



Liner shipping single service design problem with arrival time service levels

Kevin Tierney¹ · Jan Fabian Ehmke² · Ann Melissa Campbell³ · Daniel Müller⁴

Published online: 27 August 2018

© Springer Science+Business Media, LLC, part of Springer Nature 2018

Abstract

We introduce three mathematical models of increasing complexity for designing liner shipping services that guarantee the punctual arrival of vessels at a specified service level. On-time reliability is an important performance indicator for many liner carriers, but current approaches for creating new routes in liner shipping networks do not consider data-driven uncertainty. We perform an empirical analysis of vessel travel times in a real liner shipping network to develop probability distributions that we use within novel, chance-constrained mathematical models for liner shipping service design. Our models are also the first to support variable vessel speeds for service design. In our experiments, we use real-world data from 22 liner shipping routes and evaluate the designed services using a simulation procedure that demonstrates the effectiveness of our approach for reducing lateness. We show that our models can be effectively used for decision support at a tactical level not only for designing services, but also potentially for negotiating maximum demand transit times and prices with customers.

Keywords Maritime transportation · Liner shipping optimization · Service design · On-time guarantee

Electronic supplementary material The online version of this article (<https://doi.org/10.1007/s10696-018-9325-y>) contains supplementary material, which is available to authorized users.

✉ Kevin Tierney
kevin.tierney@uni-bielefeld.de

¹ Decision and Operation Technologies Group, Bielefeld University, Bielefeld, Germany

² Management Science Group, Otto von Guericke University Magdeburg, Magdeburg, Germany

³ Department of Management Sciences, Tippie College of Business, University of Iowa, Iowa City 52242, USA

⁴ Decision Support and Operations Research Lab, University of Paderborn, Paderborn, Germany

1 Introduction

The liner shipping industry is situated at the center of global trade, providing efficient and secure freight transportation on over 5000 seagoing vessels (United Nations Conference on Trade and Development (UNCTAD) 2015). Each year, more and more freight is transported with standardized steel containers in liner shipping networks. In 2015, over 171 million TEU¹ of freight was carried across the oceans, representing over 1.6 billion tons of goods (United Nations Conference on Trade and Development (UNCTAD) 2015).

Liner shipping differentiates itself from other forms of maritime transportation, such as tramp or industrial shipping, due to the periodic and cyclical nature of the routes in liner shipping networks. In liner shipping, these routes, called *services*, visit a sequence of ports within specified time windows on a periodic (usually weekly) basis. Once a vessel reaches the last port in the sequence, it travels to the first port in the sequence and starts again. As it may take more than a week to finish a round-trip, it is often necessary to deploy more than one vessel to fulfill the periodicity at the ports. The periodicity and reliability of liner shipping services have become key selling points, resulting in liner shipping services forming the backbone of many modern supply chains (Notteboom and Rodrigue 2008). However, the planning of efficient and reliable liner shipping services is challenging, as there are many sources of delay that can send a vessel off-schedule, which can have expensive repercussions throughout the supply chain (Notteboom 2006). Sources of delays for vessels include, for example, bad weather, port congestion, equipment breakdowns, labor disputes and medical emergencies. These can cause a vessel to be late for its weekly time window or even have to cancel its stop at a particular port. Moreover, it is not only important that vessels arrive at ports when scheduled, but that the shipped containers (*demands*) reach their destination in the time period that was guaranteed by the liner carriers, as arrival uncertainty causes extra costs in the supply chain (Vernimmen et al. 2007).

Due to the periodicity and cyclical nature of liner shipping services, there are interdependencies between the expected travel times of a service, the number of vessels assigned to a service, and the resulting reliability of a service. When designing services for their networks, liner carriers often add buffer time before each port call to reduce the effect of such delays and increase the reliability of schedules. The amount of buffer is usually determined based on the experience of planners or simple rules of thumb. Liner carriers often have no analytical basis in the historical data for each port for decision making, nor is the order of the service optimized to provide a certain amount of reliability. Analytical tools are missing that allow for an exploration of costs and reliability of liner shipping services considering the specific structure of the schedules and the complex trade-offs.

This paper focuses on the planning of a single service within a liner shipping network that explicitly incorporates uncertainty in travel times. The ports that should be

¹ A single twenty foot equivalent unit (TEU) represents one twenty foot container, with two TEU representing the commonly found forty foot container.

visited and the berthing windows identifying when they should be visited are given as inputs for all models. In addition to optimizing the route a service takes to ports, we also consider the number of vessels necessary for running the service in conjunction with the planned speed of the vessels. Our models provide an arrival time service level to planners in a similar fashion to the service levels in Ehmke et al. (2015) for the well known vehicle routing problem with time windows, with one key difference. In Ehmke et al. (2015), lateness accumulates over the course of each vehicle's tour. In the case of liner shipping service design, during the execution of a route, planners have a range of actions they can take to get a vessel back on schedule (see Brouer et al. 2013b). We do not explicitly model recourse actions in our tactical model, but assume that they can be taken to avoid propagation of lateness. However, we analyze the impact that propagation can have on the reliability of a service through simulation.

We present three mathematical models of increasing complexity for planning liner shipping services to understand how different features influence the results. To establish these models, we begin with a data-driven investigation into the distribution of travel times between different kinds of ports in liner shipping (see Sect. 3). This provides an analytical background for how we model uncertainty in liner shipping. In Sect. 4, we present the mathematical models and how they are able to accommodate the derived travel time distributions. In Sect. 4.1, we start with analyzing the impact of different service level guarantees on the arrival time at different ports on the design of cost-minimizing services, assuming the vessels plan to travel at their design speed. This helps us quantify the impact of modeling uncertainty. In Sect. 4.2, we expand the first model to explicitly consider the impact of varying speed levels on the number of required vessels and the resulting arrival time service levels. Contrasting the model of Sect. 4.1, we can now trade a smaller number of vessels for higher speeds in service execution or vice versa. In order to provide acceptable service levels with regard to port time windows, our first two models often create large buffers in the services' schedules. Thus, they ignore delivery time guarantees on freight. Our third model, in Sect. 4.3, alleviates these limitations and helps us understand how different levels of delivery time guarantees can impact the amount of freight that can be carried when there are limitations on maximum container transit times.

In Sect. 5, we experiment with the use of the presented mathematical models in a series of computational experiments. To this end, we apply the presented travel time distributions to instances based on data from the well-known LINERLIB (Brouer et al. 2013a) and compute the optimal solutions for the different models. Considering different service levels through chance constraints in our optimization models, we analyze the trade-off between a larger fleet of vessels and adaptations of speed levels from a tactical network design perspective. Then, we assess the operational performance of these optimized services with a discrete-event simulation. The simulation evaluates the optimized services by imitating individual service runs to determine the "actual" realized cost and quality of a service, analysing the impact of propagation of lateness in the course of a round-trip. We provide a set of managerial insights based on these results in Sect. 6.

In summary, our paper provides the following contributions:

1. a data-driven investigation of distributions of travel time for liner shipping services,
2. the modeling of service guarantees on arrival times through chance constraints for service design,
3. a model for service routing optimization with variable vessel speed (tactical perspective),
4. and a simulation study to evaluate the effectiveness of the models (operational perspective) and the impact of lateness propagation when executing the services.

In the following section, we delineate our modeling approach from related literature in this field.

2 Literature review

The liner shipping service design problem with arrival time service levels has connections to a number of different problem domains, including the areas of vehicle routing, tramp shipping, and liner shipping. We now discuss the similarities and differences of this work with related work in the literature, referring to Christiansen et al. (2004, 2013) for an overview of the entire area of maritime transportation and Brouer et al. (2017) for an overview of liner shipping optimization. A bibliometric overview is given in Lau et al. (2017).

2.1 Liner shipping

We divide our discussion of related work within the area of liner shipping into two subsections. First, we discuss related problems such as network design and fleet deployment and then move into work addressing single service scheduling and routing. Given the prevalence of delays in maritime applications due to storms, labor disputes, and breakdowns, considering uncertainty is a natural extension to operations research models for tramp, industrial and liner shipping. Nonetheless, there is little literature considering stochastic elements in these problems.

2.1.1 Related liner shipping problems

Liner shipping single service design can be considered a subproblem of the overall network design problem and is considered a tactical problem in Meng et al. (2013). The network design problem is very difficult to solve with exact solution approaches (Agarwal and Ergun 2008; Álvarez 2009; Brouer et al. 2013a) and even for heuristics (Brouer et al. 2014). Due to its difficulty, network design problems usually leave out or abstract a number of important side constraints that we are able to include, such as the handling of container transit times. We note that to the best of our knowledge, there has not been any work on combining uncertainty with liner shipping network design models.

Fleet deployment models are also related (see, e.g., Powell and Perakis 1997), as these assign a heterogeneous fleet of vessels to a set of services. In contrast to these models, we do not try to size the vessel to the service we are designing. In our approach, the vessel class is an input parameter. Fleet deployment models are subject to different types of uncertainty than service/network design problems. A key source of uncertainty for these types of problems is the amount of demand that is shipped each week on each service. A chance-constrained model takes this into consideration in Meng and Wang (2010), finding a minimal cost allocation of vessels to services. We do not take demand uncertainty into account in our model, but this would be a logical extension of our work. Transit time restrictions for demands were recently considered in terms of the cargo allocation problem, which assigns containers to routes in a network in Guericke and Tierney (2015), as well as for a time-constrained multicommodity flow problem in Karsten et al. (2015). We solve a simplified version of this problem as a part of our single service design. We do not consider transshipments, making our cargo allocation easier in comparison.

2.1.2 Single service scheduling

Several papers determine vessel speeds and schedules for one (or more) service(s) under uncertainty given a pre-defined port sequence. The key difference between these works and our model is that we determine both the port sequence and the schedule. Delays at ports are considered in Qi and Song (2012), in which delayed arrivals are penalized in the objective function, reflecting a potential loss of goodwill. The authors discuss how to compute the service level at each customer, acknowledging that prior delays can accumulate. Furthermore, the paper focuses on special cases, such as with 100% service levels for all customers. A drawback of this work is that the authors do not use real data for their port time distributions, instead assuming uniform and normal distributions. In a parallel work, Wang and Meng (2012a), the authors also seek vessel speeds and a schedule under uncertainty for a new liner shipping service. They create a mixed-integer non-linear stochastic programming model for the problem that minimizes ship and fuel costs while maintaining a given service level. The problem is solved using a cutting-plane algorithm.

