

Robust planning and disruption management in roll-on roll-off liner shipping

Andreas Fischer^a, Håkon Nokhart^a, Henrik Olsen^a, Kjetil Fagerholt^{a,b}, Jørgen Glomvik Rakke^c, Magnus Stålhane^a

^a*Norwegian University of Science and Technology,
Department of Industrial Economics and Technology Management
Alfred Getz veg 3, NO-7491,
Trondheim, Norway*

^b*Norwegian Marine Technology Research Institute (MARINTEK),
Otto Nielsens veg 10, NO-7052
Trondheim, Norway*

^c*Wallenius Wilhelmsen Logistics,
P.O Box 33, NO-1324
Lysaker, Norway*

Abstract

This paper presents different strategies for handling disruptions in fleet deployment in roll-on roll-off liner shipping, which basically consists of assigning a fleet of vessels to predefined voyages at minimum cost. A new mathematical model of the problem is presented, including a set of robust planning strategies, such as adding slack and rewarding early arrivals. To solve real-life instances a rolling horizon heuristic is proposed. A computational study, where we also propose some recovery planning strategies, is conducted, and simulation results show that adding robustness significantly reduces the actual cost of the plan and the total delays of the voyages.

Keywords: Maritime transportation; fleet deployment; robust planning; disruption management;

1. Introduction

This paper studies different strategies for handling disruptions in roll-on roll-off (RoRo) liner shipping. The RoRo segment of liner shipping is characterized by transporting cargo such as cars, trucks, farming equipment, and other types of cargo that contain wheels or can be placed on trailers, and thus can be rolled onto and off of the vessels. The vessels themselves are special purpose vessels, similar to gigantic ferries, where the rolling cargo is stacked on multiple decks onboard the vessels, with each deck having different height and weight restrictions.

This study is done in cooperation with Wallenius Wilhelmsen Logistics (WWL), which is one of the world's largest operators in the RoRo segment. The planning process undertaken by WWL can be divided into strategic, tactical, and operational levels of planning. At the strategic level decisions regarding which trade routes they should service, the frequency of service on these trade routes, and the composition of the fleet must be taken. Once this is decided, the company faces a fleet deployment problem (FDP) at the tactical level, where each vessel is assigned to service

a sequence of trade routes, while taking the frequency requirements into account. Finally, at the operational level, decisions regarding stowage onboard the vessels, weather routing, and disruption management, is handled.

In this paper we focus on creating robust solutions to the FDP at the tactical level, in order to minimize the effect of disruptions at the operational level. In addition, we study recovery strategies that may be employed once a disruption has occurred that makes the existing plan undesirable. The FDP faced by WWL may be described as optimally assigning individual vessels to trade route voyages during a planning horizon stretching from a few months and up to a year. A trade route is a long distance maritime shipping route, and consists of a pre-determined sequence of port calls with sailing legs in between. E.g. a ship performing/sailing/servicing the trade route *Europe - North America*, visits a set of ports in Europe, where cargo is loaded, before sailing across the Atlantic to North-America where a set of unloading ports are visited, and the ship is emptied before starting on a new trade route voyage, possibly with some ballast (empty) sailing in between. According to monthly demand and contractual agreements, each trade route has to be serviced with a certain frequency, and a vessel completing one instance of a trade route once is denoted as a voyage. For each vessel, we thus get a schedule which consists of a sequence of voyages, combined with a start time for each voyage.

The term *fleet deployment* is defined by [16] as the “...allocation of ships to routes, their general scheduling, and the chartering of vessels, if any, to complement the owned fleet in the fulfillment of the transportation missions”. Fagerholt et al. [5] describe the FDP as determining an optimal way of servicing voyages defined for the planning horizon with the shipping company’s fleet of vessels. Here, an optimal deployment is seen with respect to minimizing the costs. Most literature on the FDP studies the problem faced by container liner shipping companies. Gelareh and Meng [6] present a nonlinear mathematical model for the FDP with speed optimization in container liner shipping. The model is linearized as a mixed integer linear programming model and solved for randomly generated numerical instances. In Meng and Wang [13], the authors present a chance constrained programming model to solve the short-term liner ship fleet planning problem with cargo demand uncertainty for a liner container company. The demand uncertainty is modeled with a normal distribution, and for each liner route operated by the company, the chance constraints ensure that a minimum level of service is associated with each route. Wang et al. [22] revisit the problem discussed in [13], but propose a joint chance constrained programming model, and show that the service level provided has significant effect on the total cost. Liu et al. [12] formulate a joint optimization model for the container flow and fleet deployment problems. Finally, Wang and Meng [21] present a cost minimization model for the FDP with container transshipment operations. The model allows containers to be delivered from its origin port to its destination port by the use of more than one vessel.

The FDP in container shipping separates itself from RoRo liner shipping in two respects: (1) The container shipping lines are cyclic with each ship assigned to a single service/line/port rotation for the entire planning horizon, while the trade routes in RoRo shipping usually go from one part of the world to another, and each ship is assigned to a sequence of voyages on different trade routes. (2) In container shipping they usually consider deployment of ship types, and not the specific deployment of individual ships. The latter also means that they are not taking into account the different initial positions the ships may have at the start of the planning horizon. This may be a reasonable assumption in container liner shipping due to (1), but is a too crude approximation in RoRo liner shipping.

To the authors' knowledge, Fagerholt et al. [5] and Andersson et al. [2] are the only studies that deal with the FDP in RoRo liner shipping. Fagerholt et al. [5] present a very simplified model that is solved by a multi-start local search heuristic. Later, Andersson et al. [2] extend their model, and propose a rolling horizon heuristic to solve an integrated fleet deployment and speed optimization problem. While studying the same problem as in this paper, their model differs from ours by simplifying the modeling of the cargo onboard the vessel. Their model has one cargo capacity per vessel, considers only one type of cargo, and has only one capacity restriction per voyage, thus implicitly assuming that all cargo transported on a voyage is onboard the vessel concurrently at some point during the voyage. The model presented in this paper divides the capacity of each vessel into groups dependent on the the height and weight limits of each cargo deck, considers more than one type of cargo, and takes into account the fact that several loading/unloading regions may be visited on a given trade route. In addition, we model the variable speed of a vessel in greater detail, by allowing the speed of a vessel on a trade route voyage to differ from the speed on the subsequent ballast sailing.

The FDP is usually presented as a deterministic problem in the academic literature, however, in reality shipping companies operate in an environment that is highly uncertain and constantly changing. Hence, the execution of a predetermined fleet deployment plan is often subject to changes, or disruptions, caused by unforeseen events. According to Brouer et al. [3], common events leading to disruptions in maritime transportation are strikes in ports, bad weather, congestion in passageways, and mechanical failures. Other, but less frequent events include piracy and crew strikes on vessels. All types of events may lead to disruptions, such as delays, or in some cases the need to charter an additional vessel. The occurrence of disruptions are common in a global shipping network; a study conducted by Notteboom [15] showed that as many as 70 to 80 % of the global shipping lines were disrupted in some way.

There are significant economic impacts associated with disruptions in liner shipping, and Kjeldsen [11] highlights that this has an effect on two fronts. First, there are costs to the shipping company such as increased bunker cost, increased port costs, charter costs of extra ships and cargo

space, and intangible costs (e.g., goodwill, loss of customers). Kjeldsen [11] exemplifies this by showing that increasing the service speed of an 6,600 TEU container vessel from 18 to 24 knots may increase fuel consumption by up to 130 tons per day. With the fuel prices of USD 650 per ton at the time the research was conducted, the higher speed costs USD 84,500 extra for each day it has to be maintained. In addition to the cost to the liner shipping company, there are significant costs associated for the customers whose cargo are onboard the vessels subject to disruption. Estimates given by Notteboom [15] show that a vessel carrying 4,000 TEUs traveling from the Far East to Belgium may lead to extra costs for its customers of at least EUR 57,000 per day it is delayed.

Yu and Qi [24] classify approaches to handle disruption management into two stages: in-advance planning and real-time re-planning. The purpose of in-advance planning is to find an optimal plan while taking future uncertainties into account. This may be referred to as robust planning or robust optimization. Here, future uncertainties can be modeled by a set of scenarios, and the goal of robust optimization is to generate a plan that is "good" for many of the possible future outcomes [24]. Real-time re-planning aims to revise the original plan previously found whenever needed during the period of execution. This may for instance be when an event occurs that leads to a disruption which causes the original plan to become undesirable.

There have been only a few papers published on incorporating robustness in tactical planning within maritime transportation. Christiansen and Fagerholt [4] present a ship scheduling problem concerned with the pickup and delivery of bulk cargoes within given time windows. Here, ports are closed at night and during weekends, and loading/discharging may take several days. The objective is to make robust schedules that are less likely to result in ships staying in ports over the weekend, with weather and port service times treated as uncertain. Robustness is introduced in the mathematical model by imposing penalty costs for arrivals at "risky" times, i.e. close to weekends. In Halvorsen-Weare and Fagerholt [8], robust solutions to a supply vessel planning problem are created. The model is based on a previous study, Halvorsen-Weare et al. [9], but several robustness strategies are applied to the vessel routes to make them more resilient to disruptions such as delays and bad weather. A vessel routing and scheduling problem under uncertainty in the liquefied natural gas business is studied in Halvorsen-Weare et al. [10]. In the problem considered, an LNG producer is responsible for transportation from production sites to customers all over the world. The aim is to create vessel schedules that are more robust against uncertainty in sailing times due to changing weather. Four robustness strategies are proposed and evaluated in a simulation-optimization framework. In addition, a re-route optimization procedure is called whenever plans are needed to be adjusted. There are a few additional articles that considers robustness in maritime planning problems, including Zeng et al. [25], Kjeldsen [11] and Alvarez et al. [1].

