unloaded in another (e.g. Europe). However, some contracts can also be intra-regional, i.e. both the loading and unloading ports are in the same region. Each contract c is a transportation arrangement for a certain product type. Let $\mathcal{P}$ be the set of product types, and $\mathcal{P}_p^S$ be the set of product types that can be stored in the same space as product type p. E.g. cars can be stored on decks facilitated for storing breakbulk, but not the other way around. See for example

Pantuso et al (2016) for more details. Further, let K^V_{kp} denote the capacity for product p on ship k. The demand for the whole planning period for product type p for contract c (measured in square meters) is given by D_{cp}, where each contract contains only one product type. Each contract c has a given loading port $l(c)$ and a given unloading port $u(c)$ associated with it. Let $\mathcal{C}^L_i$ and $\mathcal{C}^U_i$ be the sets of cargoes that are to be loaded and unloaded at port i, respectively. Each time that contract c is serviced, i.e. picked up by a vessel, the quantity picked up must be within the interval $[\underline{Q}_{cp}, \bar{Q}_{cp}]$. These limits ensure that impractical pickups are omitted, such as e.g. loading 15 cars on one and 985 cars on another vessel voyage, fulfilling a total demand of 1000 cars. The shipping company is committed to service all contracts, but the total quantity of each contract can be split among multiple voyages. Finally, let $\mathcal{C}^T$ denote the set of contracts with transit time restrictions, where the maximum transit time of a corresponding contract c is given by T^T_c.

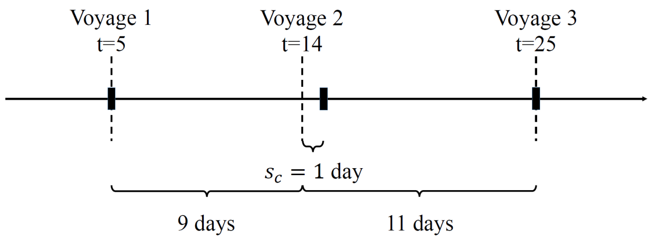

Fig. 2 A contract is to be serviced three times a month (30 days). Three voyages serve the contract, where the time of each service is shown in the figure. The desired spread of ten days between each pickup is illustrated by the black squares. The slack variable s_c gives the maximum deviation from the desired spread, i.e. one day in this example.

Some of the transportation demand arises from contracts that require their pickups to be fairly evenly spread or separated throughout the planning horizon. Let $\mathcal{C}^E$ be the set of such contracts. Each such contract c specifies a lower and an upper limit on the number of pickups within the planning horizon, i.e. the number of partial cargoes the contract may be split into, which is a decision to be made within the STSRSP. Let $\underline{P}_c$ and $\bar{P}_c$ denote these lower and upper pickup limits, respectively. Furthermore, the partial cargoes should then be fairly evenly spread or separated in time. From a modeling perspective, it is hardly possible to model this using time windows, as the number of partial cargoes is not given a priori. Instead, we have implemented the fairly evenly spread requirements using separation constraints. The desired spread for the partial cargoes for a given contract c is given by T^{PH}/b_c, where T^{PH} is the point in time up to which all voyages along the trade must begin, i.e. the length of planning horizon, and $b_c \in [\underline{P}_c, \bar{P}_c]$ is the number of partial cargoes chosen, i.e. the number of pickups. Given a planning horizon of 30 days and three partial cargoes, the desired spread is ten days, as illustrated in Figure 2. As shown in the example, a perfect spread could have the pickup times for the three partial cargoes at days 5, 15, and 25. Note that servicing a contract on days 0, 15, and 30, is not recognized as evenly spread, as it would generate problems when rolled out over several planning horizons.

Thus, two vessel voyages that both service partial cargoes from contract c comply with the evenly spread requirements if the following constraint is satisfied:

$$\frac{T^{PH}}{b_c} - s_c \leq t_{l(c)m} - t_{l(c)k} \leq \frac{T^{PH}}{b_c} + s_c$$

where s_c is a spread slack variable for contract c, which gives the maximum deviation from the desired spread for each contract. The variables $t_{l(c)m}$ and $t_{l(c)k}$ give the times at the pickup port of contract c for vessels m and k, respectively. These separation constraints are non-linear, since b_c is a variable. They should only be active if both vessels m and k carry contract c, and if m is the next vessel after k, servicing contract c. In the example in Figure 2, the slack variable must be at least 1 to satisfy the constraint. Alternatively, one could delay vessel voyage 2's visit to this given port with one day, but that may increase other costs, such as the charter time costs. The sum of s_c for all contracts reflects the evenly spread service level at company level, which is bounded by a preset threshold L. Here, a lower threshold gives a higher service level, and vice versa.

There are costs associated with sailing a voyage along the trade, i.e. port visit costs, fuel costs, and time charter costs. Let C_i^V be the cost of calling port i. The piecewise linear approximation method proposed by Andersson et al (2015) is used for modeling the speed-dependent fuel consumption for each vessel. Let $\mathcal{S}$ be the set of discrete speed alternatives indexed by s, ordered from low to high. The cost of sailing from a node (port) i to node (port) j for vessel k using speed alternative s is denoted C_{ijks}^S, where the corresponding sailing time is given by T_{ijks}^S. The vessels can be available at the start of the planning horizon or become available during the planning horizon, due to duties on other trades. Let T_k^A be the time vessel k becomes available at its origin. Furthermore, let T_{ik}^P represent the time used on piloting at port i by ship k. The time used to handle, i.e. load or unload, one unit of product type p on vessel k is given by T_{kp}^H. Let C_k^C be the daily charter rate for vessel k. The total time charter cost for performing a voyage with a given vessel k is given by the vessel's daily charter rate multiplied by the number of days spent on the voyage, which again depends on the port times and the speed-dependent sailing times between the ports chosen to be visited along the voyage. A list of all notations used in the models is given in the Appendix.

Figure 3 a) shows a small example of the STSRSP for a trade from US to Japan with five ports and with only three contracts (only for illustrational purposes, as a realistic number of contracts is much larger). Contract 1 has a total demand for the planning period (assumed to be one month) of 1000 units, which can be split in two, three or four partial cargoes to be loaded in Baltimore and unloaded in Yokohama. Contract 2 from Tacoma to Yokohama has a demand of 1500 and must be serviced three or four times, while contract 3, which is an intra-regional contract from Baltimore to Port Hueneme (both ports in the US), has a demand of 300 to be split in one or two partial cargoes.

Figure 3 b) shows a possible solution to the example problem, which consists of three vessel voyages along the trade. Voyage 1 visits only two ports, starting in Tacoma on day 7 and ending in Yokohama on day 19. Voyage 2 starts in Baltimore on day 2, then visiting Port Hueneme, Tacoma and Yokohama on days 14, 17 and 29, respectively, while voyage 3 starts in Baltimore on day 15 and visits Tacoma on day 28, before ending up in Yokohama on day 39. In this solution, contract 1

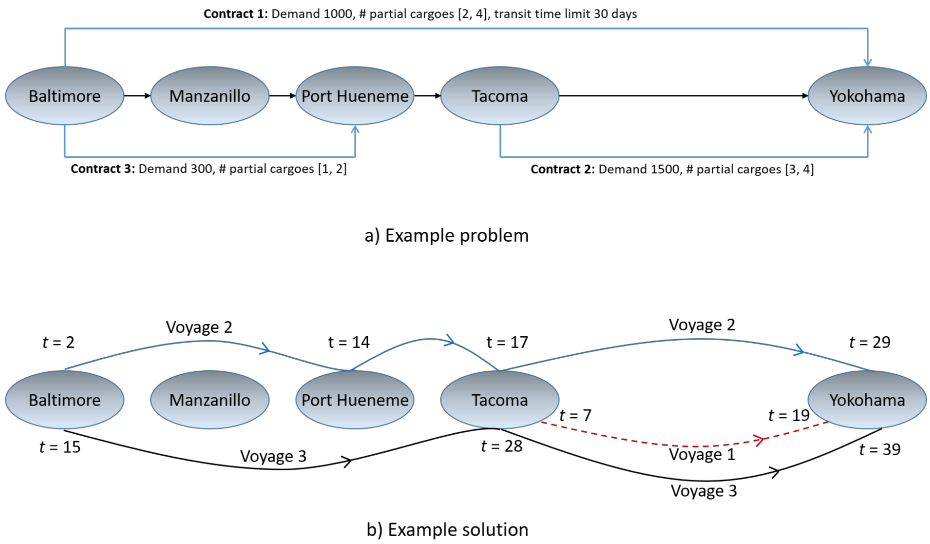

Fig. 3 An example of the STSRSP for a trade from US to Japan with five ports and only three contracts (a) and parts of a possible solution to the problem (b).