The most relevant work in the area of liner shipping in terms of focusing on delays is Lee et al. (2015). The paper assumes the routes are given and looks at the impact of different steaming speeds for executing the route, given the stochastic nature of port operations. Buffers are planned into liner shipping schedules in order to maintain a specified arrival service level. In contrast, not only do we perform route planning, we also do not limit ourselves to only considering port delay. While port delays make up a large percentage of the sources of delay (Notteboom 2006), a number of other sources exist that we are able to account for in our model.

A feeder network design problem is presented in Santini et al. (2017) that routes and schedules several services within a geographically limited area. The authors use an expanded time space graph so that vessel speeds can be precomputed and assigned as costs directly to the arcs. The number of vessels used on the routes is fixed to a value K , meaning the set of speeds for vessels and amount of buffer that

Table 1 A categorization of related work within the area of liner shipping service/schedule design

Article	Routing	Lateness dist.	Speed opt.	Method
Wang and Meng (2012a)	✗	Any truncated	Non-linear	Non-lin. stoch. prog.
Wang and Meng (2012b)	✗	Uniform & normal	Non-linear	Non-lin. stoch. prog.
Qi and Song (2012)	✗	Uniform/normal	Non-linear	Sim. stoch. approx.
Song and Dong (2013)	✓	✗	Non-linear	Heuristic decomposition
Plum et al. (2014)	✓	✗	✗	Branch-cut-and-price
Lee et al. (2015)	✗	Normal/any	Non-linear	Markov chains
Song et al. (2015)	✗	Trunc. normal	Non-linear	NSGA-II (Deb et al. 2002)
Reinhardt et al. (2016)	✗	✗	Disc. secants	MILP
Wang and Wang (2016)	✗	✗	Non-linear	Polynomial time algorithm
Santini et al. (2017)	✓	✗	Disc. graph	Branch and price
This paper	✓	Log-logistic 3P	Disc. secants	MILP

can be inserted into the schedule, is limited. Thus, Santini et al. (2017) does not consider arrival time service levels.

The paper that motivates the basic form of our model is Plum et al. (2014). The goal in this work is to design a single service given a set of ports that must be visited. A set of container demands between ports must be satisfied while taking into account the capacity restrictions of the vessels. Our first deterministic model (Phase 1) resembles this work very closely, adding only the optimization of the number of ships on the service. Furthermore, our models extend the work of Plum et al. (2014) with variable vessel speed and arrival time guarantees, which represent important characteristics for the liner shipping industry. A summary of the features of our models in relation to the literature is given in Table 1.

2.2 Handling delays in related routing problems

There are several related problems in which delays have been considered. Ehmke et al. (2015) is among the most recent; they explicitly consider the interplay of stochastic travel times and time windows in the vehicle routing problem, ensuring service levels in routing through chance constraints. For them, the largest challenge stems from the computation of the combined arrival time distribution along a route, which we ignore in our approach since we assume statistical independence between the individual sequences on a route.

Tramp and industrial shipping involve the transportation of bulk or liquefied goods (but rarely containers) in a vehicle routing-like fashion. Tramp/industrial shipping problems differ greatly from liner shipping problems in terms of the schedule structure. Whereas liner shipping has a periodic, fixed schedule like a public bus network, tramp/industrial shipping more resembles taxis, in which ships sail wherever there is a demand to be satisfied. However, despite their differences, all of these types of shipping involve vessels that can be delayed in the same way, as well as have similar cost profiles for sailing. Speed optimization has become a standard feature of tramp shipping models (Norstad et al. 2011).

The most recent work in this area is Agra et al. (2015), which considers a maritime inventory routing problem for liquefied natural gas with stochastic sailing and port times. The authors determine routes and the quantity of gas to load/unload a priori and model the problem as a stochastic program with recourse using scenarios. In contrast, we compute distributions over the arrival times of vessels at ports. In Agra et al. (2016), this work is extended to use a log-logistic distribution to model delays as in our work. Uncertain sailing times are considered in Halvorsen-Weare and Fagerholt (2013), motivated by a real energy company, as well as uncertain production rates for a natural gas provider. In Halvorsen-Weare et al. (2013), the same problem is solved as in Halvorsen-Weare and Fagerholt (2013), but with robustness strategies added. The sailing times are fitted to a log-logistic probability distribution, based on information gathered for a gas tanker in Kauczynski (1994).

A method for creating replenishment schedules for offshore installations is proposed in Halvorsen-Weare and Fagerholt (2011). The paper defines four weather states and their impact on sailing speed and service time in ports at an offshore delivery location. A key difference with our work is that the authors assume a set of routes already exist, and they must choose which routes to use.

3 Data exploration

A key challenge that arises in guaranteeing arrival time service levels for a liner shipping service is determining what distribution underlies the travel time of the vessels. As discussed in the previous section, the maritime literature has several suggestions on appropriate travel time distributions, including the log-logistic distribution (Halvorsen-Weare et al. 2013). However, it is not clear a priori that this distribution is the best fit for travel times of liner shipping services, and furthermore, it is not clear that this distribution is appropriate for trips between all ports worldwide. To this end, we gathered data from operating liner shipping service routes, including profiles of the vessels used and the actual transit times of the vessels between ports. We first describe the data we use for distribution fitting in more detail, including a frank discussion of the strengths and weaknesses of our dataset, and finish with a presentation of the distributions we found.

3.1 Data sources and integration

We have taken care to make our dataset as realistic as possible. To do this, we use service information from the liner carrier COSCO and combine this with AIS (positioning) data from MarineTraffic regarding the vessels on 25 selected liner shipping services. We note that we have no relationship with COSCO and use data regarding their network because they offer detailed arrival time information for their services. We gathered positioning data from the vessels' transponders for the year 2014 and use this data to check whether the vessels are on time or not by comparing it with the planned schedule from COSCO (COSCO Shipping Lines 2017). In total, our dataset contains 40 ports, 118 vessels, 125 port to port connections, and records for

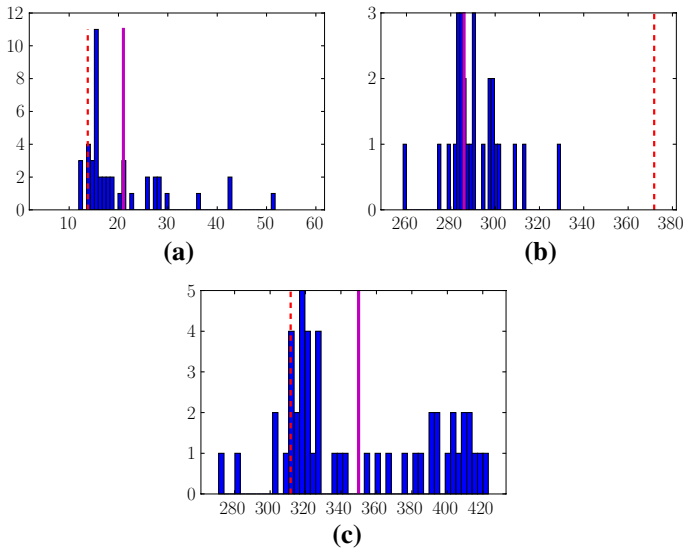


Fig. 1 Histograms of travel times (in hours) of several port to port pairs. The x-axis shows the varying travel times in hours between the ports. The y-axis gives the frequency of a particular travel time. The scheduled transit time is shown with a dashed red line and the travel time at the vessel's design speed with a solid magenta line. **a** Rotterdam, NL to Le Havre, FR, **b** Long Beach, US to Ningbo, CN, **c** Singapore, SG to Suez, EG

1872 transits between ports. This provides us with a list of transits between the ports visited by the 25 services and an indication of how late (or early) the vessel was as compared to the planned travel time.

Figure 1 shows three sample transits and their associated histograms, chosen because they represent a range of connection types: an intra-region connection (a), a Pacific Ocean connection (b) and an Indian Ocean connection (c). While it is not possible to generalize from the data of only a few connections, it is clear that even on short connections, such as crossing the English channel, significant delays above the scheduled sailing time are possible. Furthermore, we can see that some schedules are planned with enough buffer that even delays of a couple of days do not cause lateness, as in the case of the backhaul from Long Beach to Ningbo. Finally, some schedules are tightly planned, as in the case of sailing from Singapore to the southern side of the Suez canal. Here, significant delays of up to almost five days can occur. We note that it is not unusual for connections to be assigned a speed faster than the vessel's design speed as in Fig. 1a, c. These are likely connections where shippers wish to have fast transit times and the carrier must meet these requests to carry shipper's cargo.

The data we use from the AIS transponders is the best data we can obtain, as carriers we have spoken to do not keep track of more accurate statistics regarding travel time. There are, however, weaknesses that we want to outline. An obvious issue is that we do not know what recourse actions were taken to avoid or reduce lateness during operations. In particular, if a vessel is running late, a carrier may instruct

the captain to speed up to stay on schedule. We note that AIS data does include the speed of the vessel, but even with the speed, we cannot assess whether a speed-up was part of the original service schedule or not. Another source of error in our data is that extremely late vessels may simply skip port calls to save time, meaning we do not end up knowing how late they actually were. Furthermore, if no vessel performs a particular port to port transit, we will have no data for it, and this makes it difficult to know what distribution to use to model travel time for such a connection when considering the creation of a new service. We attempt to counteract some of the weaknesses in our approach by clustering the port to port connections and by computing distributions between regions rather than individual ports. This alleviates the problem for many connections.

3.2 Distribution fitting

We compute distributions based on intra and inter-region transit times, normalized according to the average travel time for each region-to-region pair. To create distributions for any pair of ports worldwide, we aggregate and generalize the operational data as follows. First, to generalize the port-to-port observations, we follow a common idea in the network design literature (e.g. Mulder and Dekker 2014) and cluster ports into 16 regions with the well known k -means algorithm, adjusting the ports between some regions slightly by hand. The resulting clusters provide the input for our distribution fitting. Second, between two regions or within a single region, for each pair of ports, we normalize the travel times based on the mean empirical travel time between (or within) the regions containing the ports. That is, given ports i and j from regions r and s , we compute the average travel time between the regions and divide the travel time between i and j by this value. Since not every pair of regions has sufficient data to allow for an empirical fit of a distribution, we also compute a distribution across all data that can be used for such port pairs.

We use the software EasyFit (MathWave Technologies 2016) to generate a list of distributions for each inter and intra region pair to narrow down the type of distribution to use across our data. In 17 out of 28 pairs with enough port to port transfers to fit a distribution, the three-parameter log-logistic (ll3p) distribution is one of the top three best fitting distributions, and in all other cases was still one of the best fitting distributions. Other good fits included the Cauchy distribution and the four-parameter Burr distribution. Given the previous use of the ll3p distribution in the literature, we select it for the remainder of this work.