Despite having a robust initial plan, disruptions may cause the initial FDP to become undesirable. In these situations, it is necessary to have a set of recovery strategies in order to minimize the

impact of these disruptions on the remainder of the plan. Kjeldsen [11] presents a mathematical model for simultaneous rescheduling of vessels and cargo in container liner shipping in the event of disruptions. The objective of the model is to construct a set of vessel schedules and cargo routings that allow resumption of the scheduled service at the end of the planning period at minimal costs. To solve the problem, they propose a large neighborhood search heuristic. Brouer et al. [3] study the vessel schedule recovery problem (VSRP) for container liner shipping, and present a MIP model for handling disruptions. The VSRP evaluates a current disruption scenario and selects the action that balances increased bunker consumption, impact on cargo, and customer service level by using multi-criteria optimization in the objective function. A computational study conducted on four real-life cases reveals potential cost savings of up to 58 % compared to the real-time solution chosen by the case company. Further, Qi [17] takes the ideas of Brouer et al. [3] and develops a model that handles recovery of multiple vessels in a network after a major disruption. Two different decisions are handled by the model: vessel routing and speeding decisions, and the container flow. The model extends the work done by Brouer et al. [3], which assumes that containers will be transported by their original vessels in the event of a disruption, by including the possibility of container re-routing.

The purpose of this paper is to study the effect of introducing a set of robustness strategies to the FDP faced by a RoRo liner shipping company. The paper includes the following contributions: (1) We present a new mathematical model of the FDP in RoRo shipping that models speed decisions, cargo types, and ship capacity restrictions on each voyage more realistically than previous models in the literature. (2) We test several robustness strategies from the literature that have not previously been applied together, and apply them to the FDP in RoRo liner shipping. (3) We propose several recovery strategies to be used in real-time planning when the initial plan needs be re-optimized. The effect of both the robustness and recovery strategies are tested in a simulation framework, where the execution of the FDP is simulated with disruptions randomly occurring throughout the planning period.

The remainder of the paper is structured as follows: In Section 2 a detailed description of the FDP is given, followed by a mathematical formulation. Then a set of strategies to add robustness to the FDP is presented in Section 3, before a rolling horizon heuristic is presented in Section 4. In Section 5 we describe a simulation framework to test the effect of the proposed robustness strategies, before a comprehensive computational study is given in Section 6. Finally, some concluding remarks are given in Section 7.

2. Problem description and mathematical model

Let \mathcal{R} denote the set of trade routes operated by the company, and let n_r be the number of times trade route r should be sailed during the planning horizon. For each trade route r , let

(r, i) denote a voyage on that trade route, for $i \in \{1, \dots, n_r\}$. The fleet deployment problem can then be defined on a graph $G = (\mathcal{N}, \mathcal{A})$, where each node $(r, i) \in \mathcal{N}$ represents a voyage, and arc $((r, i), (q, j)) \in \mathcal{A}$ represents a vessel sailing voyage (r, i) , then sailing ballast between the end of trade route r and the beginning of trade route q , before sailing voyage (q, j) . The planning period may be divided into a set of months \mathcal{M} , where each month consists of a set of consecutive days. Let E_{ri} and L_{ri} be the earliest and latest day voyage (r, i) is preferred to begin, and let \mathcal{N}_m be the set of voyages (r, i) where E_{ri} is a day in month $m \in \mathcal{M}$. If the voyage starts after L_{ri} , a penalty cost C^P is incurred each day that the start of a voyage is delayed, up to the upper limit for delays denoted by the parameter D^{MAX} .

The voyages are serviced by the shipping company's fleet of vessels, represented by the set \mathcal{V} . Each vessel $v \in \mathcal{V}$ has an initial position denoted by the node $(o(v), 1)$, an artificial end node $(d(v), 1)$, and a set of nodes $\mathcal{N}_v \subseteq \mathcal{N} \cup \{(o(v), 1), (d(v), 1)\}$ representing the voyages that can be sailed by the vessel. Since a vessel can be undertaking a voyage at the start of the planning period, the parameter $E_{o(v)1}$ denotes the earliest day vessel v is available. Due to port or cargo compatibility restrictions, not all vessels can service all trade routes, and we denote the subset of arcs in the graph that can be traversed by vessel v , \mathcal{A}_v . In addition, each vessel may sail a given voyage at different speeds. As shown by [2], the cost of sailing a given distance is a non-linear convex function of the speed of the vessel. We model this cost function by using a piece-wise linear function, where the break points correspond to a set of sailing speed options, \mathcal{S}_v , from which the vessel can choose any convex combination when sailing a given voyage. As shown in Figure 1, the bunker consumption estimated by such a convex combination gives only a slight overestimate of the actual fuel consumption, while using a linear combination of the minimum and maximum speed gives a much larger overestimate. Let T_{vris} denote the number of days it takes vessel v to sail voyage (r, i) at speed option s , and C_{vris} be the cost associated with this sailing. Similarly, T_{vriqjs}^B is the number of days it takes vessel v to perform the ballast sailing between voyages (r, i) and (q, j) at speed s with a corresponding cost of C_{vriqjs}^B . Note that the sailing times and sailing costs may vary for different voyages on the same trade route to capture seasonal variations in travel times and cost.

The cargoes loaded on a vessel are transported along the trade route that is predefined by the shipping company. Each trade route consists of an alternating sequence of sailing legs and port visits along the trade route. To keep track of the cargo onboard a vessel, the term *balance categories* is introduced. Let \mathcal{B} be the set of balance categories, indexed by b . Balance categories group cargo volumes of different customers with respect to their starting and ending points. For one leg of a trade route, e.g. when a vessel servicing the trade route Europe-US-East Asia sails from the last port in Europe to the first port in the US, the vessel may be loaded with cargo from the balance categories that is (1) shipped from Europe to the US and balance categories that is




Figure 1: Illustration of why a piece-wise linear function of the fuel consumption function gives a better estimate than using a linear function. Source: Andersson et al. [2]

(2) shipped from Europe to East Asia. In other words, each route can service a predefined subset of balance categories, \mathcal{B}_r , where $\mathcal{B}_r \subseteq \mathcal{B}$.

Unlike container shipping, where containers have more or less standardized dimensions, there is a greater variety in the dimensions of the products transported in the RoRo segment. The cargoes are transported on wheels and can vary from small cars to large break bulk cargo. In the mathematical model this is handled by introducing a set of cargo segments, \mathcal{P} . A cargo segment $p \in \mathcal{P}$ may, for example, refer to small cars. The subset \mathcal{P}_b contains the cargo segments associated with balance category b . A product is fully defined through its cargo segment p and balance category b . For each month $m \in \mathcal{M}$, a demand D_{bpm} is to be transported. This demand must either be covered with the company's vessels or by renting space onboard a competitor's vessel, known as *space chartering*. The cost of space chartering is C_{bpm}^{SC} per unit, and we assume there is an upper monthly space chartering limit \bar{Q}_{bpm}^{SC} .

Each vessel has a range of decks where the cargo can be stored during transportation. The decks have different capacities and properties related to the dimension and weight of the goods. Let \mathcal{K} be the set of *capacity classes*, or deck types. Each capacity class k refers to a certain deck type. Further, the sets $\mathcal{K}_p \subseteq \mathcal{K}$ contain the set of capacity classes that can carry a cargo segment



Figure 2: Illustration of capacity classes. Source: WWLBreakbulk [23]

p . Let Q_{vrk} be the capacity in each capacity class k when vessel v services trade route r . Figure 2 illustrates the connection between decks and capacity classes.

To keep control of the total volumes of all cargo segments loaded on the decks of a vessel on the different sailing legs of a voyage, a set of *capacity groups* Ξ indexed by ξ is introduced. A capacity group ξ contains the balance categories associated with a given sailing leg of a voyage. Let this set be denoted by \mathcal{B}_ξ , and for each capacity group the sum of loads of products (b, p) , where $b \in \mathcal{B}_\xi$, should be within the capacity of the vessel. The connection between balance categories and capacity groups is illustrated in Figure 3, with a voyage sailing from Europe to East Asia via North America. A ship sailing such a voyage will usually carry cargoes from Europe to both North America and Asia, as well as from North America to Asia. Since all cargo is normally transported between regions/continents, it is in practice sufficient to define capacity groups only for these sailing legs.

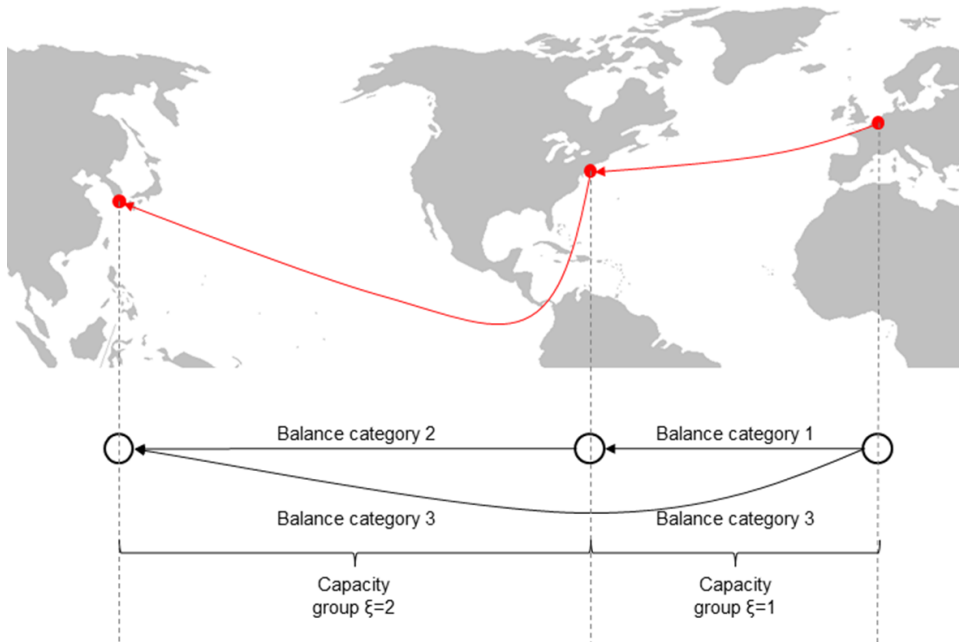


Figure 3: Illustration of a voyage with two deep sea sailing legs and the corresponding balance categories and capacity groups.