(from Baltimore to Yokohama) is serviced twice, i.e. on days 2 and 15, contract 2 (from Tacoma to Yokohama) is serviced three times, i.e. on days 7, 17 and 28, while contract 3 (from Baltimore to Port Hueneme) is serviced once on day 2. We can see that the services of each of the contracts are fairly evenly spread. Evaluating the evenly spread requirements for contract 2, we understand that the desired spread is 10 days, as it is serviced three times. From the solution, we see that the first two pickups are perfectly spread, with pickup on days 7 and 17. The second pickup and the third pickup are on days 17 and 28, respectively. Here, the spread is 11, which is one day off the desired spread. Thus, for contract 2, the spread slack variable equals one day.

In this example there is no contract associated with the port Manzanillo (which is unrealistic for a practical case as long as it is included in the trade). It is therefore no need to visit that port on any of the voyages along this trade. It should also be emphasized that we have neither shown how each contract's demand is distributed among the voyages nor the chosen speeds used on the different sailing legs for the voyages, which are also important decisions that have to be made in the STSRSP.

So to summarize, the objective of the STSRSP is to minimize the total cost while satisfying all transportation and service requirements within the planning horizon. The decisions to be made are which vessels to use, the routing (i.e. which ports to visit along each voyage), the assignment of contracts to vessels, the size of each partial cargo to transport at each port visit, as well as the sailing speeds for all sailing legs along each voyage (i.e. the schedule).

3 Mathematical models

In this section, we present two arc-flow models representing the STSRSP. The first model, named *Vessel-model*, is based on the typical ship routing and scheduling formulation for industrial and tramp shipping, see e.g. Christiansen et al (2007), with some notable differences. Firstly, in our approach, each node represents a physical port and not a pickup or delivery of a cargo, as is commonly used in the literature. Secondly, time window constraints are replaced by transit time constraints. Finally, separation constraints are added to ensure that the partial cargoes for each contract are evenly spread. The second model, referred to as the *Voyage-model*, is an alternative formulation of the STSRSP, which generates the sailing routes and schedules for each voyage and assign vessels to them. Both models represent the exact same problem, though modeled in two different ways. The key difference between the two formulations is how the arc-flow variables are defined. In the Vessel-model, we use decision variables to decide whether a vessel sails an arc or not, which is the most intuitive way to model the problem. In the Voyage-model, the corresponding decision variables are connected to voyages, i.e. if an arc is used on a voyage or not. Additional decision variables are included to assign a vessel to each voyage.

The explicit use of voyages reduces the Voyage-model's readability to some extend and increases the overall number of constraints and variables in the model. On the other hand, the voyage ordering enables symmetry breaking and eases the modeling of the evenly spread requirements. In Section 4.2, we continue the comparison of the models, based on the computational results. Sections 3.1 and 3.2 present the Vessel and the Voyage models, respectively.

3.1 Vessel-model

In order to present the Vessel-model we define the following additional notation. The binary decision variable x_{ijk} defines whether vessel k sails from node i to node j or not. The variable w_{ijks} represents the weight of speed alternative s for vessel k on the arc (i, j). As explained in Section 2, we use the linear approximation method proposed by Andersson et al (2015) to model the speed. In this method, the speed is determined through the use of continuous speed variables representing the weight of each discrete speed alternative. This means that a vessel k sailing arc (i, j) could use a sailing speed which results from an interpolation of two discrete speed alternatives. Suppose speed alternatives 1 and 2 are 14 and 16 knots, respectively. If $w_{ijk1} = 0.4$ and $w_{ijk2} = 0.6$ in the solution, vessel k sails arc (i, j) with a sailing speed of $14 \cdot 0.4 + 16 \cdot 0.6 = 15.2$ knots. The variable l_{ijkp} equals the load of product type p on vessel k on the arc (i, j). Furthermore, binary variable δ_{kc} is 1 if vessel k serves contract c, 0 otherwise, and q_{kcp} represents the quantity of product p included in contract c that is picked up by vessel k. The time variable t_{ik} defines the start of service at node i for vessel k.

With this notation, the STSRSP can be formulated as follows:

$$\min z = \sum_{k\in\mathcal{K}} \sum_{(i,j)\in\mathcal{A}_k} \sum_{s\in\mathcal{S}} C^S_{ijks} w_{ijks} + \sum_{k\in\mathcal{K}} \sum_{(i,j)\in\mathcal{A}_k} C^V_i x_{ijk} + \sum_{k\in\mathcal{K}} C^C_k (t_{d(k)k} - t_{o(k)k}) \tag{1}$$

$$\sum_{j \in \mathcal{N}_k^P \cup d(k)} x_{o(k)jk} = 1, \qquad \forall k \in \mathcal{K} \tag{2}$$

$$\sum_{i \in \mathcal{N}_k} x_{ijk} - \sum_{i \in \mathcal{N}_k} x_{jik} = 0, \qquad \forall k \in \mathcal{K}, j \in \mathcal{N}_k^P \tag{3}$$

$$\sum_{i \in \mathcal{N}_k^P \cup o(k)} x_{id(k)k} = 1, \qquad \forall k \in \mathcal{K} \tag{4}$$

$$x_{ijk} = \sum_{s \in \mathcal{S}} w_{ijks}, \qquad \forall k \in \mathcal{K}, (i,j) \in \mathcal{A}_k \tag{5}$$

$$0 \le l_{ijkp} \le K_{kp}^V x_{ijk} - \sum_{p' \in \mathcal{P}_p^S} l_{ijkp'}, \qquad \forall k \in \mathcal{K}, (i,j) \in \mathcal{A}_k, p \in \mathcal{P} \tag{6}$$

$$\sum_{j \in \mathcal{N}_k} l_{jikp} + \sum_{c \in \mathcal{C}_i^L} q_{kcp} - \sum_{c \in \mathcal{C}_i^U} q_{kcp} = \sum_{j \in \mathcal{N}_k} l_{ijkp}, \qquad \forall k \in \mathcal{K}, i \in \mathcal{N}_k^P, p \in \mathcal{P} \tag{7}$$

$$\sum_{j \in \mathcal{N}_k^P} l_{o(k)jkp} = 0, \qquad \forall k \in \mathcal{K}, p \in \mathcal{P} \tag{8}$$

$$\underline{P}_c \le \sum_{k \in \mathcal{K}} \delta_{kc} \le \bar{P}_c, \qquad \forall c \in \mathcal{C}^E \tag{9}$$

$$\delta_{kc} \le \sum_{i \in \mathcal{N}_k^P \cup o(k)} x_{il(c)k}, \qquad \forall k \in \mathcal{K}, c \in \mathcal{C} \tag{10}$$

$$\delta_{kc} \le \sum_{i \in \mathcal{N}_k^P} x_{iu(c)k}, \qquad \forall k \in \mathcal{K}, c \in \mathcal{C} \tag{11}$$

$$\underline{Q}_{cp} \delta_{kc} \le q_{kcp} \le \bar{Q}_{cp} \delta_{kc}, \qquad \forall k \in \mathcal{K}, c \in \mathcal{C}, p \in \mathcal{P} \tag{12}$$

$$\sum_{k \in \mathcal{K}} q_{kcp} = D_{cp}, \qquad \forall c \in \mathcal{C}, p \in \mathcal{P} \tag{13}$$

$$t_{o(k)k} = T_k^A, \qquad \forall k \in \mathcal{K} \tag{14}$$

$$t_{ik} + T_{ik}^P x_{ijk} + \sum_{c \in \mathcal{C}_i^L \cup \mathcal{C}_i^U} \sum_{p \in \mathcal{P}} T_{kp}^H q_{kcp} + \sum_{s \in \mathcal{S}} T_{ijks}^S w_{ijks} \le t_{jk}, \qquad \forall k \in \mathcal{K}, (i,j) \in \mathcal{A}_k \tag{15}$$

$$t_{u(c)k} - t_{l(c)k} \le T_c^T + M_c^T (1 - \delta_{kc}) \qquad \forall k \in \mathcal{K}, c \in \mathcal{C}^T \tag{16}$$

$$t_{jk} - M_{jk}^S (1 - x_{o(k)jk}) \le T^{PH}, \qquad \forall k \in \mathcal{K}, j \in \mathcal{N}_k^P \tag{17}$$

$$x_{ijk} \in \{0,1\}, \qquad \forall k \in \mathcal{K}, (i,j) \in \mathcal{A}_k \tag{18}$$

$$\delta_{kc} \in \{0,1\}, \qquad \forall k \in \mathcal{K}, c \in \mathcal{C} \tag{19}$$

$$0 \le w_{ijks} \le 1, \qquad \forall k \in \mathcal{K}, (i,j) \in \mathcal{A}_k, s \in \mathcal{S} \tag{20}$$