Figure 2 shows histograms and the best fitting normal and log-logistic probability distribution functions for several region-to-region pairs. The x -axis provides the normalized travel time, with the y -axis reflecting the frequency of a particular value. There is a clear trend across all of the regions we consider. The normal distribution has less probability mass than the ll3p distribution around on-time transits at time 0, which is a good indicator of why the log-logistic distribution is often the one with the best fit. In many cases, the log-logistic distribution puts more mass near on-time arrivals and is skewed to the right, reflecting a large number of vessels arrive late and across a wide range of values. For a few region-to-region pairs, the log-logistic

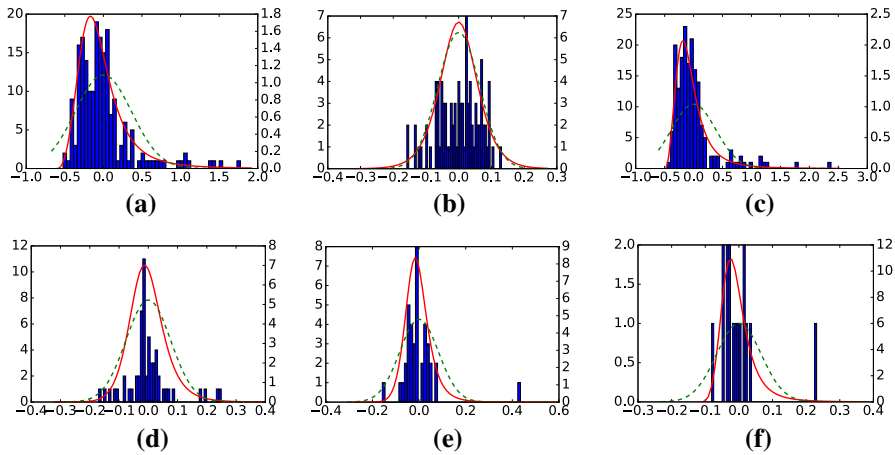


Fig. 2 Histogram and probability distribution functions for several region-to-region pairs. Blue bars show the histogram, a solid red line shows the fit of the lI3p distribution, and the dashed green line the fit of a normal distribution. **a** Intra Europe, **b** China/Japan to Southern Mediterranean, **c** Intra Western North America, **d** Western North America to China/Japan, **e** Central/South America to North America, **f** Indian Ocean to Southern Mediterranean

distribution looks almost like the normal distribution, as is the case in Fig. 2b, and this further emphasizes the importance of fitting specific parameters for different parts of the world, instead of only one set of parameters.

4 Mathematical models

We construct an arc flow model that extends the one presented in Plum et al. (2014) to find a cyclical route through ports given fixed port time windows. We decompose the modeling of the extended problem into three phases, each one with added complexity. This helps us understand the impact of different problem features and also their impact on the solution time. Unlike in Plum et al. (2014), all of our models contain chance constraints requiring vessels arrive at ports on-time with a specified arrival time service level. In the first phase, we address the basic liner shipping service design problem, requiring all demand to be transported without the maximum transit times on vessels sailing at their design speeds used in Plum et al. (2014). The design speed is used to provide schedules that can serve as a reasonable baseline for Phase 2 and Phase 3. In the second phase, we relax the fixed sailing speed requirement and allow the vessel to vary its speed, taking this into account in the objective function. In the third phase, we impose maximum container transit times, but allow demand to be rejected if it is not possible to meet the transit time limitations. The objective function in this phase switches to profit maximization from cost minimization in Phases 1 and 2. As mentioned in Sect. 1, our models do not consider propagation of lateness over the course of the service. This is because route planners have a range of actions they can take to get a vessel back on schedule (see Brouer et al.

2013b) if needed, including skipping ports, negotiating new time windows, etc. We do not explicitly model these recourse actions in our tactical model, because we do not know which ones could be negotiated in each case. We assume, though, that such actions can be taken to avoid propagation of lateness, but also create expenses that a company would want to avoid if possible through on-time arrivals.

4.1 Phase 1: design speed model

In Phase 1, we impose the following assumptions:

1. All ports are assigned a fixed port time window per week in which they are called. The vessel must arrive before the port time window starts (planned arrival) and may not leave until the window ends (planned departure).
2. Vessels must wait until the start of their time window if they arrive early at a port.
3. Time windows must be satisfied with a confidence level of α (arrival time service level). For example, α may be a value such as 75% or 90%.
4. Travel times between ports are statistically independent.
5. All vessels must plan on traveling at their design speed between ports.
6. All pickups and deliveries of demand must be served, assuming that the vessel capacities are sufficiently large.
7. Forty foot containers are broken down into two twenty foot containers, and their flow can be modeled in a continuous fashion.
8. All ports must be called exactly once.
9. The objective is to minimize the cost of the single route and the number of vessels required.
10. The service frequency is weekly.

Note that the requirement that all ports must be called exactly once means we cannot design butterfly or conveyor belt style routes. We note our model can be adjusted to create such routes and in its current form can support having a port specified multiple times. According to Song and Dong (2013), cycle routes (i.e., those where all ports are visited once) are the most common topological structure and can be found in about 45% of services. We assume that one of the ports is arbitrarily identified as the start port for the service, and it is also the end port since the route is cyclical. The starting port is assumed, without loss of generality, to be a port that has a time window that is fully contained during the week, i.e., the time window does not intersect with Monday at 12:00 AM. We further note that delays have some correlation with the amount of loaded cargo (Hasheminia and Jiang 2017) and we do not take this into account in this model.

We now define the mixed-integer linear program for the Phase 1 model.

Parameters

$P = \{1, \dots, p, p + 1\}$ Set of ports to call; ports 1 and $p + 1$ represent the initial port (and its return)

A	Set of arcs (i, j) ; we assume that the network is complete, with no arcs into the initial port or out of the ending port
A'	All of the arcs of A with the addition of a single arc connecting the ending port to the initial port
K	Set of demands
o_k	Origin of demand $k \in K$
d_k	Destination of demand $k \in K$
a_k	Amount of cargo of demand $k \in K$
u	Capacity of a vessel in TEU
t_i^s, t_i^e	Time window start and end (respectively) within a week for port $i \in P$, where time 0 is Monday at 12:00 AM. The values for the time window start and end time represent hours. The latest arrival time has a domain $t_i^s \in \{0, \dots, 167\}$, whereas the earliest departure time $t_i^e \in \{t_i^s, \dots, 335\}$. This allows port time windows to start and end in consecutive weeks.
c_{ij}^f	Fixed sailing cost for arc $(i, j) \in A$
t_{ij}^α	Time for sailing arc $(i, j) \in A$ at the vessel's design speed to provide a service level of α
c^v	Charter cost per vessel per week

Variables

w_i	Service week at port $i \in P$, where $w_1 = 0$
x_{ij}	Indicates whether arc $(i, j) \in A$ is included in the service
τ_i^s	Start service time of a vessel at port $i \in P$ (hours from beginning of service)
τ_i^e	End service time of a vessel at port $i \in P$ (hours from beginning of service)
τ_i^a	Arrival time at port $i \in P$ (hours from beginning of service)
f_{kij}	Flow of demand $k \in K$ on arc $(i, j) \in A$

Objective and constraints

$$\min c^v w_{p+1} + \sum_{(i,j) \in A} c_{ij}^f x_{ij} \tag{1}$$

$$\text{subject to } \sum_{(i,j) \in A} x_{ij} = 1 \quad \forall j \in P \setminus \{1\} \tag{2}$$

$$\sum_{(i,j) \in A} x_{ij} = 1 \quad \forall i \in P \setminus \{p+1\} \tag{3}$$

$$\tau_i^s = t_i^s + 168w_i \quad \forall i \in P \quad (4)$$

$$\tau_i^e = t_i^e + 168w_i \quad \forall i \in P \quad (5)$$

$$\tau_i^a \leq \tau_i^s \quad \forall i \in P \quad (6)$$

$$\tau_i^e + t_{ij}^\alpha \leq \tau_j^a + M(1 - x_{ij}) \quad \forall (i, j) \in A \quad (7)$$

$$w_1 = 0 \quad (8)$$

$$\sum_{j \in P \setminus \{o_k\}} f_{k o_j} = a_k \quad \forall k \in K \quad (9)$$

$$\sum_{j \in P \setminus \{d_k\}} f_{k j d_k} = a_k \quad \forall k \in K \quad (10)$$

$$\sum_{(j,i) \in A'} f_{kji} = \sum_{(i,j) \in A'} f_{kij} \quad \forall k \in K, i \in P \setminus \{o_k, d_k\} \quad (11)$$

$$\sum_{k \in K} f_{kij} \leq ux_{ij} \quad \forall (i, j) \in A' \quad (12)$$

$$x_{p+1,1} = 1 \quad (13)$$

$$x_{ij} \in \{0, 1\} \quad \forall (i, j) \in A' \quad (14)$$

$$\tau_i^a, \tau_i^s, \tau_i^e \geq 0 \quad \forall i \in P \quad (15)$$

$$f_{kij} \geq 0 \quad \forall (i, j) \in A', k \in K \quad (16)$$

$$w_i \in \mathbb{Z}^+ \quad \forall i \in P \setminus \{1\} \quad (17)$$

The objective function (1) includes the cost for vessels and the sailing cost for the complete service. The variable w_{p+1} contains the week of the last port visit. Since we are assuming a weekly service frequency, the last week also specifies the number of vessels. Constraints (2) and (3) require that the service enters and leaves each port exactly once, except for the initial and ending port. The time the vessel starts and ends its visit at each port is determined in Constraints (4) and (5). The term $168w_i$ converts the week selected for customer i to the appropriate hour from the start of service. Constraints (6) restrict the arrival time to be before the start of the visit at a port.

We use a chance constraint to ensure that the service route has enough buffer at each port to achieve the level of punctuality requested. Constraints (7) model a general constraint that can be used with any probability distribution in which the inverse CDF can be computed. These constraints force the arrival time at the next port to always be greater than the departure time at the previous port plus the travel time with an adequate buffer. We compute a bound for M as follows. First, we compute a maximum path length (worst case tour) traveling salesman problem (TSP) through all of the ports.² Given this “slow” TSP path, we then add extra buffer based on our service level formulas to the path if required to reach a particular service level, and further add buffer that considers when the vessel arrives at a port and how long it must wait until the port is available.