Let the variable y_{vriqj} be 1 if vessel v traverses arc $((r, i), (q, j))$, and 0 otherwise, while x_{vris} is the weight assigned to speed option s for vessel v sailing voyage (r, i) , and x_{vriqjs}^B is the weight of speed option s for the ballast sailing between voyage (r, i) and voyage (q, j) by vessel v . Let t_{ri} denote the start day of voyage (r, i) , and let d_{ri} denote the number of days voyage (r, i) is delayed compared to the latest preferred starting day, L_{ri} .

Servicing every contracted voyage in the planning period may in some cases not be possible by using the company's fleet. To ensure feasibility, we introduce the binary variable s_{ri} , which is equal to one if a contracted voyage is not serviced by the company's fleet, and is penalized with a high cost, C^S . When using this option, it is assumed that the voyage is covered by other transportation options or companies. For these cases, a corresponding capacity for each capacity class is denoted as Q_k^S , which for simplicity is calculated as

$$Q_k^S = \max_{v \in \mathcal{V}, r \in \mathcal{R}} Q_{vrk}, \quad \forall k \in \mathcal{K}.$$

Finally, the variables l_{ribpk} describe the loaded volumes of product (b, p) on capacity class k on voyage (r, i) . The volume of product (b, p) covered with space charter in month m is denoted by z_{bpm} . The mathematical model of the fleet deployment problem for the RoRo liner shipping company may then be formulated as follows:

$$\begin{aligned} \min z = & \sum_{v \in \mathcal{V}} \sum_{(r,i) \in \mathcal{N}_v} \sum_{s \in \mathcal{S}_v} C_{vris} x_{vris} + \sum_{v \in \mathcal{V}} \sum_{((r,i),(q,j)) \in \mathcal{A}_v} \sum_{s \in \mathcal{S}_v} C_{vriqjs}^B x_{vriqjs}^B + \\ & \sum_{(r,i) \in \mathcal{N}} C^P d_{ri} + \sum_{b \in \mathcal{B}} \sum_{p \in \mathcal{P}} \sum_{m \in \mathcal{M}} C_{bpm}^{SC} z_{bpm} + \sum_{(r,i) \in \mathcal{N}} C^S s_{ri} \end{aligned} \quad (1)$$

$$\sum_{v \in \mathcal{V}} \sum_{(q,j) \in \mathcal{N}_v} y_{vriqj} + s_{ri} = 1, \quad (r, i) \in \mathcal{N}. \quad (2)$$

$$\sum_{(q,j) \in \mathcal{N}_v} y_{vo(v)1qj} = 1, \quad v \in \mathcal{V}, \quad (3)$$

$$\sum_{(q,j) \in \mathcal{N}_v} y_{vriqj} - \sum_{(q,j) \in \mathcal{N}_v} y_{vqjri} = 0, \quad v \in \mathcal{V}, (r, i) \in \mathcal{N}_v, \quad (4)$$

$$\sum_{(r,i) \in \mathcal{N}_v} y_{vrid(v)1} = 1, \quad v \in \mathcal{V}, \quad (5)$$

$$\sum_{(q,j) \in \mathcal{N}_v} y_{vriqj} - \sum_{s \in \mathcal{S}_v} x_{vris} = 0, \quad v \in \mathcal{V}, (r, i) \in \mathcal{N}_v, \quad (6)$$

$$y_{vriqj} - \sum_{s \in \mathcal{S}_v} x_{vriqjs}^B = 0, \quad v \in \mathcal{V}, ((r, i), (q, j)) \in \mathcal{A}_v. \quad (7)$$

$$t_{ri} + \sum_{s \in \mathcal{S}_v} (T_{vriqjs}^B x_{vriqjs}^B + T_{vris} x_{vris}) - t_{qj} \leq M_{vriqj} (1 - y_{vriqj}), \quad v \in \mathcal{V}, ((r, i), (q, j)) \in \mathcal{A}_v. \quad (8)$$

$$t_{ri} - d_{ri} \leq L_{ri}, \quad (r, i) \in \mathcal{N}. \quad (9)$$

$$\sum_{(r,i) \in \mathcal{N}_m} \sum_{k \in \mathcal{K}_p} l_{ribpk} + z_{bpm} = D_{bpm}, \quad m \in \mathcal{M}, b \in \mathcal{B}, p \in \mathcal{P}_b. \quad (10)$$

$$\sum_{b \in \mathcal{B}_\xi} \sum_{p \in \mathcal{P}_b} l_{ribpk} - \left(\sum_{v \in \mathcal{V}} Q_{vrk} \sum_{(q,j) \in \mathcal{N}_v} y_{vqjri} + Q_k^S s_{ri} \right) \leq 0, \quad (r, i) \in \mathcal{N}, k \in \mathcal{K}, \xi \in \Xi. \quad (11)$$

$$0 \leq z_{bpm} \leq \bar{Q}_{bpm}^{SC}, \quad b \in \mathcal{B}, p \in \mathcal{P}_b, m \in \mathcal{M}. \quad (12)$$

$$y_{vriqj} \in \{0, 1\}, \quad v \in \mathcal{V}, ((r, i), (q, j)) \in \mathcal{A}_v, \quad (13)$$

$$s_{ri} \in \{0, 1\}, \quad (r, i) \in \mathcal{N}_v, \quad (14)$$

$$x_{vris} \geq 0, \quad v \in \mathcal{V}, (r, i) \in \mathcal{N}_v, s \in \mathcal{S}_v, \quad (15)$$

$$x_{vriqjs}^B \geq 0, \quad v \in \mathcal{V}, ((r, i), (q, j)) \in \mathcal{A}_v, s \in \mathcal{S}_v, \quad (16)$$

$$t_{ri} \geq E_{ri}, \quad (r, i) \in \mathcal{N}, \quad (17)$$

$$l_{ribpk} \geq 0, \quad (r, i) \in \mathcal{N}_v, b \in \mathcal{B}_r, p \in \mathcal{P}_b, k \in \mathcal{K}, \quad (18)$$

$$0 \leq d_{ri} \leq D^{MAX}, \quad (r, i) \in \mathcal{N}. \quad (19)$$

The objective function (1) minimizes costs related to operating the fleet of vessels and fulfilling the constraints below. The costs of sailing voyages and ballast between voyages with the speed profile determined by x_{vris} and x_{vriqjs}^B are included in the two first expressions. The third expression calculates the total delay costs, the fourth expression the space chartering costs, and the fifth term the cost of not servicing voyage (r, i) with a vessel from the company's fleet. Constraints (2) force all contracted voyages to be serviced, while constraints (3) – (5) ensure that each vessel travels a continuous path through the problem defining network. Further, constraints (6) and (7) ensure that a convex combination of speed options is chosen for each ship on each voyage and ballast sailing, and constraints (8) keep track of the start of each voyage. Note that constraints (8) are formulated using the big-M method, where $M_{vriqj} = L_{ri} + D^{MAX} - E_{qj} + \max_{s \in \mathcal{S}_v} \{T_{vriqjs}^B + T_{vris}\}$. Constraints (9) count the number of days a voyage is delayed, while constraints (10) ensure that the volume transported and space chartered in a given month equal the monthly demand. It should be noted that constraints (10) assume that the cargo transported on a given voyage (r, i) covers the cargo demand in the month of its earliest start time (E_{ri}) regardless of when the voyage is actually started. This is a reasonable assumption since the time windows for starting a voyage in RoRo shipping are normally fairly narrow (typically at most one week and rarely span two subsequent months). Finally, constraints (11) limit the total flow of cargoes in a capacity group on a voyage to be within the capacity of the vessel sailing it, and constraints (12)-(19) define the variable domains.

3. Adding robustness strategies to the fleet deployment model

The mathematical model (1)-(19) defines the FDP for a RoRo liner shipping company. However, the solutions obtained when solving this model may sometimes become infeasible or undesirable if disruptions occur during their execution. As stated in Notteboom [15] it is reasonable to divide such disruptions in liner shipping into events happening in port and events happening when sailing between ports. For WWL, typical disruptions in ports are congestion at the port, delay in receiving the necessary paperwork from the port authorities, and sometimes disputes with the port regarding bunkering that necessitates the presence of a surveyor to reach a settlement. These port disruptions can delay the departure from the port by anything from a couple of hours and up to several days. When sailing between ports typical disruptions are rough weather conditions and congestions in passageways. Rough weather conditions force all vessels in the area to sail at a lower speed, while congestion in a passageway causes vessels to wait longer than planned before passing through it. In both cases the effect of the disruption is that the affected voyages takes longer than anticipated, causing the vessels to finish their voyages later than planned.