The objective function (1) is to minimize the total cost; the sum of the sailing costs, the costs associated with visiting ports, and the time charter costs. Constraints (2)-(4) describe the network flow on a route for each vessel k. Constraints (5) describe the relationship between the flow variables and the speed variables, such that the weights of the speed alternatives add up to 1 if the vessel k sails between node i and j, and 0 otherwise. When a vessel k is not used, i.e. $x_{o(k)d(k)k} = 1$, then $\sum_{s\in\mathcal{S}} w_{o(k)d(k)ks} = 1$. To ensure that this sailing does not generate any costs, the corresponding cost parameters are set to $C^S_{o(k)d(k)ks} = 0$ for all s. Constraints (6) ensure that the capacity limit of each product type on vessel k is respected. The load balance constraints (7) ensure that the load on vessel k in node j equals the load in the previous node i adjusted for the quantities loaded and unloaded in node i. Constraints (8) define the initial load on vessel k to be 0. Constraints (9) ensure that the number of pickups of an evenly spread contract c is within the required interval. Constraints (10) and (11) ensure that a vessel k visits both the loading and unloading port of contract c if it serves that contract. Constraints (12) ensure that the quantity picked up from contract c is within the prescribed bounds, while constraints (13) require that the contracted demand over the planning horizon must be serviced completely. Constraints (14) define the time vessel k can start sailing. Constraints (15) ensure that the time of starting service at a node j must be greater than or equal to the start of service at the previous node i, plus the speed-dependent sailing time between the nodes, the piloting time in node i, and the cargo handling time. Constraints (16) ensure that the transit time restrictions are respected, where an upper bound on M^T_c is given by the maximum time a ship may use from $l(c)$ to $u(c)$, i.e. the sum of the maximum sailing, piloting, handling, and waiting times. Constraints (17) require that each vessel that sails must visit the first port within the planning horizon, where an upper bound on M^S_{ik} is given by the maximum time ship k may use from its origin $o(k)$ to port i. Constraints (18) and (19) put binary restrictions on the arc-flow and pickup variables, respectively. Constraints (20) ensure that the speed variable w_{ijks} takes values between 0 and 1.

Evenly spread constraints

The model above does not account for the evenly spread requirements for the partial cargoes under contracts $\mathcal{C}^E$, which represents the planning situation where the service level is completely disregarded. We define the following additional notation to include evenly spread requirements: Let the binary variable z_{kmc} define whether vessel m is the next vessel after vessel k, picking up contract c or not. If $z_{kmc} = 1$, we will refer to the pair of vessels (k, m) as a spread pair. For example, if vessels k, m, and l pick up contract c in the given order, the spread pairs for contract c are given by (k, m) and (m, l). Thus, if a contract is picked up n times, there will exist $(n - 1)$ spread pairs for contract c. Furthermore, let ϕ_{nc} be 1 if contract c is picked up n times during the planning horizon, and 0 otherwise. The variable s_c defines the maximum number of days contract c deviates from the evenly spread requirement, i.e. the maximum deviation from the perfect spread over all spread pairs. If we use the example from Figure 2 where the perfect spread is ten days, we see that the time between service of the first two partial cargoes (first spread pair) is nine days, while it is 11 between the second spread pair. The maximum deviation from the perfect spread is therefore one day, so $s_c = 1$. Finally, let L be the upper limit on the total deviation in days for all contracts from the evenly spread requirement. This can be considered as a measure for the service level regarding the evenly spread requirements.

The evenly spread feature can be added to the STSRSP model (1)-(20) through the following constraints:

$$\sum_{k\in\mathcal{K}}\sum_{m\in\mathcal{K}\setminus\{k\}} z_{kmc} \geq \sum_{k\in\mathcal{K}} \delta_{kc} - 1, \qquad \forall c \in \mathcal{C}^E \qquad (21)$$

$$\sum_{m\in\mathcal{K}\setminus\{k\}} z_{kmc} \leq \delta_{kc}, \qquad \forall k \in \mathcal{K}, c \in \mathcal{C}^E \qquad (22)$$

$$\sum_{m\in\mathcal{K}\setminus\{k\}} z_{mkc} \leq \delta_{kc}, \qquad \forall k \in \mathcal{K}, c \in \mathcal{C}^E \qquad (23)$$

$$\sum_{n=\underline{P}_c}^{\bar{P}_c} n\phi_{nc} = \sum_{k\in\mathcal{K}} \delta_{kc}, \qquad \forall c \in \mathcal{C}^E \qquad (24)$$

$$\sum_{n=\underline{P}_c}^{\bar{P}_c} \phi_{nc} = 1, \qquad \forall c \in \mathcal{C}^E \qquad (25)$$

$$\sum_{n=\underline{P}_c}^{\bar{P}_c} \frac{T^{PH}\phi_{nc}}{n} - s_c - M_c^E(1 - z_{kmc}) \leq t_{l(c)m} - t_{l(c)k}, \quad \forall k \in \mathcal{K}, m \in \mathcal{K}\setminus\{k\}, c \in \mathcal{C}^E \qquad (26)$$

$$\sum_{n=\underline{P}_c}^{\bar{P}_c} \frac{T^{PH}\phi_{nc}}{n} + s_c + M_c^E(1 - z_{kmc}) \geq t_{l(c)m} - t_{l(c)k}, \quad \forall k \in \mathcal{K}, m \in \mathcal{K}\setminus\{k\}, c \in \mathcal{C}^E \qquad (27)$$

$$\sum_{c\in\mathcal{C}^E} s_c \leq L \qquad (28)$$

$$z_{kmc} \in \{0,1\}, \qquad \forall k \in \mathcal{K}, m \in \mathcal{K}\setminus\{k\}, c \in \mathcal{C}^E, \qquad (29)$$

$$\phi_{nc} \in \{0,1\}, \qquad \forall c \in \mathcal{C}^E, n = \underline{P}_c..\bar{P}_c \qquad (30)$$

$$s_c \geq 0, \qquad \forall c \in \mathcal{C}^E \qquad (31)$$

Constraints (21) ensure that the sum of spread pairs for contract c is greater than or equal to the number of pickups minus 1. These could have been modeled as equality constraints, but preliminary testing showed that the model was easier to solve with the constraints as presented here. (22) and (23) require a vessel k to pick up contract c in order be included in a spread pair, while also ensuring that a vessel is present in at most two spread pairs for each contract. Constraints (24) and (25) ensure that ϕ_{nc} is 1 if contract c is picked up n times. Constraints (26) and (27) ensure that vessels k and m have evenly spread arrival times at the loading port of contract c if vessel m is the next vessel after vessel k, picking up contract c. Here, the spread slack variable s_c may take a positive value to correct for the deviation from the desired spread and M_c^E is bounded by the latest time contract c may be serviced plus T^{PH} minus the earliest time contract c may be serviced. The sum of deviations for all contracts are limited by the service level constraint (28). Constraints (29) and (30) put binary restrictions on the spread pair and pickup-counter variables, respectively. Constraints (31) ensure that the spread deviations are non-negative.

3.2 Alternative formulation: Voyage-model

In the Vessel-model formulation, vessels are routed, which is the common way to model similar problems. Here, we present an alternative formulation, where we generate the sailing routes and schedules for all voyages and assign vessels to them. Let $\mathcal{V}$ be the set of possible voyages during the planning horizon, where $|\mathcal{V}| = |\mathcal{K}|$. Let binary variable y_{vk} define whether vessel k sails voyage v or not. The binary variable x_{ijv} defines whether voyage v use the arc between nodes i and j or not, w_{ijvks} represents the weight of speed alternative s for vessel k on the arc (i, j) on voyage v, and l_{ijvp} equals the load of product type p on voyage v on the arc (i, j). δ_{vc} is 1 if voyage v serves contract c, and 0 otherwise, and q_{vcp} represents the quantity of product p in contract c that is picked up on voyage v. The time variables t_{iv} define the start of service at node i on voyage v.