To compute the value of t_{ij}^α , consider, for example, the case of the normal distribution where we want to guarantee a confidence level of α that we arrive on time. Assume we have a parameter z^α that provides the number of standard deviations associated with a service level of α . We can then set t_{ij}^α to be the mean travel time between i and j plus z^α times the standard deviation. Consider now the case of a non-normal distribution such as the three-parameter log-logistic distribution. Given the scale parameter σ , shape parameter ξ and location parameter μ , we want to find the value of t_{ij}^α providing the minimum required travel time between i and j enabling a service level α . We can compute this from the inverse CDF as follows:

$$F^{-1}(\alpha; \sigma, \xi, \mu) = \left(\frac{1}{\alpha} - 1\right)^{-1/\sigma} \left(\xi + \mu \left(\frac{1}{\alpha} - 1\right)^{1/\sigma}\right).$$

If we are interested in comparing the results with a deterministic model, as we do in our experiments, we can simply set t_{ij}^α to a fixed value t_{ij} representing the travel time at the design speed with no added buffer.

The visit to the initial port is specified to be in the initial week by Constraints (8). Constraints (9), (10), and (11) are used to handle the flow of demands between their supply and destination points as well as between transshipment points. These are modeled in a similar way to Plum et al. (2014), but adapted to require that all demand is served. The capacity of the vessel is enforced by Constraints (12) on each arc, but only if a vessel is sailing on the arc. Should an arc not be used in the vessel path, the flow capacity is set to zero. Constraint (13) connects the last port to the first port, ensuring that the service’s route is complete. Lastly, Constraints (14), (15), (16), and (17) define the variables appropriately.

² First, we note that this version of the TSP is NP-complete, but since our problems do not contain many ports, it is easy to solve. Second, in later phases when speed can be adjusted we use the slowest vessel speed allowed.

4.2 Phase 2: optimized speed model

While the Phase 1 model covers a number of basic properties of liner shipping service design, in practice, planners can adjust the speed of the vessels to reduce sailing costs or increase the probability of meeting tight deadlines. Hence, we adjust the fourth assumption to the following:

4. Vessels may sail between their minimum and maximum speed, and the cost of doing so will be reflected in the objective function.

We implement variable vessel speeds with several new sets of variables and parameters. Since the distance between all ports is known in advance, the model can simply decide the sailing duration for each leg of the service and the speed is known (duration/distance). The duration is used in a piecewise linearization to determine the approximate bunker costs. Three new sets of variables are required:

- γ_{ij} The cost of sailing between i and j for arc $(i,j) \in A$
- ρ_{ij} The duration for sailing between ports i and j for arc $(i,j) \in A$
- β_{ij}^α The amount of buffer added into the schedule between ports i and j for arc $(i,j) \in A$ to enforce a service level α

We also need to add further parameters:

- δ The speed of a vessel
- δ^* The design speed of a vessel
- c^B Cost per ton of bunker fuel
- ϕ_{ij}^g Slope of secant g of the bunker consumption cost function approximation for arc $(i,j) \in A$
- ω_{ij}^g y -intercept of secant g of the bunker consumption cost function approximation for arc $(i,j) \in A$
- t_{ij} The sailing time at the vessel's design speed between i and j for arc $(i,j) \in A$
- $t_{ij}^{Min}, t_{ij}^{Max}$ The minimum or maximum sailing time between i and j for arc $(i,j) \in A$

Since the fuel consumption function of vessels is roughly cubic (Brouer et al. 2013a), we use a piecewise linearization of the function in our Phase 2 model. One approximation for the non-linear bunker consumption function (Brouer et al. 2013a) is given by

$$B(\delta) = \left(\frac{\delta}{\delta^*} \right)^3 B(\delta^*),$$

where δ is the vessel's speed, δ^* is the vessel's design speed and $B(\delta^*)$ is the bunker consumption in tons of fuel per hour at the design speed. We use the secant approximation presented in Reinhardt et al. (2016) to linearize this function and note that, due to the convexity of the function, no binary variables are necessary. We now show their approximation, adjusting the notation to fit our model. First, we modify the bunker consumption function to accept a travel duration between two ports with

$$\hat{B}(\rho_{ij}) = \left(\frac{t_{ij}}{\rho_{ij}} \right)^3 B(\delta^*)c^B,$$

where ρ_{ij} is the travel duration, t_{ij} is the sailing time at the vessel's design speed as previously defined, and c^B is the cost per ton of bunker fuel. This function can be approximated with a given number of secants, θ . Each secant g is defined with the function

$$\hat{c}_{ij}^g(\rho_{ij}) = \phi_{ij}^g \rho_{ij} + \omega_{ij}^g,$$

where ϕ_{ij}^g is the slope of the secant and ω_{ij}^g is the y-intercept. Then, we can add the following linear constraints to our model for computing the sailing cost:

$$\gamma_{ij} \geq \phi_{ij}^g \rho_{ij} + \omega_{ij}^g x_{ij} \quad \forall (i, j) \in A, 0 \leq g < \theta. \tag{18}$$

Since the variables x_{ij} indicate whether an arc is being used or not, we ensure that the sailing costs on the arc are only constrained when the arc is being used. Furthermore, we replace the objective function (1) with the following:

$$\min c^v w_{p+1} + \sum_{(i,j) \in A} \gamma_{ij}. \tag{19}$$

To handle the arrival time service levels with varying speeds, we modify Constraints (7) to include the decision variable ρ_{ij} :

$$\tau_i^e + \rho_{ij} + \beta_{ij}^\alpha \leq \tau_j^a + M(1 - x_{ij}) \quad \forall (i, j) \in A. \tag{20}$$

In other words, the travel time from Phase 1 is replaced with the variable specifying the travel duration on the arc, if the arc is used. We then require constraints ensuring that the minimum duration according to the chosen service level β_{ij}^α is enforced:

$$\rho_{ij} + \beta_{ij}^\alpha \geq t_{ij}^\alpha x_{ij} \quad \forall (i, j) \in A. \tag{21}$$

For a deterministic model, the term β_{ij}^α can be set to zero. We then constrain the minimum and maximum duration of the voyage using the parameters t_{ij}^{Min} and t_{ij}^{Max} with the following constraint:

$$t_{ij}^{Min} x_{ij} \leq \rho_{ij} \leq t_{ij}^{Max} x_{ij} \quad \forall (i, j) \in A. \quad (22)$$

These constraints ensure that vessels do not sail faster or slower than is allowed should the arc be chosen. Finally, we impose bounds on the travel time, travel costs, and buffer variables:

$$\rho_{ij} \geq 0, \gamma_{ij} \geq 0, \beta_{ij}^{\alpha} \geq 0 \quad \forall (i, j) \in A. \quad (23)$$

4.3 Phase 3: optimized speed with maximum transit times model

For our Phase 3 model, we require that demands are not carried for longer than their maximum transit time. Requiring that all demand be served in combination with transit times limitations can lead to infeasibility, especially when a high level of on-time arrival is desired. A further assumption of our Phase 3 model is that the demands at a port are the same each week. This is, of course, not the case in practice, but it is sufficient for tactical level planning. In Phase 1 and 2, we wanted to see the impact of service levels and speed optimization on serving the same demands, so we relaxed this transit time requirement. Here, to enforce the maximum transit times, we will relax the fifth assumption of Phase 1, namely that all demands must be carried. We now have the following assumptions in addition to the Phase 2 assumptions (minus assumption 5 from Phase 1):

10. All demands have a maximum transit time. This is a common assumption in practice.
11. Demands can be rejected if their maximum transit times cannot be met.
12. The objective is to maximize the profit, i.e. the revenue from carrying cargo minus the cost of sailing and cost of vessels.

These changes require two new variables:

- y_k This equals 1 if demand k is transported or 0 if it is not transported for $k \in K$
- δ_k This equals 0 when the service start time of the destination is greater than the end service time of the cargo source for demand k for $k \in K$. However, due to the cyclical structure of liner shipping routes, it is possible that the start service time of the destination is actually less than the end service time of the source if the path of the cargo crosses the (arbitrary) “end” of the service. When this occurs, this variable will be 1.

We also need to add further parameters:

- r_k The revenue of delivering a demand k
- l_k The maximum transit time of a demand k

For our new profit maximizing objective, we add the revenue of delivering a demand, r_k , to the objective function, and subtract the vessel and sailing costs from it:

$$\max \sum_{k \in K} r_k y_k - c^v w_{p+1} - \sum_{(i,j) \in A} \gamma_{ij}. \tag{24}$$

We next replace Constraints (9) and (10), which control the start and end of the flow of demand k , with the following:

$$\sum_{j \in P \setminus \{o_k\}} f_{koj} = a_k y_k \quad \forall k \in K \tag{25}$$

$$\sum_{j \in P \setminus \{d_k\}} f_{kj d_k} = a_k y_k \quad \forall k \in K. \tag{26}$$

Constraints (25) and (26) set the amount of containers to be carried for a particular demand to 0 when y_k is 0, ensuring that the demand is either taken in its entirety or not at all.

Given the maximum transit time of a demand k , l_k , we use the following constraints to restrict the transit time:

$$\tau_{d_k}^s \geq \tau_{o_k}^e - M \delta_k \quad \forall k \in K \tag{27}$$

$$\tau_{o_k}^s \geq \tau_{d_k}^e - M(1 - \delta_k) \quad \forall k \in K \tag{28}$$

$$\tau_{d_k}^s - \tau_{o_k}^e \leq l_k + M \delta_k + M(1 - y_k) \quad \forall k \in K \tag{29}$$

$$\tau_{d_k}^s - \tau_{o_k}^e + 168 w_{p+1} \leq l_k + M(1 - \delta_k) + M(1 - y_k) \quad \forall k \in K \tag{30}$$

Constraints (27) and (28) set the value of δ_k depending on the schedule of the vessel. When $\delta_k = 0$ (the “normal” case), the destination of demand k is scheduled later than the origin. However, when $\delta_k = 1$, the destination comes after the “end” of the service, meaning the scheduled time of the origin is actually later than the destination in the model. Constraints (29) and (30) set the transit time limitation depending on the value of δ_k . We then add one more set of constraints to the model to connect the flow of each demand to its corresponding 0/1 variable:

$$f_{kij} \leq a_k y_k \quad \forall k \in K, (i, j) \in A. \tag{31}$$

Finally, we add bounds constraints for the new variables:

$$y_k \in \{0, 1\} \quad \forall k \in K \tag{32}$$

$$\delta_k \in \{0, 1\} \quad \forall k \in K. \tag{33}$$

5 Computational experiments

We will first describe how we create our instances and experimental design in Sect. 5.1. For our Phase 1 and 2 models, we simulate operations of the created schedules to evaluate their performance with regard to cost and service quality. We describe our simulation in Sect. 5.2. Then, we present our results of tactical planning for the three models in Sect. 5.3. Finally in Sect. 5.4, we present sample schedules generated by our different models.