Finishing a voyage later than planned is unfortunate. Since a voyage sails from a loading region to an unloading region, and since (most of) the cargo is in most cases already onboard the ship when disruptions occur, it is not considered practical to abort or alter an ongoing voyage. Therefore these delays are usually accepted by the customers. However, what the shipping company wants to avoid is that these delays cause a knock-on effect on to subsequent voyages. Thus, what is most important when a disruption occurs is to ensure that as few as possible of the future voyages are affected. To achieve this, the case company has mainly two recourse actions. The first is to adjust the speed of the vessel to reclaim some, or all, of the lost time. Note that this is only possible if the original voyage was planned at a speed lower than the maximum speed. The effect of such an action is also diminishing the closer to the end of a voyage the vessel is. The second recourse action is to re-plan the coming voyages for the fleet of vessels, so that the delayed vessel continues onto another voyage than originally planned once the delayed voyage is completed, while some other vessel in the fleet sails the originally planned subsequent voyage. Note that these types of vessel-voyage swaps become more difficult the tighter the original plan is, since some slack in the plan is necessary for this to be possible. Such swaps may also affect the amount of cargo that can be carried on a voyage since the vessels may have different capacities to carry different types of cargo.

In this section we propose a set of strategies to make the fleet deployment plans more robust and better prepared for disruptions. For modeling purposes it is assumed that both port and sailing disruptions affect the values of the parameters T_{vris} and T_{vriqjs}^B in the mathematical model.

3.1. Adding extra sailing time

As presented in Halvorsen-Weare et al. [10], planning some extra sailing time for each voyage is a straight-forward way of introducing slack in the vessel schedules. This is done by multiplying the sailing times parameters T_{vriqjs}^B and T_{vris} by some value greater than 1. When considering the planned extra sailing times against the actual sailing times of the voyages, the result is that we have introduced slack in the vessel schedules.

3.2. Rewarding early arrivals

Rewarding properties that introduce slack in solutions is another way of creating more robust schedules. Halvorsen-Weare and Fagerholt [8] present three possible ways of introducing robustness in a routing problem. They reward: (1) each day a vessel is idle, (2) each vessel that has one idle day during the week, or (3) each vessel that sails no more than two voyages during the week. In this paper we use a similar approach by rewarding vessels that arrive before the start of the time window, E_{ri} . To accomplish this, we need to introduce a few additional variables and parameters to the mathematical model. Let t_{ri}^R be the number of days rewarded for early arrival before starting voyage (r, i) . Since slack is assumed to give diminishing returns as it increases, an upper limit, R^{MAX} , is used on how many days of slack can be rewarded. Furthermore, let δ_{ri}^R be 1 if voyage (r, i) can be started early enough to qualify for a reward, and 0 otherwise. The reward given per time unit of early arrival is denoted as Λ^R .

Figure 4 illustrates how rewarding early arrivals works in practice. If a vessel arrives at the start of voyage (r, i) before the start of the time window E_{ri} , any slack up to R^{MAX} days will be awarded with a factor of Λ^R per time unit (see scenario (a)). Any arrivals after E_{ri} , shown by scenario (b), do not qualify for any rewards.

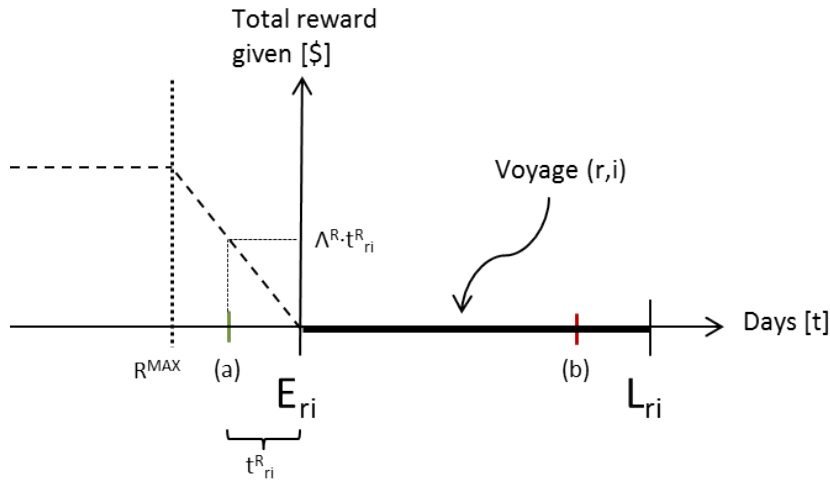


Figure 4: Illustration of (a) rewarding and (b) not rewarding an early arrival. In scenario (a), the vessel arrives at the start of voyage (r, i) t_{ri}^R days before the time window starts, which results in a reward of $\Lambda^R \cdot t_{ri}^R$. In scenario (b), the vessel arrives late in the time window and no reward is given.

The following sets of constraints must be added to the mathematical model:

$$t_{ri} + \sum_{s \in \mathcal{S}_v} (T_{vriqjs}^B x_{vriqjs}^B + T_{vris} x_{vris}) \quad v \in \mathcal{V}, ((r, i), (q, j)) \in \mathcal{A}_v, \quad (20)$$

$$+ t_{qj}^R - E_{qj} \delta_{qj}^R \leq M_{vriqj}^R (2 - y_{vriqj} - \delta_{qj}^R),$$

$$0 \leq t_{ri}^R \leq R^{MAX} \delta_{ri}^R, \quad (r, i) \in \mathcal{N}, \quad (21)$$

$$\delta_{ri}^R + s_{ri} \leq 1, \quad (r, i) \in \mathcal{N}. \quad (22)$$

$$\delta_{ri}^R \in \{0, 1\}. \quad (23)$$

Constraints (20) calculate the amount of slack between two voyages, where the parameter $M_{vriqj}^R = L_{ri} + D^{MAX} + R^{MAX} + \max_{s \in \mathcal{S}_v} \{T_{vriqjs}^B + T_{vris}\}$. Further, constraints (21) limit the number of slack days rewarded, while constraints (22) ensure that we only reward voyages sailed by the vessels in the fleet, and constraints (23) define the domain of the δ_{ri}^R variables. In addition, the following term is subtracted from the objective function (1):

$$\sum_{(r,i) \in \mathcal{N}} \Lambda^R t_{ri}^R \quad (24)$$

3.3. Penalizing risky voyage start times

A similar approach to rewarding properties that provide robustness is to penalize characteristics of a solution that reduce robustness. This method was used by Christiansen and Fagerholt [4], where a pickup and delivery ship scheduling problem with multiple time windows is studied, and robust solutions are created by penalizing arrivals that are close to the end of harbor opening hours, i.e. close to nights and weekends. In this paper we suggest a similar approach for creating robust fleet deployment, where we penalize start times close to the end of time windows. This may alter vessel schedules or increase planned sailing speeds, both of which are strategies to increase slack in the solution and thereby its robustness.

To extend the mathematical formulation (1)-(19) to penalize risky start times, we introduce additional variables and parameters. Let variable t_{ri}^P represent the number of days penalized for risky start time of voyage (r, i) . For each voyage we apply an upper limit P_{ri}^{MAX} to the number of days penalized, which is the smaller of a global upper limit P^{MAX} and the width of the voyage's time window $|L_{ri} - E_{ri}|$. The penalty given per time unit of risky start time is denoted as Λ^P .

Figure 5 illustrates how risky start times are penalized. If a vessel plans to start a voyage at the time given by (a), no penalty is given. However, if (b) is the planned start time, this is within the later part of the time window. We classify this as a risky start time as it is exposed to even minor delays, and a penalty with a factor of Λ^P is suffered.

To model the penalizing of risky start times, the original model defined by (1)-(19) is extended by adding the following constraints.

$$t_{ri} - d_{ri} - (L_{ri} - P_{ri}^{MAX}) - t_{ri}^P \leq 0, \quad (r, i) \in \mathcal{N}, \quad (25)$$

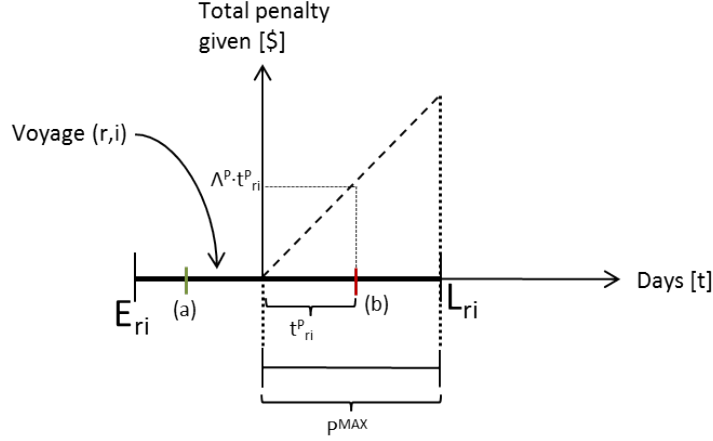


Figure 5: Illustration of (a) not penalizing and (b) penalizing the start time of a voyage. In scenario (a), the vessel is planned to start voyage (r, i) more than P^{MAX} days before the end of the time window. This start time is not considered risky and is therefore not penalized. In scenario (b), the start time of voyage (r, i) is t_{ri}^P days into the risky area of start times for the voyage, which results in a penalty of $\Lambda^P \cdot t_{ri}^P$.

$$0 \leq t_{ri}^P \leq P^{MAX}, \quad (r, i) \in \mathcal{N}. \quad (26)$$

Constraints (25) set the number of days penalized for risky start time of voyage (r, i) to be greater than or equal to the start time of the voyage minus the time when penalties start occurring, while constraints (26) set the appropriate range of t_{ri}^P .