In the Vessel-model, the cost of chartering a vessel was calculated by multiplying a vessel's charter rate by the time spent on its voyage. Here, we split the chartering cost calculations into two terms, where four aspects constitute the total time of a voyage, i.e. sailing time, piloting time, handling time and waiting time. Let C^{SC}_{ijks} be the sailing and chartering cost corresponding to the piloting and sailing time from node i to j with vessel k using speed alternative s. Furthermore, let the variable t^{HW}_k represent the total time used for handling and waiting by vessel k during the voyage. This split has a positive impact on the linear relaxation of the Voyage-model (in contrast to what it had for the Vessel-model).

With this notation, the Voyage-model is given by:

$$\min z = \sum_{k \in \mathcal{K}} \sum_{(i,j) \in \mathcal{A}_k} \sum_{v \in \mathcal{V}} \sum_{s \in \mathcal{S}} C^{SC}_{ijks} w_{ijvks} + \sum_{(i,j) \in \mathcal{A}} \sum_{v \in \mathcal{V}} C^{V}_{i} x_{ijv} + \sum_{k \in \mathcal{K}} C^{C}_{k} t^{HW}_{k} \tag{32}$$

$$\sum_{k \in \mathcal{K}} \sum_{j \in \mathcal{N}^P \cup \{d(k)\}} x_{o(k)jv} = 1, \qquad \forall v \in \mathcal{V} \tag{33}$$

$$\sum_{i \in \mathcal{N}} x_{ijv} - \sum_{i \in \mathcal{N}} x_{jiv} = 0, \qquad \forall v \in \mathcal{V}, j \in \mathcal{N}^P \tag{34}$$

$$\sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{N}^P \cup \{o(k)\}} x_{id(k)v} = 1, \qquad \forall v \in \mathcal{V} \tag{35}$$

$$x_{ijv} = \sum_{k \in \mathcal{K}} \sum_{s \in \mathcal{S}} w_{ijvks}, \qquad \forall (i,j) \in \mathcal{A}, v \in \mathcal{V} \tag{36}$$

$$\sum_{j \in \mathcal{N}} \sum_{v \in \mathcal{V}} x_{o(k)jv} = 1, \qquad \forall k \in \mathcal{K} \tag{37}$$

$$\sum_{i \in \mathcal{N}} \sum_{v \in \mathcal{V}} x_{id(k)v} = 1, \qquad \forall k \in \mathcal{K} \tag{38}$$

$$\sum_{s \in \mathcal{S}} w_{ijvks} \leq y_{vk}, \qquad \forall k \in \mathcal{K}, (i,j) \in \mathcal{A}_k, v \in \mathcal{V} \tag{39}$$

$$\sum_{v \in \mathcal{V}} y_{vk} = 1, \qquad \forall k \in \mathcal{K} \tag{40}$$

$$\sum_{k \in \mathcal{K}} y_{vk} = 1, \qquad \forall v \in \mathcal{V} \tag{41}$$

$$0 \leq l_{ijvp} \leq \sum_{k\in\mathcal{K}} K^V_{kp} y_{vk} - \sum_{p'\in\mathcal{P}^S_p} l_{ijvp'}, \qquad \forall (i,j)\in\mathcal{A}, v\in\mathcal{V}, p\in\mathcal{P} \quad (42)$$

$$l_{ijvp} \leq M^C_p x_{ijv}, \qquad \forall (i,j)\in\mathcal{A}, v\in\mathcal{V}, p\in\mathcal{P} \quad (43)$$

$$\sum_{j\in\mathcal{N}} l_{jivp} + \sum_{c\in\mathcal{C}^L_i} q_{vcp} - \sum_{c\in\mathcal{C}^U_i} q_{vcp} = \sum_{j\in\mathcal{N}} l_{ijvp}, \qquad \forall i\in\mathcal{N}, v\in\mathcal{V}, p\in\mathcal{P} \quad (44)$$

$$\sum_{k\in\mathcal{K}}\sum_{j\in\mathcal{N}} l_{o(k)jvp} = 0, \qquad \forall v\in\mathcal{V}, p\in\mathcal{P} \quad (45)$$

$$\underline{P}_c \leq \sum_{v\in\mathcal{V}} \delta_{vc} \leq \overline{P}_c, \qquad \forall c\in\mathcal{C} \quad (46)$$

$$\delta_{vc} \leq \sum_{i\in\mathcal{N}} x_{il(c)v}, \qquad \forall v\in\mathcal{V}, c\in\mathcal{C} \quad (47)$$

$$\delta_{vc} \leq \sum_{i\in\mathcal{N}} x_{iu(c)v}, \qquad \forall v\in\mathcal{V}, c\in\mathcal{C} \quad (48)$$

$$\underline{Q}_{cp}\delta_{vc} \leq q_{vcp} \leq \overline{Q}_{cp}\delta_{vc}, \qquad \forall v\in\mathcal{V}, c\in\mathcal{C}, p\in\mathcal{P} \quad (49)$$

$$\sum_{v\in\mathcal{V}} q_{vcp} = D_{cp}, \qquad \forall c\in\mathcal{C}, p\in\mathcal{P} \quad (50)$$

$$t_{o(k)v} = T^A_k y_{vk}, \qquad \forall v\in\mathcal{V}, k\in\mathcal{K} \quad (51)$$

$$t_{iv} + T^P_{iv} x_{ijv} + \sum_{c\in\mathcal{C}^L_i\cup\mathcal{C}^U_i}\sum_{p\in\mathcal{P}} T^H_{kp} q_{vcp} + \sum_{k\in\mathcal{K}}\sum_{s\in\mathcal{S}} T^S_{ijks} w_{ijvks} \leq t_{jv} \qquad \forall (i,j)\in\mathcal{A}, v\in\mathcal{V} \quad (52)$$

$$t_{l(c)v} + T^T_c + M^T_c(1-\delta_{vc}) \geq t_{u(c)v}, \qquad \forall v\in\mathcal{V}, c\in\mathcal{C} \quad (53)$$

$$t_{jv} - M^S_{jk}(1 - x_{o(k)jv}) \leq T^{PH}, \qquad \forall j\in\mathcal{N}^P, v\in\mathcal{V}, k\in\mathcal{K} \quad (54)$$

$$t^{HW}_k \geq t_{d(k)v} - t_{o(k)v} - \sum_{(i,j)\in\mathcal{A}_k} T^P_{iv} x_{ijv} - \sum_{(i,j)\in\mathcal{A}_k}\sum_{s\in\mathcal{S}} T^S_{ijks} w_{ijvks} - M^L_k(1-y_{vk}), \quad \forall v\in\mathcal{V}, k\in\mathcal{K} \quad (55)$$

$$\sum_{k\in\mathcal{K}} t^{HW}_k \geq 2\cdot\sum_{c\in\mathcal{C}}\sum_{p\in\mathcal{P}} \min_{\forall k\in\mathcal{K}}(T^H_{kp}) D_{cp}, \quad (56)$$

$$\sum_{i\in\mathcal{N}^P}\sum_{k\in\mathcal{K}} x_{o(k)i(v+1)} \leq \sum_{i\in\mathcal{N}^P}\sum_{k\in\mathcal{K}} x_{o(k)iv}, \qquad \forall v\in\mathcal{V}\setminus\{|\mathcal{V}|\} \quad (57)$$

$$x_{ijv}\in\{0,1\}, \qquad \forall (i,j)\in\mathcal{A}, v\in\mathcal{V} \quad (58)$$

$$\delta_{vc}\in\{0,1\}, \qquad \forall v\in\mathcal{V}, c\in\mathcal{C} \quad (59)$$

$$0 \leq w_{ijvks} \leq 1, \qquad \forall k\in\mathcal{K}, (i,j)\in\mathcal{A}_k, v\in\mathcal{V}, s\in\mathcal{S} \quad (60)$$

The objective function (32) replaces (1) and constraints (33)-(36) correspond to (2)-(5). Constraints (37) and (38) state that each vessel must sail at most one time out from its origin and to its destination. These constraints are redundant with (33) and (35), but included to tighten the linear relaxation. Constraints (39) ensure that a vessel k cannot sail a leg (i,j) on voyage v unless the vessel is assigned to the given voyage. Constraints (40) and (41) ensure that each vessel is assigned to a voyage and each voyage is sailed by one vessel, respectively. Constraints (42) and

(43) are equivalent to (6), while constraints (44)-(54) replace (7)-(17) in the Vessel-model. Constraints (55) set the time each vessel uses on handling and waiting, where an upper bound on M_k^L is given by the latest time ship k may arrive at its artificial destination $d(k)$. Constraint (56) defines a lower bound on the minimum time used to handle the contracts, included to tighten the formulation. Symmetry breaking constraints (57) ensure that voyages that do not visit any ports, i.e. unused voyages, are placed last in the voyage ordering.