5.1 Instances

For our experiments, we want to compare the results from deterministic and stochastic variants (with different service levels) of each model across a variety of instances. This will help us understand the impact of different service levels, the ability to optimize speed in the planning phase, and the impact of the maximum transit times. We also want to understand how the underlying features of the created services, such as the locations of the ports being considered, impact the results.

The instances are constructed using real demand data and vessel information from the LINERLIB (Brouer et al. 2013a), which we combine with empirical travel time distributions constructed from COSCO liner shipping services (as discussed in Sect. 3). We build 44 instances based on 22 existing COSCO services as follows. For each COSCO service, we make a “standard” instance and a “tight” instance using the service’s port calls. The standard instance matches the time windows from the COSCO data, while the tight instance modifies the time windows to provide a schedule with time windows as close together as possible. That is, all buffer between ports i and j in the service’s port sequence is removed such that a vessel would have to sail at maximum speed to be on time. However, we are only able to remove buffer time up to a certain point, because the sum of all travel and port times has to be divisible by the number of hours in a week. We therefore remove all buffers between all port calls except for the return trip of the vessel to the origin port. For each instance, we extract relevant demands from the LINERLIB, including the amount of containers and the maximum transit time. For Phase 3, we also create a “relaxed” version that allows the cargo to be delivered up to 1.5 times later than the original transit time. This creates a total of 88 instances for Phase 3.

Since the LINERLIB data has been designed for the liner shipping network design problem, i.e., for problems that may contain several services, it is not perfectly suited towards single service design. This is because a significant amount of cargo is transported through hubs rather than along direct routes. To ensure our instances have a realistic amount of cargo, we aggregate ports into 25 clusters (similar to the aggregation performed in Mulder and Dekker 2014). We then assume each of those clusters has some feeder services that will carry cargo to and from ports on the service we are designing. In some cases, especially with cargo originating from Asia, this can result in too much demand for any vessel to carry. We then randomly

Table 2 Instance properties for the instances tested in this work. We note that on psw5, the vessels from the carrier are slightly faster than the vessel we use from the LINERLIB, hence one less vessel is required

Service	Ports	Demands	Minimum vessels	Service vessels	Vessel type
abx	9	13	7	8	Post Panamax
aesa	12	14	10	13	Super Panamax
awe1	7	20	8	10	Super Panamax
awe2	7	12	9	10	Super Panamax
awe3	10	37	8	10	Super Panamax
awe4	7	18	9	11	Super Panamax
awe8	9	33	9	11	Super Panamax
cen	7	8	6	7	Super Panamax
ces	9	27	9	10	Super Panamax
ese	11	18	7	8	Super Panamax
fall	11	18	9	11	Super Panamax
fax	5	6	7	7	Super Panamax
fwas	9	19	9	12	Super Panamax
fwax	10	22	12	12	Super Panamax
ne2	10	31	9	10	Super Panamax
ne6	11	29	10	11	Super Panamax
ne7	10	28	9	10	Super Panamax
psw1	4	6	5	6	Super Panamax
psw5	6	8	7	6	Super Panamax
tas1	7	16	4	4	Super Panamax
wsa	10	34	9	10	Super Panamax
wsa2	11	30	9	10	Super Panamax

reduce the size of the demands. We subtract the amount of time required to carry the cargo to the hub from the transit time limit, along with three days for transshipment as done in Guericke and Tierney (2015).

An overview of our instances³ is given in Table 2. The service name corresponds to the COSCO service name. Note that our services in some cases do not exactly correspond with the original services since we remove duplicate port calls. In addition to the number of ports and port-to-port demands, we compute the minimum number of vessels necessary to sail on the service using the original schedule and provide the number of vessels the carrier assigned to each service. The vessel type dictates the design speed, the minimum and maximum speed, as well as the (fixed) bunker costs in tons per day at design speed (which we get from LINERLIB). For vessels of the Post Panamax and Super Panamax classes, these values are as follows: design speed 16.5/17.0 knots, minimum speed 12 knots, maximum speed 23/22 knots, bunker costs 82.2/126.9 tons/day. We also assume a fuel price of \$400 per ton.

³ We provide instance data at the following repository: https://bitbucket.org/dotbielefeld/lsssd_instances.

We solve all instances with Gurobi version 7.0.2 (Gurobi Optimization 2015) on eight Intel Xeon E5506 CPUs at 2.13 GHz with a maximum runtime of 24 h. For each instance, we solve all of the models in Sect. 4. We experiment with modeling the underlying travel time data as a normal distribution and as an ll3p distribution we identified to be a better fit. We experiment with both to see if modeling travel time data with a normal distribution, which is easy to use and fit, offers a similar quality performance as the more accurate distribution. For both distributions, we experiment with service levels of 70 and 90%, and, for the ll3p distribution, also for a service level of 95% because of its long tail. We also test a deterministic version with no added buffer. This yields 8 runs for each instance and for each model. The output of each run is an optimized schedule that contains the ordering of ports of a service (a service design), the associated number of vessels that minimize the particular objective, the optimized objective value, and the scheduled arrival and departure times at ports. In Phase 2 and Phase 3, the optimized speed level between each pair of ports is also reported. We note that Gurobi is able to solve most of our models in relatively little CPU time, although some are not solved after 24 h. A detailed look at the runtime is provided in the “[Runtime](#)” section of the Appendix.

5.2 Simulation

For each service that is created by Phase 1 and 2 models, we simulate the resulting schedule to evaluate the impact of the varying service levels on “actual” costs and reliability, imitating operations of individual service runs. The purpose of the simulation is to analyze the impact of the stochastic travel time information and the possibility of adapting speeds to fulfill service levels in liner shipping network design. From a tactical perspective, the optimization models can create liner shipping services considering uncertainty through chance constraints as presented in Sect. 4. From an operational perspective, the simulation now tries to realize the optimal plans from the tactical level. We imitate that by drawing random travel times as an input for the realization of a service. We let the simulation react with speed adaptation according to the same degrees of freedom and cost parameters as our mathematical models, which causes differences to the planned speeds and to the planned variable sailing costs. Then, we can directly compare the planned costs (from optimization) with the “actual” realized costs (from simulation) for the different optimization models.

To create this simulation, for each arc of a service, we sample the travel time between ports using the ll3p distribution that has been fit to the particular OD pair as discussed in Sect. 3. If the sampled travel time for an arc would cause a delay at the destination port, the simulation increases the speed of the vessel (and reduces the travel time) to make up for the delay, inducing higher sailing costs. If the sampled travel time is less than what is required to arrive at the destination port on time, the simulation decreases the speed (and increases the travel time) to reduce fuel

consumption and save money. Following the degrees of freedom of our optimization model in Phase 2, the simulated speed and travel times are limited by the minimum and maximum speeds as defined by the vessel type from the LINERLIB. As a consequence, in case of speed ups, the simulated vessel may not be able to make up for all of the delay at the destination port. The remaining lateness is then propagated to the next arc until the service terminates. We propagate lateness in the simulation because, as mentioned earlier, it is hard to anticipate which recourse action would be taken at each port where lateness might occur. By letting it accumulate in the simulation, we can easily evaluate the quality of the tactical planning and get an idea of the worst case performance of the service.

More formally, the simulation creates individual services runs from an optimized schedule (including order of ports, scheduled departure and arrival times at ports, and optimized speeds [only Phase 2]) as provided by Phase 1 and Phase 2 optimization models. The discrete-event simulation process is as follows:

1. For each arc between ports i and j , a travel time t_{ij}^{rand} is sampled from the appropriate travel time distribution. For the ll3p distribution, if $t_{ij}^{rand} > 10\mu_{ij}$, where μ_{ij} is the empirical mean travel time from i to j , then t_{ij}^{rand} is set to $10\mu_{ij}$. We truncate the ll3p-generated travel times because the fitted ll3p distributions often yield a very long tail, inducing unrealistically long sampled travel times. Even if such a long travel time may occur in practice, e.g., due to a storm or a port strike, we would not be able to deal with these by means of a more reliable schedule or an adaption of speeds on a tactical planning level. In these cases, ports would need to be skipped and other significant changes would need to be made, which we do not want to consider here to allow for a fair comparison between tactical planning and operations.
2. The delay at port j when sailing from i to j is computed as $e_{ij} := -(t_j^s - t_{ij}^{rand} - t_i^e)$, assuming that departure is at t_i^e .
3. To determine the arc-specific speed δ_{ij} that a vessel would need to sail to arrive at port j on time, the simulation computes the sailing duration as follows:

$$\rho_{ij} = t_j^s - (t_i^e + e_{ij}).$$

Here, we assume that today's sailing time between two ports i and j is random, but known when leaving port i so we can react with an optimized speed level. Implementing the idea of minimum and maximum sailing times from Constraints (22), the sailing duration is restricted by minimum and maximum sailing times:

$$t_{ij}^{Min} \leq \rho_{ij} \leq t_{ij}^{Max}.$$

Finally, the arc-specific speed can be derived from the distance between the ports Δ_{ij} , resulting in

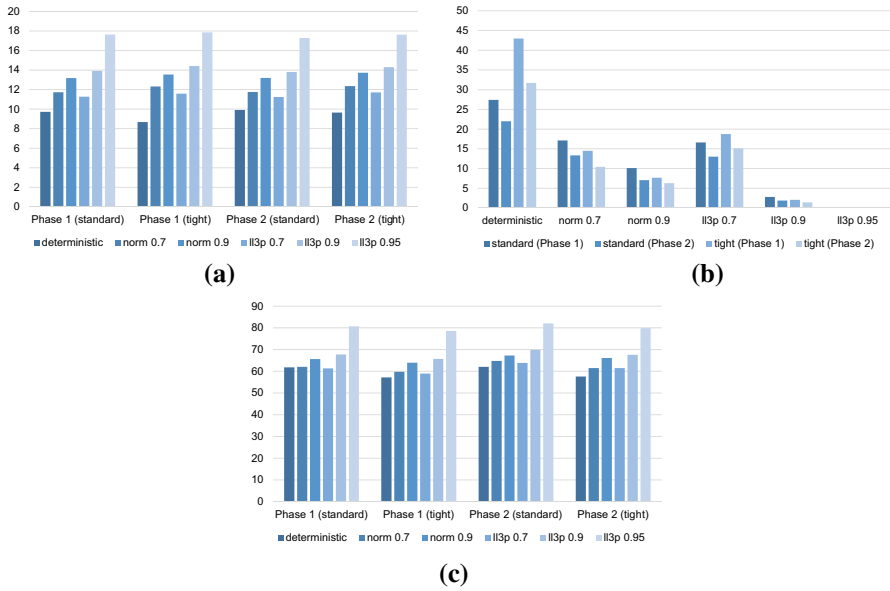


Fig. 3 Visualization of simulation results. **a** Number of vessels from optimization in Phases 1 and 2, **b** analysis of hours late for different travel time distributions, **c** total simulated costs

$$\delta_{ij} = \Delta_{ij} / \rho_{ij}.$$

4. In case of a large positive delay, the resulting speed up may still not be enough to achieve an on-time arrival, and the remaining delay will then be propagated to the next arc.