Also, to allocate penalties, the following term is added to the objective function (1):

$$\sum_{(r,i) \in \mathcal{N}} \Lambda^P t_{ri}^P \quad (27)$$

4. Rolling horizon heuristic

The problem presented in Section 2 is an extension of the vehicle routing problem with time windows (VRPTW) which is an NP-hard problem ([20]), and thus, by restriction, the problem studied in this paper is also NP-hard. In fact, finding a feasible solution to the VRPTW (and thus also to the problem presented in this paper) is an NP-complete problem (as proved by Savelsbergh [19]). Because the problem is NP-hard and the number of ships and trade routes is large for realistic instances of the FDP, it is not possible to solve realistic instances to optimality within a reasonable amount of time.

A popular approach to obtain good solutions to hard planning problems with long planning horizons is to use a rolling horizon heuristic (RHH). This heuristic approach was successfully implemented by Andersson et al. [2] to solve a maritime fleet deployment problem for a 6-10 months planning horizon. Due to the problem similarities, it is reasonable to believe that a rolling horizon heuristic has the potential of being an excellent solution method for the problem studied in this paper.

4.1. The rolling horizon heuristic mechanisms

The main idea behind a RHH is to divide the problem into a set of subproblems along the time axis, and then solve the subproblems iteratively in chronological order, while taking the decisions made in the previous subproblems into account. Each subproblem usually includes a primary period and a forecasting period. The decisions made in the primary period are partly or fully fixed in the next iteration, while the forecasting period is included to ensure that the decisions made in the primary period do not have adverse affects later in the planning horizon. When solving the subproblem, the properties of the original problem are kept intact in the primary period, while properties in the forecasting period are often simplified. This simplification can, for example, be a relaxation of integer variables. The RHH iterates through all the subproblems until all decisions for the whole planning horizon have been made.

Figure 6 illustrates the mechanism of a RHH, and is inspired by similar figures in Mercao and Fontan [14], Rakke et al. [18], and Andersson et al. [2]. Let us denote K as the number of iterations to be performed. At iteration k , the subproblem H_k is solved. Further, $S_P(H_k)$ and $S_F(H_k)$ describe the primary period S_P and the forecasting period S_F in subproblem H_k , respectively. When the computation in each iteration has been done, the solutions in $S_P(H_k)$ are fixed according to a predefined fixing strategy, and the primary and forecasting periods are updated to $S_P(H_{k+1})$ and $S_F(H_{k+1})$. The new subproblem is chosen by moving forward a given number of time units equal to the length of the primary period, and then subproblem H_{k+1} is solved in a similar manner. In the last period of the horizon, no forecasting section is used, and all decisions made are final.




Figure 6: The rolling horizon heuristics mechanism.

4.2. Rolling horizon heuristics applied to the fleet deployment problem

As stated in Section 2, the time periods in the mathematical model are divided into a set of months, \mathcal{M} . In our implementation of the RHH we have chosen to let K be equal to the number of months in the planning problem, and to let each subproblem H_k consist of a central period of one month, and a forecasting period of the following two months. Let $\mathcal{M}(H_k)$ be the set of months in subproblem k , including the months prior to month k . Subproblem H_k is then the subset of the

mathematical model defined by constraints (1)–(19) consisting of the variables and constraints that are defined for the months $m \in \mathcal{M}(H_k)$, and the voyages $(r, i) \in \bigcup_{m \in \mathcal{M}(H_k)} \mathcal{N}_m$. However, to make the model easier to solve we simplify the forecasting period by relaxing the binary requirements on all variables belonging to those months.

Once a subproblem is solved, the values of the routing variables y_{vriqj} from the central period are fixed to their current value, while the variables related to speed (x_{vriqj}), load (l_{ribpk}), start times (t_{ri}), and a vessel’s final destination ($y_{vrid(v)1}$) may be adjusted in later subproblems. This gives the model the opportunity to speed up vessels to arrive earlier at the start of a given voyage, re-distribute loads, and move start times within their time windows, in earlier months to make better decisions in later months. Thus, the final value of these variables are not decided until the RHH is solved for the last month.

5. A simulation framework for evaluating disruption management strategies

To test the effectiveness of the proposed robustness strategies suggested in Section 3 as well as some recovery strategies, a simulation framework was developed. The framework simulates the execution of the FDP and introduces realistic disruptions by randomly adding events. We consider two types of events: 1) events affecting the sailing times of a voyage and 2) events that affect the time in a port. To capture the possibilities of taking recovery actions during the execution of a plan, the framework combines simulation with optimization by calling re-planning procedures when certain conditions during the simulation are met. We consider two types of recovery strategies: 1) adjusting the speed of a vessel, and 2) re-solving the FDP. The first strategy is used by individual ship captains who can decide to deviate from the planned speed on a voyage in order to keep the schedule. The second strategy involves the shipping company planners as this recovery strategy may affect the sailing plan of many ships simultaneously.

A flow chart for the developed simulation program is presented in Figure 7. The input solution is obtained by solving the FDP for a given problem instance. This may or may not include any of the robustness strategies defined in Section 3. We also predefine the event scenarios, i.e. the set of events occurring during each simulation. Each event is associated with a given sailing leg, or a given port, happens in a specific time period, and causes a delay of a given number of days for any ship sailing that leg or visiting that port within the event window.

A simulation starts on day $t = 0$, and iterates through each day t in the planning horizon to see if any disruption happens at that day. If such an event occurs on day t , we update vessel schedules with any impacts from the event. This means that sailing parameters are adjusted if any voyage or ballast sailings are affected, and extra port time is added in case some vessels are on their way to the relevant port(s). These new schedules are then evaluated, and the re-planning procedure is called if the new schedules satisfy a certain criterion. The re-planning procedure treats all parts of

the schedules up to day t as fixed, and tries to mitigate any occurred disruptions as effectively as possible. This is done by re-solving the mathematical model, with parts of the model fixed, using the RHH. The new solution to the FDP may have taken recovery actions, such as increasing the speed of the vessels on some voyages (including the remainder of their current voyage), swapping future voyages between vessel schedules, and/or reallocate cargo loads on future voyages. After re-optimizing, the vessel routes and schedules are updated with the effects of the recovery actions. Then, we check if there are more days left in the planning period. If there are, we update $t = t + 1$ and move on to the next day. If this is the last day, we terminate the simulation. Finally, the resulting solution and a set of simulation statistics, such as the total simulated costs, are collected as output from the simulation.

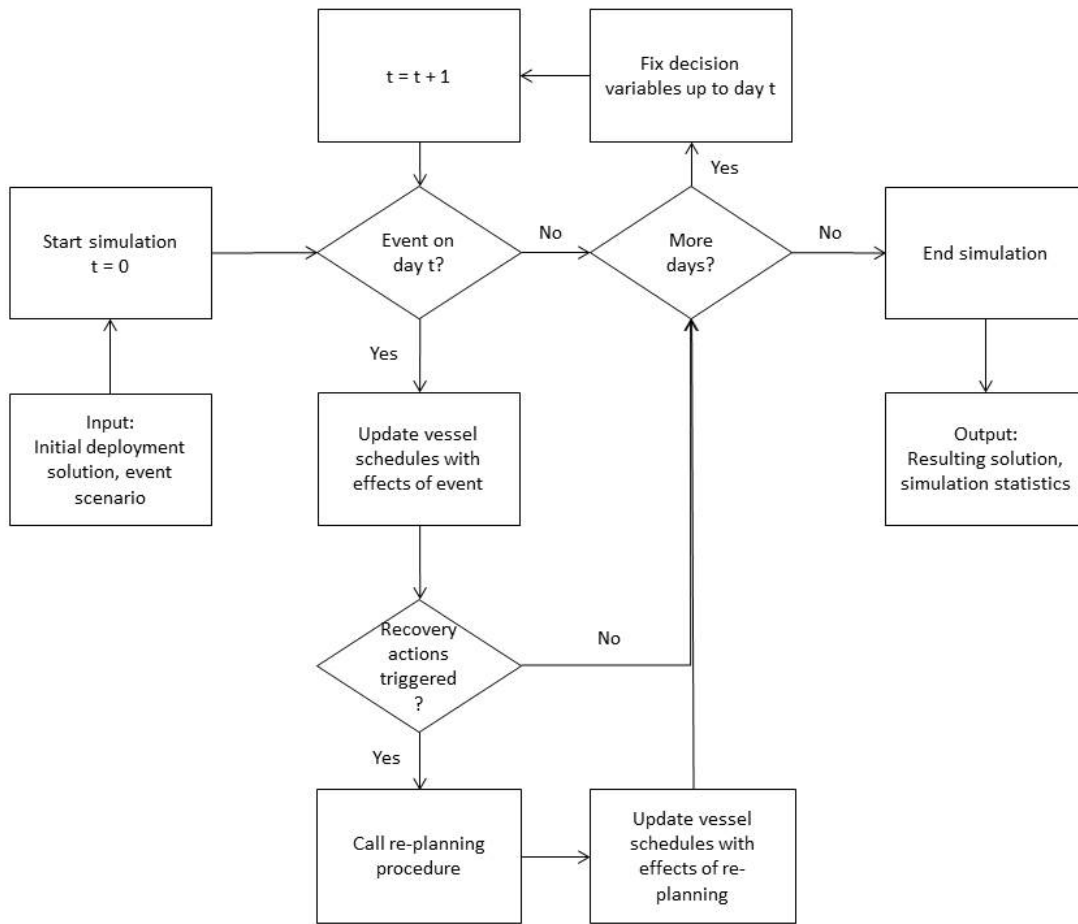


Figure 7: Flow chart for the simulation program.