Evenly spread constraints

The evenly spread constraints for the Voyage-model are almost identical to the Vessel-model's ones. The main difference is that in the Voyage-model we can utilize the inherent property of the voyage ordering, i.e. if both voyages v_1 and v_2 include a visit to port i, voyage v_1 will visit port i no later than voyage v_2. Let $\mathcal{V}_v^S$ be the set of voyages succeeding voyage v, $\mathcal{V}_v^S \subset \mathcal{V}$. For example, if $\mathcal{V} = \{1,2,3,4\}$, then $\mathcal{V}_2^S = \{3,4\}$. Let the binary variable z_{vwc} define whether voyage w is the next voyage after voyage v, picking up contract c or not, defined for all $v \in \mathcal{V}, w \in \mathcal{V}_v^S$. The evenly spread constraints for the voyage-model can be formulated as follows:

$$\sum_{v\in\mathcal{V}}\sum_{w\in\mathcal{V}_v^S} z_{vwc} \geq \sum_{v\in\mathcal{V}} \delta_{vc} - 1, \qquad \forall c \in \mathcal{C}^E \tag{61}$$

$$\sum_{w\in\mathcal{V}_v^S} z_{vwc} \leq \delta_{vc}, \qquad \forall v \in \mathcal{V}, c \in \mathcal{C}^E \tag{62}$$

$$\sum_{w\in\mathcal{V}\setminus(\mathcal{V}_v^S\cup\{v\})} z_{wvc} \leq \delta_{vc}, \qquad \forall v \in \mathcal{V}, c \in \mathcal{C}^E \tag{63}$$

$$\sum_{n=\underline{P}_c}^{\bar{P}_c} n\phi_{nc} = \sum_{v\in\mathcal{V}} \delta_{vc}, \qquad \forall c \in \mathcal{C}^E \tag{64}$$

$$\sum_{n=\underline{P}_c}^{\bar{P}_c} \phi_{nc} = 1, \qquad \forall c \in \mathcal{C}^E \tag{65}$$

$$\sum_{n=\underline{P}_c}^{\bar{P}_c} \frac{T^{PH}\phi_{nc}}{n} - s_c - M_c^E(1 - z_{vwc}) \leq t_{l(c)w} - t_{l(c)v}, \quad \forall v \in \mathcal{V}, w \in \mathcal{V}_v^S, c \in \mathcal{C}^E \tag{66}$$

$$\sum_{n=\underline{P}_c}^{\bar{P}_c} \frac{T^{PH}\phi_{nc}}{n} + s_c + M_c^E(1 - z_{vwc}) \geq t_{l(c)w} - t_{l(c)v}, \quad \forall v \in \mathcal{V}, w \in \mathcal{V}_v^S, c \in \mathcal{C}^E \tag{67}$$

$$\sum_{c\in\mathcal{C}^E} s_c \leq L \tag{68}$$

$$z_{vwc} \in \{0,1\}, \qquad \forall v \in \mathcal{V}, w \in \mathcal{V}_v^S, c \in \mathcal{C}^E, \tag{69}$$

$$\phi_{nc} \in \{0,1\}, \qquad \forall c \in \mathcal{C}^E, n = \underline{P}_c..\bar{P}_c \tag{70}$$

$$s_c \geq 0, \qquad \forall c \in \mathcal{C}^E \tag{71}$$

Constraints (61)-(71) replace constraints (21)-(31) for the Vessel-model.

4 Computational study

We have performed a computational study using 90 test instances, generated mainly based on data provided by the case company. We begin in Section 4.1 by introducing the test instances, before testing the performance of the Vessel and Voyage models in Section 4.2. Finally, the potential benefits of flexible planning and managerial insights are explored in Section 4.3.

The mathematical models are implemented in Mosel and solved using Xpress-IVE 1.28.12. All computational experiments have been run on a PC with Intel Core i7-6500U processor and 16 GB of RAM running Windows 10.

4.1 Test instances

To evaluate the performance and capabilities of the two proposed models, we have generated a large number of instances, representing realistic planning instances for the case company. Three different trade routes are used in this computational study: US - Japan (5 ports in total), Asia - Europe (10 ports in total), and Europe - US (15 ports in total). The trade routes are referred to based on their number of ports, i.e. small (S), medium (M), and large (L), respectively. The port visit costs are drawn from a uniform distribution in the interval [25 000 USD, 40 000 USD]. The total volume to be transported is based on data from the case company. For each trade route, two sets of instances with 50 and 100 contracts are created, which results in a total of six cases with 15 instances in each, i.e. S-50, S-100, M-50, M-100, L-50, and L-100.

Each case is further divided into three sets of instances used for studying the effect of the service level requirement, i.e. N, M, and H, representing sets of five instances with no (N), medium (M) or high (H) service level requirements, respectively. This results in five triples of instances where only the service level requirements differ among the instances within a triple. The instance specific service level requirement is defined by the service level threshold L. The threshold value for the high service level, i.e. L^H, is set by solving the STSRSP where the objective function is replaced by min $L^H = \sum_{c \in \mathcal{C}} s_c$, for each instance. This means that the service level threshold L is an instance specific parameter, not a fixed (e.g. average) value used for all instances. We do this to ensure that all instances can be solved feasibly. The medium service level threshold is set as $L^M = L^H + (L^N - L^H)/3$, where L^N is the post calculated threshold with no service level requirement.

Finally, for all instances the following parameters are drawn from a uniform distribution with the corresponding intervals given in brackets: First, the size of each contract is set, measured in square meters of goods. We associate a size parameter with each contract c, which is set in the interval [0.05, 1] for ordinary contracts and [0.2, 1.5] for the evenly spread contracts. Then, the size parameters are normalized. The total volume to be transported is then distributed among the contracts based on the normalized size parameters, which gives the demand for each contract. The service frequency requirement for each contract is chosen in the interval [1, minimum required voyages], where the minimum required voyages is calculated as the total demand among all contracts over the planning horizon divided by the capacity of the largest vessel (rounded up to the nearest integer).

The transit time limits for the contracts with such requirements are chosen in the interval [minimum sailing time · 1.2, maximum sailing time without waiting], where the minimum sailing time is the time for the direct sailing between the two ports with the highest possible speed, while the maximum sailing time without waiting corresponds to the sailing time in the case with the lowest possible speed and where all ports between the contract's loading and unloading port along the trade are visited. Loading and unloading ports are set based on the typical demand in the ports. For the evenly spread contracts, the minimum and maximum quantities picked up are given by $\underline{Q}_{cp} = 0.8D_{cp}/\bar{P}_c$ and $\bar{Q}_{cp} = 1.2D_{cp}/\underline{P}_c$, respectively. For all other contracts, $\underline{Q}_{cp} = D_{cp}/|\mathcal{K}|$ and $\bar{Q}_{cp} = D_{cp}$. The vessels' characteristics, such as capacities, speed-dependent fuel consumption functions, and estimations of time charter rates are provided by the case company. The bunker price is set to 350 USD/mt. Daily charter rates are in the interval [15 000 USD, 40 000 USD], ship capacities are in the interval [35 000 m^2, 75 000 m^2], and fuel consumptions are in the interval [25 ton/day, 90 ton/day]. A randomly selected subset of the vessels are available in each instance.

We identify each set of five instances by its name Size-Contracts-Service level characteristics, so that M-50-H means a set of five instances on the medium trade route, with 50 contracts, with a high service level requirement. All instances consist of 40% evenly spread contracts, 20% transit time contracts, and 40% of contracts without any service requirements. 90% of the contracts are inter-regional and 10% intra-regional. The planning horizon is set to 30 days, meaning that all voyages must begin within this time limit. We have set a maximum running time of three hours (10,800 seconds) for solving each of the instances.

4.2 Comparison of models

Table 1 shows the average sizes of the MIPs for the Vessel and the Voyage models, grouped by trade size. For the small instances, we see that the number of constraints and variables in the models are moderately higher for the Voyage-model, while having almost equally many binary variables. However, for the medium and large instances, the differences are notable. The continuous speed variables make up the largest share of the variables, given by w_{ijks} and w_{ijvks} for the Vessel and the Voyage models, respectively. Thus, the Voyage-model has approximately the cardinality of voyages ($|\mathcal{V}|$) as many speed variables than the Vessel-model. The Voyage-model has more constraints and binary variables on average. Note that almost half of the binary variables in the Vessel-model are present due to the evenly spread constraints, while this share is only around 20% for the Voyage-model.