Let us look at a numerical example for a schedule with $t_j^s = 1471.0$, $t_i^e = 918.0$ and scheduled travel time of 549.71. The sample travel time from ll3p distribution is $t_{ij}^{rand} = 558.786$. Next, $e_{ij} = -(1471.0 - 558.786 - 918.0) = 5.786$. Last, $\rho_{ij} = 1471.0 - (918.0 + 5.786) = 547.214$, i.e. $\delta_{ij} = 9345.0 / 547.214 = 17.08$. The simulation is coded in Java 8 and run on a 64-bit Windows 10 machine. We run 100,000 simulations per schedule.

5.3 Results

First, we present the results of computational experiments for Phase 1 and Phase 2 models including the results of the simulation. We then provide the results from experiments with the Phase 3 model.

5.3.1 Phase 1 and 2 results

Figure 3 presents the simulation results for Phases 1 and 2. Detailed results concerning variable and fixed costs are provided in the “[Detailed Phase 1 and 2 simulation results](#)” section of the Appendix. Figure 3a compares the average optimal number of vessels for the different instances and phases. For the standard instances, the number of vessels required for *ll3p* with 95% service level is almost double the number in the deterministic results (17.64 vs. 9.73). For the tight instances, the vessels required for *ll3p* with 95% service level are slightly more than double the number in the deterministic results (17.86 vs. 8.68). The average number of vessels is slightly larger for the tight instances than for the standard instances except when service levels are considered. This indicates the deterministic solutions do not have many natural buffers in the solutions.

Figure 3b displays the results of the simulations with the different travel time distributions with regard to average hours of lateness at ports. It is obvious that planning in a deterministic way leads to a significant amount of lateness on average, especially for the tight instances (up to 43 h in Phase 1 and up to 32 h in Phase 2). Interestingly, optimizing for speed levels (Phase 2) produces schedules that help reduce the average amount of lateness. However, the best option to reduce lateness is the inclusion of buffers. With *norm 0.7* or *ll3p 0.7* optimized schedules, lateness can be reduced significantly. Considering the long tail of the *ll3p* distribution, planning with a buffer based on *ll3p 0.9* makes lateness almost disappear, of course at the cost of a large number of required vessels. Here, the value of the more realistic distribution, especially its long tail, comes into play for all instances and phases.

Finally, Fig. 3c visualizes the total simulated costs, which include variable costs arising from realizing schedules by simulation and fixed vessel costs provided by optimization. This metric yields the operational costs and should match the total costs of tactical planning in an idealized setting. The total costs increase by 49% for the standard and 61% for the tight instances. For almost all instances and phases, a higher service quality comes at a higher total cost, mainly due to the larger number of required vessels. For Phase 2, the variable sailing costs decrease for all instance types, and the number of vessels remains close to the values found in the Phase 1 optimization results. The average planned tactical speed level drops from 16.5 or 17 knots in Phase 1 to as low as 12.46 knots with the *ll3p 0.95* service level. Interestingly, for standard instances in Phase 1, including buffers based on *ll3p 0.7* can reduce total simulated costs a bit compared with the deterministic schedule. Enforcing a high service level based on *ll3p 0.95* increases total costs greatly, though.

5.3.2 Phase 3 results

The main difference between Phase 3 and Phase 2 is the inclusion of revenue generation and transit time limitations for container demands. For our experiments, we focus on how much demand is carried under different service guarantees with these new limitations. Table 3 shows the percentage of total available demand carried for various *ll3p* service levels, along with the average increase in port-to-port travel time over the deterministic Phase 3 model solution. We provide results for

Table 3 The percentage of total cargo carried in the Phase 3 model and increase in port-to-port transit times over the deterministic solution for several service levels using the ll3p distribution

Service	Cat.	Cargo carried (%)								Transit time increase (%)					
		LL				LL 1.5				LL			LL 1.5		
		Det	0.7	0.9	0.95	Det	0.7	0.9	0.95	0.7	0.9	0.95	0.7	0.9	0.95
abx	S	93	86	70	69	100	97	80	76	11	956	1446	9	854	1173
	T	97	81	73	73	99	97	80	76	93	133	237	125	1132	289
aesa	S	78	32	22	20	97	78	38	36	44	102	163	21	109	161
	T	78	22	22	22	97	78	36	32	74	134	187	37	111	179
awe1	S	39	15	14	14	79	40	39	22	42	107	149	102	169	188
	T	39	16	16	15	42	39	26	26	26	45	96	162	206	253
awe2	S	15	15	15	13	29	16	16	15	12	-1	35	30	38	69
	T	15	15	15	13	36	18	16	15	11	-1	32	34	56	70
awe3	S	19	11	6	6	36	19	15	10	67	125	192	47	75	132
	T	18	10	8	8	38	16	13	10	103	136	167	68	102	158
awe4	S	31	11	8	8	38	31	21	21	63	173	278	128	134	197
	T	21	20	12	8	36	31	21	19	13	93	257	76	109	202
awe8	S	11	8	8	2	34	10	11	8	252	307	422	148	201	271
	T	11	6	5	2	35	18	12	10	238	122	473	51	99	143
cen	S	79	16	7	7	100	72	65	58	68	90	144	20	55	93
	T	72	11	11	7	100	69	62	40	99	134	145	36	76	200
ces	S	90	47	45	43	93	83	65	59	161	201	247	73	176	150
	T	62	28	6	6	93	76	72	69	55	143	192	30	51	78
ese	S	70	15	14	10	85	85	42	28	33	87	98	0	42	57
	T	63	18	14	14	85	68	61	60	122	153	172	54	63	91
fall	S	88	44	44	7	100	83	77	47	75	109	193	33	75	145
	T	79	28	18	18	100	77	67	45	86	113	212	31	78	160
fax	S	99	35	33	33	100	99	67	64	28	123	260	52	111	226
	T	100	33	33	1	100	99	67	34	74	107	196	40	115	165
fwas	S	98	69	35	15	99	99	74	57	19	210	328	5	68	100
	T	98	76	36	11	99	100	74	57	37	105	332	140	256	375
fwax	S	60	30	20	12	81	79	57	39	72	94	188	71	329	357
	T	61	30	20	12	81	80	73	73	86	76	154	37	166	266
ne2	S	81	41	38	38	98	87	74	60	155	207	268	34	150	230
	T	83	47	45	38	97	83	75	59	35	215	188	25	138	173
ne6	S	77	12	8	7	98	97	70	53	79	179	213	41	232	391
	T	76	27	20	7	98	83	71	53	58	83	186	47	242	401
ne7	S	82	47	58	38	98	84	77	63	30	245	215	28	109	208
	T	83	37	28	28	98	86	67	63	71	96	128	26	166	208
psw1	S	45	1	1	0	100	40	40	39	49	76	92	27	43	60
	T	40	1	1	1	100	40	39	32	76	93	93	42	59	79
psw5	S	64	22	22	2	96	74	64	62	66	114	138	45	89	145
	T	29	2	2	2	94	33	25	24	209	340	528	61	125	181

Table 3 (continued)

Service	Cat.	Cargo carried (%)								Transit time increase (%)					
		LL				LL 1.5				LL			LL 1.5		
		Det	0.7	0.9	0.95	Det	0.7	0.9	0.95	0.7	0.9	0.95	0.7	0.9	0.95
tas1	S	79	45	39	35	85	70	70	64	73	70	109	82	82	109
	T	70	39	39	35	85	70	69	69	55	55	93	81	81	101
wsa	S	94	39	37	37	97	86	69	44	163	268	284	50	143	228
	T	94	20	17	4	98	86	65	43	169	207	514	43	214	350
wsa2	S	85	57	18	18	95	88	57	50	22	115	154	227	447	470
	T	87	58	18	18	95	86	53	51	57	161	205	49	225	131

two different settings of maximum transit times. The “LL” results use the maximum transit time for container demands as listed in the LINERLIB, whereas “LL 1.5” uses the LINERLIB maximum transit times multiplied by 1.5. As in Phases 1 and 2, we evaluate each service with a standard (S) and a tight (T) instance. We provide these longer durations because the ones in the LINERLIB are tailored to Maersk Line’s network, which is somewhat different than COSCO’s network. The longer durations also show how much more demand can be carried when the restrictions are relaxed.

Our results show that when using the transit times in the LINERLIB, even a service guarantee of 0.7 results in less than half of the containers of the deterministic solution being carried on 12 services. The increase in transit times between ports to achieve service levels is why cargo is not being carried, as in many cases the transit time is over 50% higher for a service level of 0.7. Raising the maximum transit times by 50% allows significantly more containers to be transported, resulting in only one service where both the S and T variants at the 0.7 service level carry less than half of the containers of the deterministic version. We note that in the case of *awe2* at service level 0.9, the average transit times slightly decrease over the deterministic solution due to the selection of a different route.

The message for carriers from these results is clear: implementing a punctuality guarantee requires them to either charge higher freight rates to shippers or to convince shippers to accept higher maximum transit times. However, this work provides a mechanism for quantitatively assessing how much extra revenue or transit time would be necessary to run a profitable service. For example, on the *abx* service in the standard case, the carrier only needs to convince their customers to accept 11% longer transit times (on average around the service) for a 70% service guarantee. Our model is particularly useful in this respect in comparison to models that do not optimize the vessel route, as such models will not yield any insights when adjusting transit time limitations or cargo revenues.