6. Computational Study

In this section different methods of handling disruption management in RoRo liner shipping, i.e. the combination of robustness and recovery strategies, is tested. We begin in Section 6.1 by

introducing the test instances, before testing the RHH presented in Section 4 in a static setting and comparing it with results of solving the full mathematical model as a MIP using commercial software in Section 6.2. The main part of the computational study, presented in Section 6.3, tests how the various robustness and disruption management strategies performs in a dynamic and stochastic setting mimicking the real world. The MIP and RHH models were solved using Gurobi Optimizer version 6.0 [7], and were run on a HP DL165 G6 computer with 2 x AMD Opteron 2431 2.4 GHz processor, and 24 GB of RAM.

6.1. Description of the problem instances

The test data used for testing the RHH’s performance consists of five different problem instances provided by WWL, with different size and complexity. The simplest and smallest problem instance contains only two vessels and five voyages to be serviced over a planning period of three months, while the largest and most complex problem instance has 24 vessels, 222 voyages and a planning period of nine months. A summary of the characteristics of the different problem instances is presented in Table 1. The problem instances are named so that an instance named "S5_V52_T7" has five ships, 52 voyages and a planning period of seven months. In addition, the table states the number of speed profiles, cargo segments, and the width of the time windows for start of the voyages, for each instance.

Table 1: Problem instances characteristics

Instance	Vessels	Voyages	Months	Speed profiles	Cargo segments
S2_V5_T3	2	5	3	1	1
S5_V52_T7	5	52	7	3	4
S5_V77_T8	5	77	8	2	4
S16_V109_T9	16	109	9	3	4
S24_V222_T9	24	222	9	3	4

The width of the time windows has a great impact on the complexity of the problems as increased width increases the number of possible combination of voyages each vessel can sail, and thus increases the solution space. Each voyage has one target day for when it should preferably start that will remain the same, while the earliest and latest start dates on each side of the target day is modified as described below. It is assumed that the voyages in the first months have less flexibility than the voyages in the later months of the planning period. Therefore, the first months consistently have tighter time windows in all modifications of the test instances described in the following:

Modification 1 (M1) For all the months in the planning period, the earliest and latest start of the voyages are the target day ± 1 day, resulting in a time window of three days.

Modification 2 (M2) For the two first months in the planning period, the earliest and latest start dates of the voyages are the target day ± 1 day, resulting in a time window of three

days. For the rest of the months in the planning period, the earliest and latest start dates of the voyages are the target day ± 2 days, resulting in a time window of five days.

Modification 3 (M3) For the two first months in the planning period, the earliest and latest start dates of the voyages are the target day ± 1 day, resulting in a time window of three days. For the two subsequent months the earliest and latest start dates of the voyages are the target day ± 2 days, resulting in a time window of five days. For the rest of the planning period the earliest and latest start dates of the voyages are the target day ± 3 days, resulting in a time window of seven days.

We apply all three modifications to each of the five test instances described in Table 1, creating a total of 15 instances. We denote each problem instance by adding M1, M2 or M3 to the end of each problem instance’s name.

6.2. Testing the rolling horizon heuristic

To test the quality of the RHH we have tested it on the instances presented in Section 6.1, and compared the solutions with those obtained by solving the full MIP model defined by constraints (1)–(19) directly using Gurobi. In this part of the testing no robustness strategies are included. The results are summarized in Table 2. The table reports the number of unserved voyages (UV), i.e. the voyages that are not serviced by the shipping company’s own fleet, the gap between the best feasible solution found by each method and the best lower bound obtained by the MIP solver (Gap), and the solution time (Time). For the MIP solver an upper time limit of 10,000 seconds was set, while for the RHH, the maximum solution time allowed to solve each subproblem was set to $\frac{10,000}{\# \text{ Months}}$ seconds. The reason for this is to ensure that the total computing time of the RHH does not exceed 10,000 seconds in total.

Table 2: Results of solving the different problem instances both as a MIP and using the RHH

Instance	MIP			RHH		
	UV	Gap(%)	Time (s)	UV	Gap(%)	Time (s)
S2_V5_T3_M1	0	0.00	0	0	0.00	0
S5_V52_T7_M1	0	0.00	17	0	0.02	7
S5_V77_T8_M1	3	0.00	29	3	0.10	15
S16_V109_T9_M1	4	0.07	10,000	4	0.74	66
S24_V222_T9_M1	15	101.00	10,000	8	20.00	1,828
S2_V5_T3_M2	0	0.00	0	0	0.00	0
S5_V52_T7_M2	0	0.00	352	0	0.09	36
S5_V77_T8_M2	3	0.38	10,000	5	45.00	75
S16_V109_T9_M2	3	0.88	10,000	3	1.55	121
S24_V222_T9_M2	61	940.00	10,000	4	4.27	3,420
S2_V5_T3_M3	0	0.00	0	0	0.00	0
S5_V52_T7_M3	0	0.00	1,145	0	0.06	40
S5_V77_T8_M3	3	0.67	10,000	3	0.72	86
S16_V109_T9_M3	3	1.90	10,000	3	2.30	210
S24_V222_T9_M3	89	1,413.00	10,000	4	4.80	5,520

For the two smallest sets of instances S2_V5_T3_MX and S5_V52_T7_MX the MIP solver is

able to find the optimal solution relatively quickly, and for the instances S5_V77.T8_MX and S16_V109.T9_MX the MIP solver is less than 2 % from the optimal solution after 10,000 seconds. However, for the largest set of instances, S24_V222.T9_MX, the MIP model cannot find any decent primal solution within the time limit.

The RHH is able to find good solutions with an optimality gap of less than 5 % for all instances except S5_V77.T8_M2 and S24_V222.T9_M1. For the latter instance the optimality gap may be explained by a poor lower bound obtained by the MIP solver, since it has not been able to explore many branch-and-bound nodes within the time limit. For S5_V77.T8_M2 the poor optimality gap comes from the two additional unserved voyages that are present in the RHH solution, compared to the MIP solution.

Comparing the solution times we see that the RHH is able to produce much higher quality solutions within the time limit, when the test instances become large. In addition, the RHH spends less than one hour (3,600 seconds) to solve all but the largest test instance, which takes 5,520 seconds to solve. The reason why the RHH spends less time than the allowed 10,000 seconds is that while the subproblems early in the planning horizon usually run until the time limit, the subproblems late in the planning period are solved to optimality much faster, since most of the variables in the problem are fixed at this point. This indicates that even better solutions might be obtained if we allocate more time to solve the subproblems early in the planning horizon.

Overall, the RHH provides good solutions to the FDP, and is capable of finding high quality solutions within the given time limits for large problem instances where solving the full mathematical model with commercial software is not possible. Thus, in the following computational testing we use the RHH both to create an initial FDP, and as a recovery tool when re-planning the FDP after disruptions have occurred during the simulation.

6.3. Testing the disruption management approaches

In this section we use the simulation-optimization framework presented in Section 5 to 1) evaluate solutions obtained by applying the alternative robustness strategies and 2) evaluate the effectiveness of using the RHH as a re-planning tool, to take recovery actions. The test instance S24_V222.T9_M2 has been used for all tests in this section since it is the largest and most realistic of the test instances presented in Section 6.1.

As stated in Section 5 we consider two types of disruptions; those affecting the time it takes to sail a voyage, and those affecting the time spent in a given port. Table 3 shows the probabilities of a disruption occurring on a given voyage or at a given port on a given day, the impact of the event, the average number of sailing and port events that occurred during the planning horizon in the simulated scenarios, and the corresponding standard deviation. These numbers do not define how many vessel schedules are affected by each event, as zero, one or several voyages may be sailing the affected trade route, or be sailing towards the affected port.

Table 3: Information regarding the values used to generate disruptions as well as data regarding the average number and standard deviation of events generated by the simulation framework. Listed under 'Impact' are the number of days added to the sailing time of a voyage if the delay happens in a port, and the number that is multiplied with the sailing time if the delay happens while sailing.

Event	Probability	Impact	Average number of events	Standard deviation (%)
Port	0.01	2 [days]	25.3	14.32
Voyage sailing	0.02	1.10 [-]	49.7	12.00

6.3.1. Evaluating the robustness strategies

We begin by testing the three robustness strategies presented in Section 3, as well as one that combines all three, and compare their solutions to those created by applying the RHH to the original mathematical model. In this part of the testing we only allow speed adjustment on individual vessels as a recovery action, and thus we test each solution's ability to absorb disruptions without changing the plan. Table 4 gives a description of these approaches and how they are denoted in the remainder of this computational study, while Table 5 shows the parameter values used for the three robustness strategies. For the COMBINED robustness strategy we slightly scale down the presented parameter values to incorporate a weighting effect. This was found to work well during the initial testing. The revised parameter values for the COMBINED strategy are presented in Table 6. It should be emphasized that setting the parameter values is a major undertaking in order to balance different criteria. The values shown in Tables 5 and 6 are carefully chosen after preliminary testing and seem to give reasonable solutions for the problem studied.