To compare the performance of the Vessel and Voyage models, we have tested both formulations on all 90 test instances. The test results are summarized in Table 2. For each of the six cases, shown as separate rows in the table and containing 15 instances each, average solution times for both models are presented, as well as the number of instances solved to optimality (# Optimal) and for which we have obtained feasible integer solutions (# Feasible) within the maximum running time, which we have set to three hours (10,800 seconds).

For the smallest instances, case S-50-*, we see that the difference in solution times are small, and both the Voyage and the Vessel-models solve all instances to optimality in reasonable time. For the instances in case S-100-*, we see signifi-

Table 1 Average number of constraints, variables (both continuous and binary), and binary variables for all instances per case after presolve, reported as: (result without evenly spread) / (result with evenly spread).

Set of instances	Vessel-model			Voyage-model		
	# Constraints	# Variables	# Binary variables	# Constraints	# Variables	# Binary variables
S-50-*	1428/2116	936/1028	269/457	2103/2530	1791/1923	421/535
S-100-*	2329/3487	1200/1630	440/836	3203/3863	2183/2434	621/833
M-50-*	2685/3829	2392/2819	506/919	7244/7698	8298/8550	1040/1274
M100-*	4212/6228	2945/3731	763/1554	8081/9145	8857/9415	1303/1753
L-50-*	5085/6508	5394/5884	928/1428	14473/14908	21359/21508	1922/2203
L-100-*	6419/7974	5918/6925	1183/2170	14972/16023	20691/21122	2098/2628

The asterisks (*) denote all instances within a certain set.

Table 2 Average solution time (in seconds) and number of instances solved per case

Set of instances	Vessel-model			Voyage-model		
	Sol time	# Optimal	# Feasible	Sol time	# Optimal	# Feasible
S-50-*	13	15	15	12	15	15
S-100-*	188	15	15	27	15	15
M-50-*	4352	11	15	2407	13	15
M-100-*	7469	5	5	5424	9	15
L-50-*	10800	0	7	9637	3	12
L-100-*	10373	1	5	9505	3	9
Average	5533	7.8	10.3	4502	9.7	13.5

The asterisks (*) denote all instances within a certain set.

cant differences in the average solution times between the two models, which may indicate that the Vessel-model is more sensitive to an increase in the number of contracts than the Voyage-model. For the medium-sized instances, i.e. cases M-50-* and M-100-*, the Voyage-model clearly outperforms the Vessel-model. The Voyage-model solves 22 of the 30 instances to optimality and finds feasible solutions to all 30, while the Vessel-model finds the optimal solutions only for 16 instances and cannot find feasible solutions for as many as 10 instances. Finally, we see that both models struggle with solving the large instances belonging to cases L-50-* and L-100-*, but also here it can be noticed that the Voyage-model performs significantly better. The Vessel-model can only find one optimal and 12 feasible solutions for the 30 instances, while the Vessel-model finds six optimal and 21 feasible solutions within the maximum running time.

These results contradict the expected outcomes based on the average number of constraints and variables in the models, see Table 1. However, despite the Voyage-model having more constraints and variables, the computational tests show that the average LP-relaxation over all 90 instances is 8.82% better when using the Voyage-model, compared to the Vessel-model. For all instances solved to optimality, the average relative MIP gap in the root node is 18.37% for the Vessel-model and 11.52% for the Voyage-model, which may explain the better performance of the Voyage-model.

Table 3 shows the average results over all instances, sorted by the service level. We see that for the 30 instances solved without the evenly spread requirement (marked with *-*-N), the Voyage-model performs slightly better than the Vessel-model. When the evenly spread constraints are active, i.e. service level medium and high (marked with *-*-M and *-*-H, respectively), the Voyage-model clearly outperforms the Vessel-model, both with regards to solution times and the number of instances where both optimal and feasible solutions are found within the

Table 3 Average results, grouped by service level.

Set of instances	Vessel-model			Voyage-model		
	Sol time	# Optimal	# Feasible	Sol time	# Optimal	# Feasible
--N	3546	21	30	2701	25	30
--M	5864	15	17	4369	20	29
--H	7188	11	15	6437	13	22

The asterisks (*) denote all instances within a certain set.

maximum running time. It is also clear that the instances where the evenly spread requirements are disregarded are significantly easier to solve. This is most likely due to the reduced number of constraints and binary variables, as shown in Table 1.

The results also show that a higher service level requirement increases the difficulty of solving the instances, which could be explained by the following two aspects. Firstly, as a high service level implies a low service level threshold L, these instances are more constrained and feasibility becomes a challenge. Secondly, when lowering the service level threshold, the objective function value in the optimal solution increases, but the computational tests show that the LP-relaxations are unaffected by the threshold L for all test instances. This means that the solver needs to close a larger MIP gap when lowering the service level threshold, which most often implies increased solution times.

Overall the Voyage-model performs better than Vessel-model. It seems like the evenly spread requirements are handled in a better way by the Voyage-model, based on the results in Table 3. Surprisingly, the Voyage-model also performs better than the Vessel-model when solved without the evenly spread constraints. This is an interesting finding. Without the evenly spread restrictions, the STSRSP has similarities with both the industrial and tramp ship routing and scheduling problem (e.g. Christiansen et al (2013)) and the vehicle routing problem with a heterogeneous fleet and time windows (e.g. Jiang et al (2014) and Koç et al (2015)). It might be that the arc-flow formulations also for these problems could benefit from routing voyages/trips instead of vessels/vehicles.

4.3 Managerial insights

The problem discussed in this paper is motivated by the fact that RoRo-shipping companies are neither limited to weekly frequencies on the trades, nor to visiting every port on each voyage. One could say that RoRo-shipping is more demand-driven, in contrast to container shipping which is more frequency-driven. In this section, we study the potential gains that can be achieved by utilizing this inherent planning flexibility. Furthermore, we discuss the trade-offs between service level with regards to the evenly spread requirements and the costs of the solutions.

The test instances presented in Section 4.1 are again used for studying the potential gains of utilizing the planning flexibility. For each instance, we also solve a version of the problem where all ports are visited along each voyage, which we denote as the All Ports Regularly (APR) approach. This closely resembles current practice where, to make the planning tractable, one lets all ports along the trade be visited on each voyage (similar to container shipping). As discussed in the introduction, this also makes it much easier to obtain solutions that respect all contract requirements regarding evenly spread in a reasonable way. Then, for each instance, we compare the cost of sailing with regular intervals (APR) with the cost of the solutions with no, medium or high service level requirements. The Voyage-model is used for these calculations. The costs of APR are found using a slightly altered version of the Voyage-model. Firstly, we assume that regular intervals are sufficient to accommodate the evenly spread restrictions. Thus, we remove the service level threshold constraints (68) from the problem, but the remaining evenly spread constraints are kept to calculate the provided service level. Secondly, we ensure that each port is visited at regular intervals, e.g. every 10th day. We set the intervals based on the length of the planning horizon divided by the number of vessels used in the best solution from the commercial solver. Thus, if three vessels serve the trade, each port is visited every tenth day. The average results for the small and medium trade instances are presented in Table 4. All costs are shown in percentage of the APR solutions to easily see the cost reduction potentials from utilizing this planning flexibility. As the solver is unable to provide reasonable optimality gaps for the large instances, these instances are omitted in this analysis.

Table 4 Planned costs and service level comparison for the small and medium trade sizes. All instances are solved to optimality for the different service level requirements and the all ports regularly approach (APR)

			Service level requirement					
Set of instances	APR		No requirement		Medium		High	
	Cost	ST/c	Cost	ST/c	Cost	ST/c	Cost	ST/c
S-50-*	100 %	0.90	92.9 %	1.97	94.3 %	0.75	101.5 %	0.28
S-100-*[1]	100 %	1.00	92.3 %	1.19	96.5 %	0.48	102.3 %	0.24
M-50-*[1]	100 %	0.63	87.5 %	2.93	89.1 %	1.00	96.7 %	0.08
M-100-*[1]	100 %	0.59	90.9 %	2.94	92.6 %	1.03	100.5 %	0.14
Average	100 %	0.81	90.9 %	2.18	92.6 %	0.80	100.5 %	0.21

Cost - objective value as a percentage of the APR objective value. ST/c - Average deviation in days from the evenly spread requirement, per evenly spread contract. 1. Three out of the five instances in each set were infeasible due to transit time restrictions for the APR computations and are therefore not included in the table.