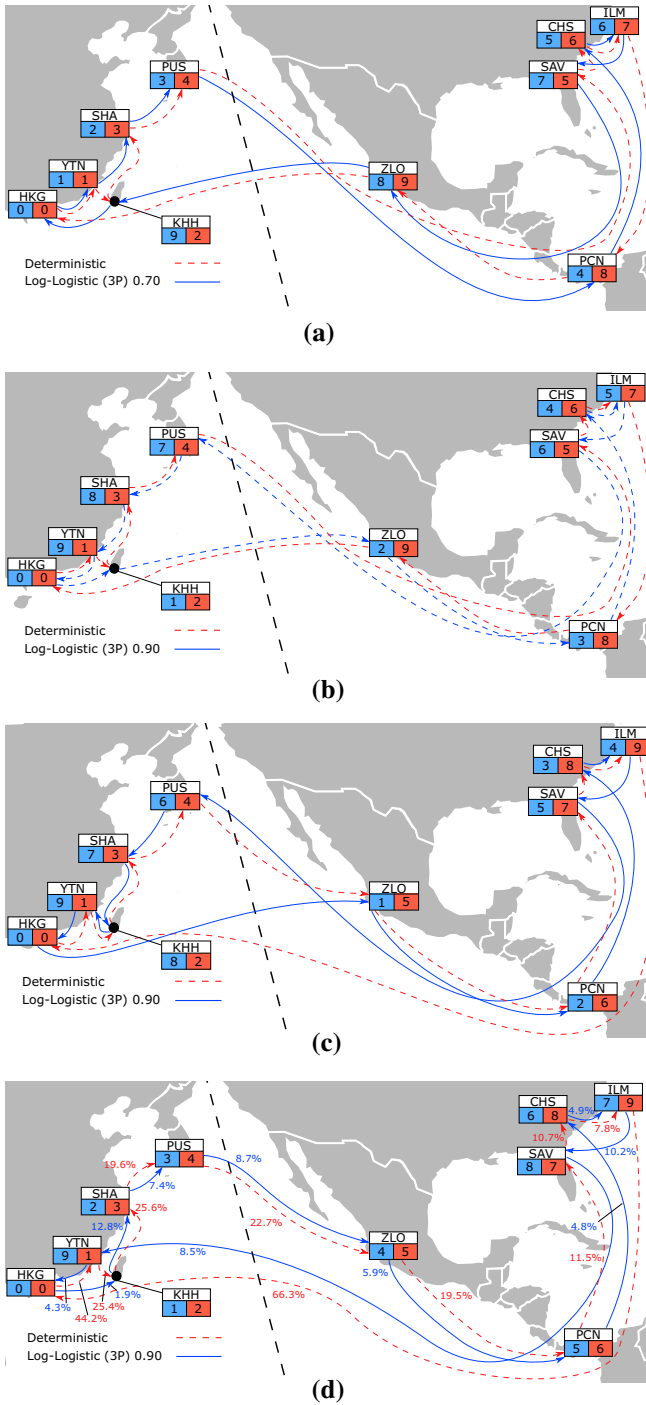


Fig. 4 Visualizations of the deterministic and 0.9 ll3p service level solutions for service awe3. **a** Phase 1 (Det/0.7), **b** Phase 1 (Det/0.9), **c** Phase 2, **d** Phase 3 (LL 1.5)

5.4 Sample service plans

Depending on the optimization model of the particular phase, the resulting schedules may differ significantly, even when total costs are similar. We present some of the created schedules to understand the impact of the different assumptions on them.

Figures 4a–d compare the optimal schedules for the different phases of service *awe3* (standard). This service connects ports in North America with ports in East Asia. In Fig. 4, the dashed line reflects the result of the corresponding deterministic optimization, and the solid blue line shows a given service level. The sequence number for each port is shown in a box with a color matching either the deterministic or 0.7/0.9/0.95 service level solution. For Phase 1 (see Fig. 4a), the optimal deterministic port order is HKG–YTN–KHH–SHA–PUS–SAV–CHS–ILM–PCN–ZLO–HKG, requiring a total of 11 vessels and sailing costs of \$2.871 million per rotation. Including buffers based on *ll3p 0.9* completely changes the order of the ports in the optimal schedule, increasing the number of vessels to 16 and slightly reducing sailing costs to \$2.837 million per rotation. Simulation of these schedules reveals a significant reduction of the speed ups required to ensure punctuality (from 40 to 4%). Interestingly, although a high service level means investing in a large number of vessels, due to slow steaming, the simulated sailing cost is so small that the simulated total costs can be reduced a bit (\$7.264 million for deterministic vs. \$7.072 million for *ll3p 0.9*).

The results of optimization of this instance for Phase 2 can be seen in Fig. 4c. Both schedules are different from their Phase 1 counterparts in the order of port visits, and the schedules operate in opposite directions. The number of vessels remains constant compared with Phase 1, but the optimized and realized speeds vary. For the deterministic solution, we have a planned average speed of 13.8 knots, and for the *ll3p 0.9* based solution a value of 12.41 knots. When optimizing speed levels, required speed ups can be reduced from 24 to 3%, reducing planned and simulated variable sailing costs. However, due to the larger number of vessels, total simulated costs increase significantly from 65.14 to 77.35, which is accompanied by an average reduction of lateness from 2.7 to 0.2 ports per trip.

The Phase 3 solutions show the drop-off in cargo on the individual legs of the solution for the 0.9 service level. Both the fronthaul and backhaul see significant reductions in containers carried, with service level 0.9 taking 2.6× less cargo on the fronthaul and 7.8× less on the backhaul. While many carriers would likely be willing to accept reductions in utilization on the backhaul in exchange for high punctuality guarantees (90% on-time would lead the industry by a significant margin), reductions in fronthaul utilization are especially bad for a carrier's profits.

6 Conclusions

We solved the liner shipping service design problem with arrival time service levels. We showed that a three parameter log-logistic distribution fits the journeys of liner ships better than other distributions, such as the uniform or normal distribution, that were previously used in the literature. Our mathematical models of the

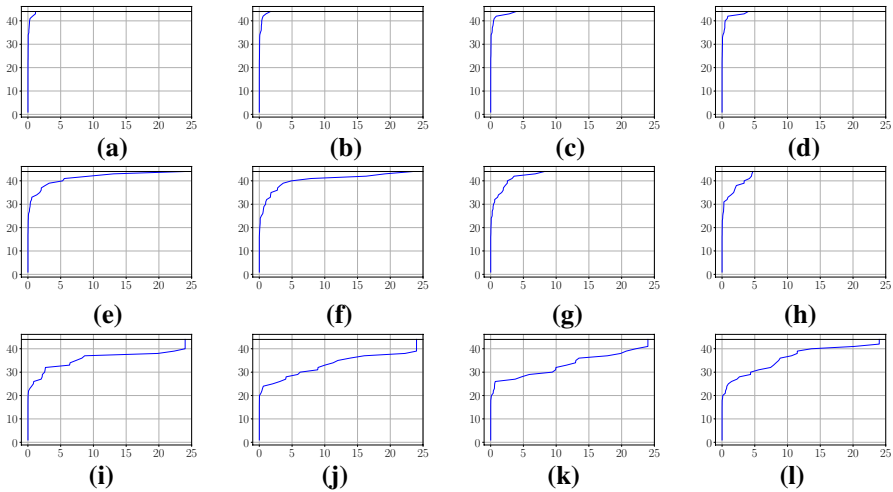


Fig. 5 The number of instances solved (y-axis) versus the runtime of Gurobi (wall time) in hours in all phases for the deterministic model and all *ll3p* models. **a** Phase 1 deterministic, **b** Phase 1 *ll3p* 0.7, **c** Phase 1 *ll3p* 0.9, **d** Phase 1 *ll3p* 0.95, **e** Phase 2 deterministic, **f** Phase 2 *ll3p* 0.7, **g** Phase 2 *ll3p* 0.9, **h** Phase 2 *ll3p* 0.95, **i** Phase 3 deterministic, **j** Phase 3 *ll3p* 0.7, **k** Phase 3 *ll3p* 0.9, **l** Phase 3 *ll3p* 0.95

service design problem use chance constraints to ensure that vessels arrive on time from a tactical planning perspective. Simulation of schedule operations allows for the comparison of planned total costs with “actual” total costs. Despite modelling a realistic version of service design, the model is nonetheless computationally tractable with a state-of-the-art mixed-integer programming solver. Our results show that services can be designed with on-time guarantees and that the route a vessel takes heavily influences the service level required. Furthermore, our model allows carriers to quantitatively assess what level of service they wish to offer and can support them during negotiations with shippers over prices and container transit time limitations.

This model provides a basis for several directions of future work. First, more detailed disruption recovery actions could be considered not only in the simulation of the model solution, but also in the model itself. Actions such as skipping ports or not fully loading/unloading containers when severely delayed could be included, however these would likely make the model significantly harder to solve. Furthermore, faster solution techniques could be considered such as branch & price or heuristic approaches. These would make the model more useful for operational decision support when negotiating berthing windows with terminals. Finally, time windows could be made optional, i.e., instead of a planner specifying time windows, the model could determine optimal time windows that the planner could then attempt to negotiate with a terminal.

Acknowledgements We thank the Paderborn Center for Parallel Computation (PC²) for the use of their high throughput cluster for the experiments in this work. We further thank Stefan Guericke for his advice and comments regarding this work.

Table 4 Simulation results for Phase 1 and 2 for different instance types

Instance	Type	Phase 1						Phase 2										
		Optimization			Simulation			Optimization			Simulation							
		Yes	Var	Total	#Late	Late/p	Faster	Spd	Yes	Var	Total	Spd	Var	#Late	Late/p	Faster	Spd	
Deterministic	Standard	9.73	24.19	61.13	24.89	2.09	27.42	45.38%	16.85	9.91	16.48	54.06	14.01	19.59	1.80	22.01	29.52%	15.13
Norm 0.7	Standard	11.73	24.52	69.03	17.55	0.65	17.11	11.93%	14.06	11.76	14.20	58.82	12.92	15.18	0.50	13.36	9.37%	13.27
Norm 0.9	Standard	13.18	24.27	74.26	15.66	0.40	10.10	6.38%	13.37	13.19	13.31	63.29	12.56	13.96	0.29	7.06	4.73%	12.78
I3p 0.7	Standard	11.27	24.31	67.14	18.49	0.77	16.62	16.46%	14.58	11.24	14.96	57.56	13.29	16.36	0.57	13.00	12.70%	13.78
I3p 0.9	Standard	13.91	24.28	76.94	15.09	0.25	2.78	4.50%	13.18	13.81	13.17	65.40	12.52	13.48	0.12	1.89	3.04%	12.64
I3p 0.95	Standard	17.64	24.59	91.28	14.02	0.01	0.11	1.55%	12.73	17.29	13.28	78.57	12.48	13.34	0.01	0.10	1.17%	12.50
Deterministic	Tight	8.68	24.22	57.14	29.16	4.22	42.96	64.73%	18.38	9.64	17.15	53.74	14.29	21.03	3.64	31.70	38.13%	15.67
Norm 0.7	Tight	12.32	24.66	71.32	18.13	0.58	14.51	12.00%	14.11	12.35	13.87	60.58	12.76	14.77	0.40	10.41	7.21%	13.07
Norm 0.9	Tight	13.55	24.36	75.74	15.90	0.37	7.67	5.41%	13.44	13.73	13.56	65.64	12.57	14.06	0.28	6.27	3.32%	12.74
I3p 0.7	Tight	11.59	24.60	68.53	19.99	0.79	18.76	19.07%	14.86	11.71	15.46	59.83	13.32	17.13	0.60	15.12	13.30%	13.88
I3p 0.9	Tight	14.41	24.29	78.87	15.30	0.15	2.08	3.91%	13.26	14.29	13.29	67.36	12.55	13.54	0.10	1.42	2.58%	12.64
I3p 0.95	Tight	17.86	24.48	92.05	14.48	0.01	0.24	1.54%	12.94	17.62	13.22	79.78	12.46	13.28	0.01	0.18	1.25%	12.49

Appendix: Additional computational results

Runtime

Our mathematical model is solved in most cases relatively quickly using the Gurobi solver. Figure 5 shows the number of instances solved at each point in time over 24 h of runtime. In Phase 1, all 44 instances can be solved within a couple of hours, with the vast majority solved in only a few minutes. The inclusion of speed optimization in Phase 2 increases the difficulty of the model, however we note that nearly all of our instances are solved within 10 h regardless of the service level desired. In Phase 3, the difficulty increases due to the extra freedom of the model to optimize the cargo intake. Here, we are able to solve over half of the instances within the first hour, with all but four instances solved before the timeout in the deterministic case, and all but two instances for $ll3p$ 0.95. For planning services that have a lifespan of months or even years, a runtime of 24 h or more is acceptable.