Table 4: List of evaluated robustness strategies

Robustness strategy	Description
BASIC	Mathematical model defined by constraints (1)–(19)
INCREASE	BASIC with increased sailing times (Section 3.1)
REWARD	BASIC with rewarding early arrivals (Section 3.2)
PENALIZE	BASIC with penalizing risky start times (Section 3.3)
COMBINED	A combination of INCREASE, REWARD and PENALIZE

Table 5: Robustness parameter values

Parameter	Description	Value
-	Sailing time adjustment	1.02
R^{MAX}	The maximum number of days rewarded for early arrival	2
Λ^R	Reward given per day of early arrival	150,000 USD
P^{MAX}	The maximum number of days penalized for risky arrival	2
Λ^P	Penalty suffered per day of risky arrival	100,000 USD
D^{MAX}	Number of days a voyage can be planned to be delayed	0
C^P	Daily delay penalty cost	200,000 USD

Table 7 shows the planned operating costs of the solutions obtained by the different robustness strategies as a percentage of the BASIC planned operating cost. These costs are calculated as the costs of the solutions given by the objective function (1) when the costs of any unserved voyages have been deducted. Also, any costs or revenues associated with robustness rewards or penalties

Table 6: Revised parameter values for the COMBINED robustness strategy

Parameter	Description	Value
-	Sailing time adjustment	1.01
Λ^R	Reward given per day of early arrival	50,000 USD
Λ^P	Penalty suffered per day of risky arrival	50,000 USD

are excluded in the planned operating costs since these also are artificial costs/rewards. This is done because all solutions had the same number of unserved voyages, and we want to capture the difference in the operating costs of the company’s vessel fleet when applying the different robustness strategies. Adding these robustness strategies may increase the planned operational costs, since the set of feasible solutions to the problem is reduced. However, we see that none of the robustness strategies add more than a modest 5 % to the planned operating costs in this instance.

Table 7: The planned costs of solutions obtained with the robustness strategies INCREASE, REWARD, PENALIZE and COMBINED expressed as a % of the BASIC planned costs.

	BASIC	INCREASE	REWARD	PENALIZE	COMBINED
Plan. cost (%)	100.00	102.84	105.00	101.42	103.67

Table 8 shows the average simulated costs (SC) and average total days of delay (D) over 10 simulation runs for each of the robustness strategies. The simulations are conducted with two different settings. For the simulations denoted by *No re-planning*, we do not allow speed adjustments or any other recovery actions. This setting makes it possible to observe how slack in the solutions can absorb the impact of disruptions directly. For the simulations denoted by *Speed re-planning*, we allow speed adjustments to mitigate the impacts of disruptions. The intention behind this setting is to mimic how we assume individual vessels respond to events in reality, without any centrally coordinated actions taken from the shipping company itself. These speed adjustments may be conducted every day during the simulation and ensure that any robustness that comes with sailing voyages at low speeds is captured.

Table 8: Average simulated costs (SC) and total days of delay (D) for the different robustness strategies over ten simulations. The simulated costs are expressed as % of the planned costs for the BASIC initial solution.

	BASIC		INCREASE		REWARD		PENALIZE		COMBINED	
	SC(%)	D	SC(%)	D	SC(%)	D	SC(%)	D	SC(%)	D
<i>No re-planning</i>	155.02	604	133.72	329	117.06	113	143.38	454	116.15	118
<i>Speed re-planning</i>	121.50	197	118.93	142	114.86	82	115.40	112	113.41	77

The simulated costs in Table 8 are all given as a percentage of the planned cost for the initial BASIC solution from Table 7. Compared to the planned costs in Table 7, the simulated costs capture the costs of disruptions, i.e. the impacts on the vessel schedules from disruptions such as increased sailing and port times. These impacts potentially include longer sailing times, higher speeds, extra space chartering of cargoes and, perhaps most importantly, the costs of voyage delays. The results show that the simulated costs are higher than the planned costs, meaning that

disruptions have occurred during the planning period and caused extra costs. Also, the results show that the simulated costs when speed adjustments are allowed are substantially lower than when no recovery actions are enabled. This highlights the importance of incorporating speed decisions in the FDP model, since correctly modeling the speed/cost trade-off causes the vessels to reduce their speed on some sailing legs, thus saving money and at the same time increasing robustness of the FDP.

By incorporating robustness strategies when planning, we observe a significant improvement in performance in terms of simulated costs and amount of delays. Despite higher planned costs than the BASIC solution, as presented in Table 7, all of the robustness solutions have lower average simulated costs. When allowing *Speed re-planning*, the COMBINED strategy reduces the average costs and delays compared with the BASIC strategy by 7 % and 61 %, respectively. It is also worth noticing that the two solutions with the highest planned costs, COMBINED and REWARD, are the solutions that provide the lowest simulated costs, both when simulating with the *No re-planning* and the *Speed re-planning* settings.

6.3.2. Evaluating re-planning as a recovery action

When including the re-planning procedure in the simulation framework, we must determine when it should be used. As liner shipping companies have different preferences, this choice may differ from company to company. For instance, one company may accept some delays to maintain its cost-efficient plan. Other companies, however, aim to reduce delays as much as possible to keep highly demanding customers satisfied. We have tested three different trigger strategies for when to use the re-planning procedure as a recovery strategy:

1. at least one day delay for a voyage is anticipated with the current plan
2. at least three days delay for a voyage is anticipated with the current plan
3. at least five days delay for a voyage is anticipated with the current plan

Table 9 shows the results obtained by combining the trigger strategies proposed above with the BASIC and COMBINED robustness strategies in the simulation framework. All the reported results are average numbers based on five simulations for each combination of trigger and robustness strategy. Each row contains the results for the three different trigger strategies previously proposed. The simulated costs of the solutions are reported as a percentage of the BASIC planned costs. The other three columns for each strategy report the average total delay incurred (D), number of times the re-planning recovery procedure is called (RP), and the number of voyage swaps performed (VS), respectively.

We observe that by including the re-planning procedure in the simulation framework, significant costs savings can be achieved. For the BASIC solution, the simulated costs are, for all triggers, just above 16 % of the planned costs, compared to roughly 21 % when only speed adjustments

Table 9: Average simulated costs (SC), total days of delay (D), number of times the recovery procedure is called (RP), and the number of performed voyage swaps (VS) in the simulation for different trigger conditions applied to the BASIC and COMBINED robustness strategies. Simulated costs of solutions are expressed as % of the BASIC planned costs. All trigger conditions are simulated five times.

	BASIC				COMBINED			
	SC(%)	D	RP	VS	SC(%)	D	RP	VS
1-day	116.35	134	27	602	111.45	55	14	379
3-day	116.11	133	11	418	111.96	60	5	293
5-day	116.83	143	5	270	112.61	66	2	266

are allowed (see Table 8). Cost savings are also obtained for the COMBINED solution. We also see that the lower threshold trigger used, the more often the recovery procedure is called, as could be expected. When simulating COMBINED solutions with a 5-day voyage delay trigger, the procedure is on average only called twice. The re-planning recovery procedure is further called approximately twice as often in the BASIC simulations than in the COMBINED simulations. This illustrates that the initial COMBINED solutions are considerably more robust to disruptions and are able to absorb most of the impacts of the incurred disruptions with their incorporated slack.

Next, we observe how little the incurred total delay varies between the different trigger strategies. For the BASIC simulations, we observe that a 3-day trigger actually seems to perform slightly better than calling the procedure every time a voyage is anticipated to experience only one day of delay or more, both in terms of delays and costs. However, the results do not show the same for COMBINED solutions. We believe that this difference may be due to the recovery procedure for BASIC solutions which does not have the same robustness incentives as when called for COMBINED solutions.

By considering the aforementioned observations, we see that including the re-planning procedure has the potential of better mitigating the impacts of disruptions than by not considering recovery actions. Despite slightly lower simulated costs, easily triggering the re-planning procedure leads to many alterations to the operations. Even though voyage swaps may only be planned changes and do not actually occur due to additional events, these are still considered as changes and must be facilitated by the case company. The results suggest that if a price for performing changes was considered, triggering the recovery procedure less often may be desirable. Different companies have their own individual preferences in the trade-off between the extent of changes, expected costs and expected delays, and therefore we do not give a final suggestion of the optimal trigger strategy.

6.3.3. Combining robustness and recovery strategies

Table 10 shows the average costs of the solutions from ten simulations of each robustness strategy combined with the re-planning procedure as a percentage of the BASIC planned operating costs. For the purpose of comparing the performance of combined robustness strategies and re-planning recovery, the recovery procedure is triggered when the expected delay of any voyage

exceeds three days. We also include the results from Table 8 for easier comparison.

Table 10: Average simulated costs (SC) and total days of delay (D) for the different robustness strategies combined with re-planning recovery. The simulated costs are expressed as % of the planned costs for the BASIC solution. Each combination of robustness strategy and re-planning recovery action is simulated 10 times.

	BASIC		INCREASE		REWARD		PENALIZE		COMBINED	
	SC(%)	D	SC(%)	D	SC(%)	D	SC(%)	D	SC(%)	D
<i>No re-planning</i>	155.02	604	133.72	329	117.06	113	143.38	454	116.15	118
<i>Speed re-planning</i>	121.50	197	118.93	142	114.86	82	115.40	112	113.41	77
<i>Full re-planning</i>	117.31	143	116.55	116	114.48	81	113.10	84	112.83	67

The results in Table 10 show that by enabling re-planning in the simulations, denoted *Full re-planning*, further cost and delay reductions are achieved for all robustness strategies. As for the *No re-planning* and the *Speed re-planning* settings, the BASIC solution has the poorest performance both in terms of delay and simulated costs. Despite the relatively high simulated costs of the BASIC solution with *Full re-planning*, it is evident that including re-planning can result in considerable savings. Of the robustness strategies, the solution based on the INCREASE strategy has the largest simulated costs and highest amount of delay, while the COMBINED strategy performs the best, reducing the costs and delays compared with the *Speed re-planning* setting by another 0.4 % and 13 %, respectively.