The results show that visiting all ports on a regular basis gives an average deviation from the desired spread of 0.81 days per evenly spread contract. If the evenly spread requirements are disregarded (No requirement), the average deviation increases to 2.18 day per evenly spread contract. While this increase is substantial, the average savings on planned costs are 9.1%, which highlights the important trade-off between service level and planned costs. For the medium service level, the results show that an equally good service level as for the APR can be provided at 7.4% lower costs. From a managerial point of view, this result shows that relaxing a strict, regular service requirement for each port (close to current practice) can greatly reduce the overall costs. For the high service level, we observe lower deviations for the evenly spread, but the costs are slightly higher than for the APR. Possible explanations for this finding are increased waiting times and higher sailing speeds.

In Table 4, we have chosen to report the average deviation from the desired spread, as the service threshold is set on a company level. However, it could be argued that the service level constraints also should be set on contract level, to ensure a fair spread balance between the contracts and prevent extreme deviations

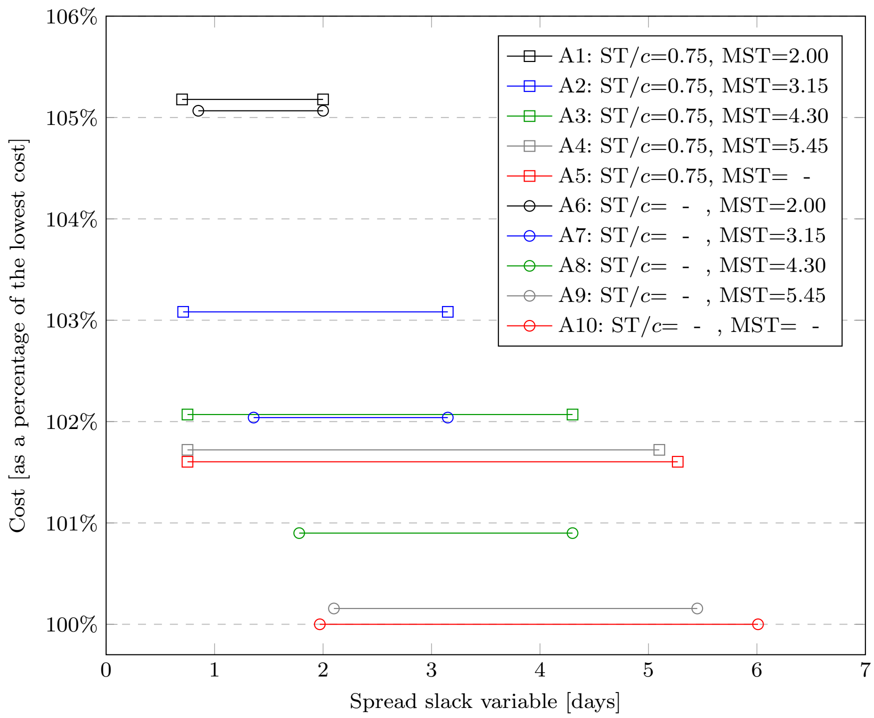

Fig. 4 Cost comparison of various threshold values for both the average slack time per contract (ST/c) and the maximum slack time (MST) for case S-50-*. The thresholds for each entry are given in the top right corner. Each entry is presented in the graph as a line, where the leftmost marker shows the resulting ST/c and the rightmost marker shows the MST.

from the desired spread. In Figure 4, we study the effects of constraining the maximum deviation with regards to costs, for case S-50-*. This is enforced by adding a new set of constraints, similar to constraints (28) and (68): $s_c \leq \text{MST}, \quad \forall c \in \mathcal{C}^E$, where the parameter MST is the maximum slack time for each contract. Each of the instances within case S-50-* is solved for five MST values, ranging from the lowest feasible integer slack time (2 days) to the lowest non-binding MST value (6.6 days), with fixed step length (1.15 days). All instances are solved twice, both with medium and no service level requirement, i.e. with and without a service level threshold L. In Figure 4, entries A1-A5 show the results when enforcing medium service level, as well as constraining the maximum deviation. For entries A6-A10, only the maximum service level is constrained. If both sets of service level constraints are disregarded, i.e. entry A10, the average slack time per contract (ST/c) and the maximum slack time are approximately two and six days, respectively. We observe a minor cost increase from reducing the MST to 5.45 days (A9), but as the MST is further lowered, the cost rapidly increases. The results for entries A6-A10 also show that reducing the MST most often has a positive effect on the ST/c. Entry A5 shows that the medium service level results both have a lower ST/c and MST, at a cost increase of 1.5 %, which may be preferable from a managerial point of view. Entries A3 and A7 highlight the intricacy of these trade-offs.

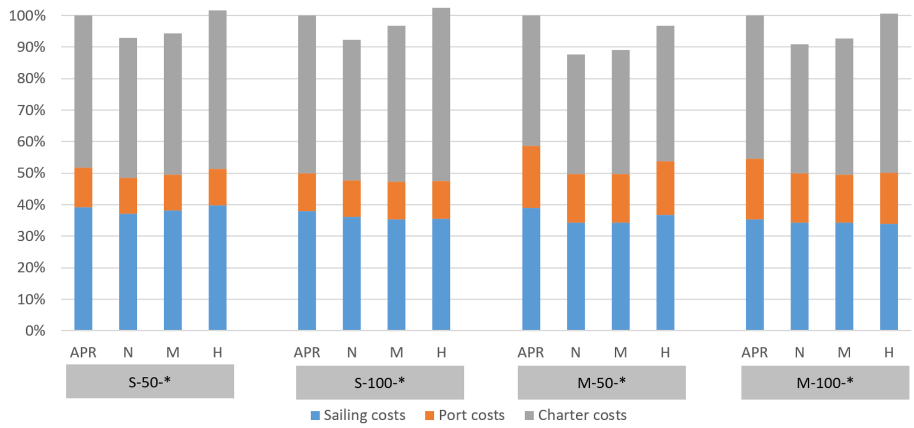

Fig. 5 Cost distributions of the solutions to the instances from Table 4

At approximately the same cost, one could either provide a low average deviation and a high maximum deviation, or vice versa.

The results summarized in Table 4 show the trade-offs between service level and the planning costs. We see that a higher service level implies higher costs. The main reason for this is that a higher service level may require more port visits to comply with the evenly spread requirements, which increases the port costs and the charter times. Additionally, we see that some voyages have increased waiting time at ports or lower/higher sailing speed to spread the contracts' pickups evenly. While this result is as expected, the trade-off costs were more significant than expected. We see that reducing the service level from high to medium reduces the average costs from 100.5% to 92.6%, while the difference in the average slack time per contract (ST/c) is close to half a day. The results highlight the importance, from a managerial point of view, of determining an appropriate service level. The results can be used to identify unfavorable/costly customers/contracts and determining the marginal cost of increasing the service level, which may be useful in future contract negotiations.

Figure 5 shows how the total costs are distributed between sailing costs, port call costs, and charter costs. The height of each column corresponds to the average total costs for the different service levels from Table 4. On average, the sailing costs constitute close to 40 % of the total costs, while the port call costs make up 15 % of the total costs. For the small instances, we see that the port costs are close to equal between APR and the three service levels. However, for the medium sized instances, it is clear that visiting all ports regularly has a considerable impact on the total costs. The figure also shows that the charter costs are the main cost component. We see that when increasing the service level, the charter costs increase more than the other cost components. This is most likely due to increased waiting times at ports to comply with the evenly spread restriction.

5 Conclusions

We have considered a new problem to the literature, which we have called the single trade ship routing and scheduling problem (STSRSP). The STSRSP deals with the planning of interrelated voyages with separation requirements along a single trade with the aim of utilizing the inherent planning flexibility in roll-on roll-off (RoRo) shipping. A given trade in RoRo-shipping is usually designed based on a large number of cargo contracts for transportation of cargoes between the different port pairs along a trade, where the contracts state that they should be serviced with a given frequency, but also that these services should be fairly evenly spread in time. Depending on the frequency and spread requirements of each cargo contract, one might not need to visit each port every time a vessel performs a voyage on the trade. Therefore, instead of aggregating requirements from cargo contracts to frequency requirements for the whole trade, which has been common to simplify planning in previous studies, we look at the frequency requirements for each contract and port along the trade. This is further utilized in the STSRSP to determine the sailing routes of each voyage along the trade, i.e. which ports to be visited and not, which ships to use, what contracts to serve, as well as the sailing speeds along the voyages, so that all contract requirements regarding frequency and spread are satisfied at minimum cost.

We have proposed two novel mixed integer programming models for the STSRSP. A central part of the models includes a new way of modeling the separation requirements for the contracts. Firstly, we showed that the less intuitive model (i.e. the Voyage-model) performs significantly better than the other one (i.e. the Vessel-model). Secondly, we demonstrated through a computational study on a large number of realistic test instances that there are significant gains that potentially can be obtained from utilizing the inherent planning flexibility. We also discussed the trade-offs between service level with regards to the evenly spread requirements and the costs of a solution.

References

Andersson H, Fagerholt K, Hobbesland K (2015) Integrated maritime fleet deployment and speed optimization: Case study from RoRo shipping. Computers & Operations Research 55:233–240

Bakkehaug R, Rakke JG, Fagerholt K, Laporte G (2016) An adaptive large neighborhood search heuristic for fleet deployment problems with voyage separation requirements. Transportation Research Part C 70:129–141

Borthen T, Loenechen H, Wang X, Fagerholt K, Vidal T (2017) A genetic search-based heuristic for a fleet size and periodic routing problem with application to offshore supply planning. EURO Journal on Transportation and Logistics https://doi.org/10.1007/s13676-017-0111-x

Brouer BD, Alvarez JF, Plum CE, Pisinger D, Sigurd MM (2013) A base integer programming model and benchmark suite for liner-shipping network design. Transportation Science 48(2):281–312

Campbell AM, Wilson WH (2014) Forty years of periodic vehicle routing. Networks 63(1):2–15

Christiansen M, Fagerholt K, Nygreen B, Ronen D (2007) Maritime transportation. Handbooks in Operations Research and Management Science, (Eds C Barnhart and G Laporte) 14:189–284

Christiansen M, Fagerholt K, Nygreen B, Ronen D (2013) Ship routing and scheduling in the new millenium. European Journal of Operational Research 228(3):467–483

Fagerholt K, Johnsen TAV, Lindstad H (2009) Fleet deployment in liner shipping: a case study. Maritime Policy & Management 36(5):397–409
Fischer A, Nokhart H, Olsen H, Fagerholt K, Rakke JG, Stålhane M (2016) Robust planning and disruption management in roll-on roll-off liner shipping. Transportation Research Part E 91:51–67
ISL (2017) Shipping statistics and market review 2017 61(9/10):39–66
Jiang J, Ng KM, Poh KL, Teo KM (2014) Vehicle routing problem with a heterogeneous fleet and time windows. Expert Systems with Applications 41(8):3748–3760
Karsten CV, Pisinger D, Ropke S, Brouer BD (2015) The time constrained multi-commodity network flow problem and its application to liner shipping network design. Transportation Research Part E 76:122–138
Kisialiou Y, Gribkovskaia I, Laporte G (2018) The periodic supply vessel planning problem with flexible departure times and coupled vessels. Computers & Operations Research 94:52–64
Koç C, Bektaş T, Jabali O, Laporte G (2015) A hybrid evolutionary algorithm for heterogeneous fleet vehicle routing problems with time windows. Computers & Operations Research 64:11–27
Meng Q, Wang S, Andersson H, Thun K (2014) Containership routing and scheduling in liner shipping: overview and future research directions. Transportation Science 48(2):265–380
Ng M (2015) Container vessel fleet deployment for liner shipping with stochastic dependencies in shipping demand. Transportation Research Part B 74:79–87
Norstad I, Fagerholt K, Hvattum LM, Arnulf HS, Bjørkli A (2015) Maritime fleet deployment with voyage separation requirements. Flexible Services and Manufacturing Journal 27(2-3):180–199
Pantuso G, Fagerholt K, Wallace SW (2016) Uncertainty in fleet renewal: a case from maritime transportation. Transportation Science 50(2):390–407
Reinhardt LB, Plum CEM, Pisinger D, Sigurd MM, Vial GTP (2016) The liner shipping berth scheduling problem with transit times. Transportation Research Part E 86:116–128
Sigurd MM, Ulstein NL, Nygreen B, Ryan DM (2005) Ship scheduling with recurring visits and visit separation requirements. Column Generation, (Eds G Desaulniers and J Desrosiers and MM Solomon) pp 225–245
Vilhelmsen C, Lusby RM, Larsen J (2017) Tramp ship routing and scheduling with voyage separation requirements. OR Spectrum 39(4):913–943
Wang S, Meng Q (2017) Container liner fleet deployment: A systematic overview. Transportation Research Part C 77:389–404
Wang S, Meng Q, Zhuo S (2013) Container routing in liner shipping. Transportation Research Part E 49:1–7

Appendix

List of notations used in the models:

Sets

$\mathcal{K}$	Set of available vessels.
$\mathcal{N}$	Set of nodes.
$\mathcal{N}_k$	Set of nodes that can be visited by vessel k.
$\mathcal{N}^P$	Set of ports along the trade route.
$\mathcal{N}_k^P$	Set of ports that can be visited by vessel k.
$\mathcal{A}$	Set of feasible arcs.
$\mathcal{A}_k$	Set of feasible arcs for vessel k.
$\mathcal{C}$	Set of given contracts for cargoes to be transported along the trade during the planning horizon.
$\mathcal{C}^T$	Set of contracts with transit time restrictions.
$\mathcal{C}^E$	Set of contracts with evenly spread requirements.
$\mathcal{C}_i^L$	Set of cargoes that are to be loaded at port i.

$\mathcal{C}_i^U$	Set of cargoes that are to be unloaded at port i.
$\mathcal{P}$	Set of product types.
$\mathcal{P}_p^S$	Set of product types that can be stored in the same space as product type p.
$\mathcal{S}$	Set of discrete speed alternatives.
$\mathcal{V}$	Set of voyages (Voyage-model only).
$\mathcal{V}_v^S$	Set of voyages succeeding voyage v (Voyage-model only).

Parameters

K_{kp}^V	Capacity for product p on ship k.
D_{cp}	Demand for the whole planning period for product type p for contract c.
$\underline{Q}_{cp}$	Minimum pickup quantity for product type p for contract c.
$\bar{Q}_{cp}$	Maximum pickup quantity for product type p for contract c.
$\underline{P}_c$	Minimum number of pickups of contract c.
$\bar{P}_c$	Maximum number of pickups of contract c.
T_{ijks}^S	Sailing time from a node (port) i to node (port) j for vessel k using speed alternative s.
T_k^A	The time vessel k becomes available at its origin.
T_{ik}^P	Piloting time at port i by ship k.
T_{kp}^H	The time used to handle, i.e. load or unload, one unit of product type p on vessel k.
T_c^T	Maximum transit time for contract c.
T^{PH}	Length of the planning horizon.
L	Evenly spread service level threshold.
$o(k)$	Initial position of vessel k.
$d(k)$	Artificial ending position of vessel k.
$l(c)$	Loading port of contract c.
$u(c)$	Unloading port of contract c.
C_k^C	Daily charter rate for vessel k.
C_i^V	Cost of calling port i.
C_{ijks}^S	Cost of sailing from a node (port) i to node (port) j for vessel k using speed alternative s. (Vessel-model only).
C_{ijks}^{SC}	Sailing and chartering cost corresponding to the piloting and sailing time from node i to j with vessel k using speed alternative s (Voyage-model only).

Variables, Vessel-model

x_{ijk}	1 if vessel k sails from node i to node j, 0 otherwise.
δ_{kc}	1 if vessel k serves contract c, 0 otherwise.
w_{ijks}	Weight of speed alternative s for vessel k on the arc (i, j).
l_{ijkp}	Load of product type p on vessel k on the arc (i, j).
q_{kcp}	Quantity of product p in contract c that is picked up by vessel k.
t_{ik}	Start of service at node i for vessel k.
z_{kmc}	1 if vessel m is the next vessel after vessel k, picking up contract c, 0 otherwise.

ϕ_{nc}	1 if contract c is picked up n times during the planning horizon, 0 otherwise.
s_c	Maximum number of days contract c deviates from the evenly spread requirement.

Variables, Voyage-model

x_{ijv}	1 if voyage v use the arc between nodes i and j, 0 otherwise.
y_{vk}	1 if vessel k sails voyage v, 0 otherwise.
δ_{vc}	1 if voyage v serves contract c, 0 otherwise.
w_{ijvks}	Weight of speed alternative s for vessel k on the arc (i, j) on voyage v.
l_{ijvp}	Load of product type p on voyage v on the arc (i, j).
q_{vcp}	Quantity of product p in contract c that is picked up on voyage v.
t_{iv}	Start of service at node i on voyage v.
t_k^{HW}	Total time used on handling and waiting by vessel k.
z_{vwc}	1 if voyage w is the next voyage after voyage v, picking up contract c, 0 otherwise.
ϕ_{nc}	1 if contract c is picked up n times during the planning horizon, 0 otherwise.
s_c	Maximum number of days contract c deviates from the evenly spread requirement.