A scatter plot of the average simulated costs and the average total delay for solutions based on the robustness strategies combined with the re-planning procedure is presented in Figure 8. In addition, a 95% confidence interval is given for both costs and delays. The figure shows that even though the COMBINED strategy gives the lowest average both in terms of costs and delays, we cannot conclude (with 95 % certainty) that it is better than the REWARD or PENALIZE strategies. Further, the plot shows that, on average, the INCREASE and BASIC strategies perform the worst, and is dominated by the COMBINED strategy.

Table 11: Average number of re-planning recovery procedure calls and voyage swaps when simulating solutions generated with the different robustness strategies with the re-planning recovery procedure enabled (*Full re-planning* setting).

	BASIC	INCREASE	REWARD	PENALIZE	COMBINED
Re-planning procedure calls	12	11	9	8	7
Voyage swaps	409	300	283	196	260

Whenever the re-planning procedure is used, changes to the previous plan may occur. Each robustness strategy affects the number of times the re-planning procedure is called during the simulation and how many changes these calls make to the existing plan in a different way. Table 11 shows the average number times the re-planning procedure is called on each robust solution, and the resulting number of total voyage swaps from these new plans. The plans based on the robustness strategies BASIC and INCREASE, that gives the highest amount of delay with the *Speed re-planning* setting, require most re-planning to be performed during the simulation. The robustness strategies with a lower total delay will have fewer voyages with delays of more than three days, and




Figure 8: Simulated costs and delays for each strategy with *Full re-planning*, given both with the average value and a 95 % confidence interval.

thus trigger fewer calls to the re-planning procedure. However, we see that the simulations of the PENALIZE solutions only trigger the re-planning procedure eight times compared to nine for the REWARD solution, despite having a 112 days of delay in the *Speed re-planning* setting compared to REWARD’s 82 days of delay. This again implies that using the PENALIZE strategy results in plans where the disruptions will cause many smaller delays, while the REWARD strategy gives vessel schedules where fewer voyages are delayed, but each delay is of a greater magnitude.

Further, we observe from Table 11 that solutions generated with the PENALIZE robustness strategy are, on average, subject to the least amount of voyage swaps. This can also be explained by how the PENALIZE robustness strategy evenly distributes slack to vessels. If the COMBINED strategy had used the same parametric values as each of the individual robustness strategies, then it would have had both less re-planning procedure calls and fewer voyage swaps. However, as seen in Tables 5 and 6, the COMBINED strategy uses smaller values for each parameter and it is therefore not surprising that the number of voyage swaps for the COMBINED strategy lies somewhere in between the other strategies. However, in terms of re-planning procedure calls, the COMBINED strategy has the fewest.

By comparing PENALIZE and COMBINED, we can see that by selecting PENALIZE and accepting operating costs that are, on average only slightly higher, considerably fewer changes are needed to be done to the deployment plan during operation. Further, if we compare the

COMBINED solutions with *Full re-planning* and *Speed re-planning* it can be seen that zero changes to the plan can be achieved by accepting, on average, only 0.4 % higher costs and 13 % more delays.

7. Concluding remarks

We have presented a set of strategies to include robustness, and handle disruptions, in the fleet deployment planning of a roll-on roll-off liner shipping company. A new mathematical model of the problem is proposed, together a set of robust planning strategies that ensure that the resulting plan is less vulnerable to disruptions. To solve real-life instances of the problem a rolling horizon heuristic is presented. A computational study is conducted to test the heuristic and the effect of the robustness strategies. To test the effect of the proposed robustness strategies we simulate the execution of a fleet deployment plan while introducing disruptions. The computational results show that none of the robustness strategies individually adds more than a modest 5 % to the planned costs of the fleet deployment. However, when the execution of the plans were simulated, the results show that when combining robustness strategies with speed adjustments, the costs are substantially lower, sometimes by as much as 7 %, and the total days of delay are reduced by up to 61 %. When also including the re-planning procedure, the impact of disruptions can be further reduced by approximately 0.4 % in costs and 13 % in total delays.

An anonymous referee noted that this paper only discusses reactive re-planning, i.e. re-planning is only performed if a disruption occurs. An extension of this could be to also include proactive re-planning when a disruption does not occur. Suppose for example that the travel time through a canal is highly uncertain and that to hedge against a long delay, extra travel time is added. Now, suppose that no disruption occurs and the vessel, after exiting the canal, is scheduled to arrive several days ahead of schedule. Then, using proactive re-planning one would also call the re-planning in this case, which would probably result in that the ship would reduce its sailing speed on the remaining travel in order to reduce costs. It should be emphasized that this could relatively easily be incorporated in our solution methodology.

Acknowledgements

We would like to thank the reviewers for their contributions to the paper, and the GreenShipRisk project funded by the Norwegian Research Council.

Bibliography

- [1] Alvarez, J. F., Tsilingiris, P., Engebretsen, E. S., Kakalis, N. M. P., 2011. Robust fleet sizing and deployment for industrial and independent bulk ocean shipping companies. *INFOR: Information Systems and Operational Research* 49 (2), 93–107.

- [2] Andersson, H., Fagerholt, K., Hobbesland, K., 2014. Integrated maritime fleet deployment and speed optimization: Case study from ro-ro shipping. *Computers & Operations Research* 55, 233–240.
- [3] Brouer, B. D., Dirksen, J., Pisinger, D., Plum, C. E. M., Vaaben, B., 2013. The vessel schedule recovery problem (VSRP) – a MIP model for handling disruptions in liner shipping. *European Journal of Operational Research* 224 (2), 362 – 374.
- [4] Christiansen, M., Fagerholt, K., 2002. Robust ship scheduling with multiple time windows. *Naval Research Logistics* 49 (6), 611–625.
- [5] Fagerholt, K., Johnsen, T. A., Lindstad, H., 2009. Fleet deployment in liner shipping: A case study. *Maritime Policy & Management* 36 (5), 397–409.
- [6] Gelareh, S., Meng, Q., 2010. A novel modeling approach for the fleet deployment problem within a short-term planning horizon. *Transportation Research Part E: Logistics and Transportation Review* 46 (1), 76–89.
- [7] Gurobi Optimization, I., 2014. Gurobi optimizer reference manual.
URL <http://www.gurobi.com>
- [8] Halvorsen-Weare, E. E., Fagerholt, K., 2011. Robust supply vessel planning. In: *Network Optimization*. Springer, pp. 559–573.
- [9] Halvorsen-Weare, E. E., Fagerholt, K., Nonås, L. M., Asbjørnslett, B. E., 2012. Optimal fleet composition and periodic routing of offshore supply vessels. *European Journal of Operational Research* 223 (2), 508–517.
- [10] Halvorsen-Weare, E. E., Fagerholt, K., Rönnqvist, M., 2013. Vessel routing and scheduling under uncertainty in the liquefied natural gas business. *Computers & Industrial Engineering* 64 (1), 290–301.
- [11] Kjeldsen, K. H., 2012. Routing and scheduling in liner shipping. Ph.D. thesis, Aarhus University, 166 pages.
- [12] Liu, X., Ye, H.-Q., Yuan, X.-M., 2011. Tactical planning models for managing container flow and ship deployment. *Maritime Policy & Management* 38 (5), 487–508.
- [13] Meng, Q., Wang, T., 2010. A chance constrained programming model for short-term liner ship fleet planning problems. *Maritime Policy & Management* 37 (4), 329–346.
- [14] Mercao, C., Fontan, G., 2003. MIP-based heuristics for capacitated lotsizing problems. *International Journal of Production Economics* 85 (1), 97–111.

- [15] Notteboom, T. E., 2006. The time factor in liner shipping services. *Maritime Economics & Logistics* 8 (1), 19–39.
- [16] Perakis, A. N., Jaramillo, D. I., 1991. Fleet deployment optimization for liner shipping part 1. Background, problem formulation and solution approaches. *Maritime Policy & Management* 18 (3), 183–200.
- [17] Qi, X., 2015. Disruption management for liner shipping. In: Lee, C.-Y., Meng, Q. (Eds.), *Handbook of Ocean Container Transport Logistics*. Vol. 220 of *International Series in Operations Research & Management Science*. Springer International Publishing, pp. 231–249.
- [18] Rakke, J. G., Stålhane, M., Moe, C. R., Christiansen, M., Andersson, H., Fagerholt, K., Norstad, I., 2011. A rolling horizon heuristic for creating a liquefied natural gas annual delivery program. *Transportation Research Part C: Emerging Technologies* 19 (5), 896–911.
- [19] Savelsbergh, M., 1985. Local search in routing problems with time windows. *Annals of Operations Research* 4 (1), 285–305.
- [20] Toth, P., Vigo, D., 2014. *Vehicle Routing: Problems, Methods, and Applications*. Vol. 18. SIAM.
- [21] Wang, S., Meng, Q., 2012. Liner ship fleet deployment with container transshipment operations. *Transportation Research Part E: Logistics and Transportation Review* 48 (2), 470–484.
- [22] Wang, T., Meng, Q., Wang, S., Tan, Z., 2013. Risk management in liner ship fleet deployment: A joint chance constrained programming model. *Transportation Research Part E: Logistics and Transportation Review* 60, 1–12.
- [23] WWLBreakbulk, 2010. Accessed: 2015-11-30.
URL <http://www.2wglobal.com/globalassets/folders/mark-v-folder-lowres-pdf.pdf>
- [24] Yu, G., Qi, X., 2004. *Disruption Management: Framework, models and applications*. World Scientific Publishing Company Incorporated, Singapore.
- [25] Zeng, Q., Yang, Z., Chen, C., 2010. Robust optimization model for resource allocation of container shipping lines. *Tsinghua Science & Technology* 15 (5), 586–594.