Detailed Phase 1 and 2 simulation results

In Table 4, we present averaged results for Phase 1 and 2 models grouped by standard and tight instances. The first column is the service level and the second column specifies the group of instances. The next three columns present results from the Phase 1 optimization (tactical perspective). The column “Ves” represents the number of vessels required for a service, “Var” is the variable sailing cost, and “Total” gives the total costs including variable and vessel costs from the optimization. Recall that all Phase 1 optimization runs assume the vessel sails at the design speed. The next five columns represent the averaged results from the 100,000 simulations of each schedule (operational perspective). This includes the variable sailing costs (“Var”), the number of late ports per service that remains in spite of speed ups (“# Late”), the average hours late per port when arrival to a port is late (“Late/p”), the percent of the time the vessel sailed above the design speed (“Faster”), and the average speed traveled in the simulation (“Spd”). The simulation assumes the use of the number of vessels selected by the optimization, which is why the total cost in the simulation is not reported. For Phase 2, the same results are reported. For the optimization, the additional column “Spd” shows the average speed level chosen by the optimization, which allows for a comparison of tactical and operational speed levels. We report the full set of Phase 1 and Phase 2 results in the electronic appendix.

References

- Agarwal R, Ergun Ö (2008) Ship scheduling and network design for cargo routing in liner shipping. *Transp Sci* 42(2):175–196
- Agra A, Christiansen M, Delgado A, Hvattum L (2015) A maritime inventory routing problem with stochastic sailing and port times. *Comput Oper Res* 61:18–30

- Agra A, Christiansen M, Hvattum L, Rodrigues F (2016) A MIP based local search heuristic for a stochastic maritime inventory routing problem. In: Paias A, Ruthmair M, Voß S (Eds) 7th International conference on computational logistics. Springer International Publishing, pp 18–34
- Álvarez J (2009) Joint routing and deployment of a fleet of container vessels. *Marit Econ Logist* 11(2):186–208
- Brouer B, Alvarez J, Plum C, Pisinger D, Sigurd M (2013a) A base integer programming model and benchmark suite for liner-shippping network design. *Transp Sci* 48(2):281–312
- Brouer B, Dirksen J, Pisinger D, Plum C, Vaaben B (2013b) The Vessel Schedule Recovery Problem (VSRP)—a MIP model for handling disruptions in liner shipping. *Eur J Oper Res* 224(2):362–374
- Brouer B, Desaulniers G, Pisinger D (2014) A matheuristic for the liner shipping network design problem. *Transp Res Part E Logist Transp Rev* 72:42–59
- Brouer B, Karsten C, Pisinger D (2017) Optimization in liner shipping. *4OR* 15(1):1–35
- Christiansen M, Fagerholt K, Ronen D (2004) Ship routing and scheduling: status and perspectives. *Transp Sci* 38(1):1–18
- Christiansen M, Fagerholt K, Nygreen B, Ronen D (2013) Ship routing and scheduling in the new millennium. *Eur J Oper Res* 228(3):467–483
- COSCO Shipping Lines (2017). <http://www.coscon.com/ourservice/toService.do>. Accessed June 2017
- Deb K, Pratap A, Agarwal S, Meyarivan T (2002) A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans Evol Comput* 6(2):182–197
- Ehmke J, Campbell A, Urban T (2015) Ensuring service levels in routing problems with time windows and stochastic travel times. *Eur J Oper Res* 240:539–550
- Guericke S, Tierney K (2015) Liner shipping cargo allocation with service levels and speed optimization. *Transp Res Part E Logist Transp Rev* 84:40–60
- Gurobi Optimization (2015) Gurobi optimizer version 7.0.2 reference manual. <http://www.gurobi.com>
- Halvorsen-Weare E, Fagerholt K (2011) Robust supply vessel planning. In: Pahl J, Reiners T, Voß S (eds) Network optimization, lecture notes in computer science, vol 6701. Springer, Berlin, pp 559–573
- Halvorsen-Weare E, Fagerholt K (2013) Routing and scheduling in a liquefied natural gas shipping problem with inventory and berth constraints. *Ann Oper Res* 203:167–186
- Halvorsen-Weare E, Fagerholt K, Ronnqvist M (2013) Vessel routing and scheduling under uncertainty in the liquefied natural gas business. *Comput Ind Eng* 64:290–301
- Hasheminia H, Jiang C (2017) Strategic trade-off between vessel delay and schedule recovery: an empirical analysis of container liner shipping. *Marit Policy Manag* 44(4):458–473
- Karsten C, Pisinger D, Ropke S, Brouer B (2015) The time constrained multi-commodity network flow problem and its application to liner shipping network design. *Transp Res Part E Logist Transp Rev* 76:122–138
- Kauczynski W (1994) Study of the reliability of the ship transportation. In: Proceeding of the international conference on ship and marine research, p 15
- Lau Y-Y, Ducruet C, Ng AK, Fu X (2017) Across the waves: a bibliometric analysis of container shipping research since the 1960s. *Marit Policy Manag* 44(6):667–684
- Lee C-Y, Lee HL, Zhang J (2015) The impact of slow ocean steaming on delivery reliability and fuel consumption. *Transp Res Part E* 76:176–190
- MathWave Technologies (2016). EasyFit Distribution Fitting Software. <http://www.mathwave.com/>
- Meng Q, Wang T (2010) A chance constrained programming model for short-term liner ship fleet planning problems. *Marit Pol Mgmt* 37(4):329–346
- Meng Q, Wang S, Andersson H, Thun K (2013) Containership routing and scheduling in liner shipping: overview and future research directions. *Transp Sci* 48(2):265–280
- Mulder J, Dekker R (2014) Methods for strategic liner shipping network design. *Eur J Oper Res* 235(2):367–377
- Norstad I, Fagerholt K, Laporte G (2011) Tramp ship routing and scheduling with speed optimization. *Transp Res Part C Emerg Technol* 19(5):853–865
- Notteboom T (2006) The time factor in liner shipping services. *Marit Econ Logist* 8(1):19–39
- Notteboom T, Rodrigue J-P (2008) Containerisation, box logistics and global supply chain: the integration of ports and liner shipping networks. *Marit Econ Logist* 10:152–174
- Plum C, Pisinger D, Salazar-Gonzalez J, Sigurd M (2014) Single liner shipping service design. *Comput Oper Res* 45:1–6
- Powell B, Perakis A (1997) Fleet deployment optimization for liner shipping: an integer programming model. *Marit Policy Manag* 24(2):183–192

- Qi X, Song D-P (2012) Minimizing fuel emissions by optimizing vessel schedules in liner shipping with uncertain port times. *Transp Res Part E* 48(4):863–880
- Reinhardt L, Plum C, Pisinger D, Sigurd M, Vial G (2016) The liner shipping berth scheduling problem with transit times. *Transp Res Part E Logist Transp Rev* 86:116–128
- Santini A, Plum CE, Ropke S (2017) A branch-and-price approach to the feeder network design problem. *Eur J Oper Res* 264(2):607–622
- Song D, Dong J (2013) Long-haul liner service route design with ship deployment and empty container repositioning. *Transp Res Part B Methodol* 55:188–211
- Song D, Li D, Drake P (2015) Multi-objective optimization for planning liner shipping service with uncertain port times. *Transp Res Part E Logist Transp Rev* 84:1–22
- United Nations Conference on Trade and Development (UNCTAD) (2015) Review of maritime transport
- Vernimmen B, Dullaert W, Engelen S (2007) Schedule unreliability in liner shipping: origins and consequences for the hinterland supply chain. *Marit Econ Logist* 9(3):193–213
- Wang S, Meng Q (2012a) Liner ship route schedule design with sea contingency time and port time uncertainty. *Transp Res Part B* 46:615–633
- Wang S, Meng Q (2012b) Robust schedule design for liner shipping services. *Transp Res Part E Logist Transp Rev* 48(6):1093–1106
- Wang S, Wang X (2016) A polynomial-time algorithm for sailing speed optimization with containership resource sharing. *Transp Res Part B Methodol Part A* 93:394–405

Kevin Tierney is a professor for Decision and Operation Technologies at Bielefeld University in Germany. He was previously an Assistant Professor at the University of Paderborn and earned his PhD from the IT University of Copenhagen in Denmark. His research interests include liner shipping optimization, metaheuristics and integrating machine learning and operations research techniques.

Jan Fabian Ehmke is a Professor for Management Science at the Otto-von-Guericke University Magdeburg, Germany, where he is a member of the Faculty of Economics and Management. He obtained a doctoral degree for his thesis on the integration of information and optimization models for routing in city logistics from the Technical University of Braunschweig, Germany. Before his current position, he was postdoc at the Technical University of Braunschweig, visiting postdoc at the University of Iowa, assistant professor at the Free University of Berlin, and university professor at Viadrina European University Frankfurt (Oder), Germany. His main research interests focus on the development of quantitative approaches at the interface of operations research and business analytics, especially for applications in transportation, mobility and logistics. He is the author or co-author of books and papers published in *EJOR*, *C&OR*, *OR Spectrum*, *Transportation Research Parts C and E*, and other journals.

Ann Melissa Campbell is the Tippie Professor of Management Sciences at the Tippie College of Business at the University of Iowa. She holds a PhD in Industrial Engineering from Georgia Institute of Technology. Her research activities are focused on the development of solution techniques for new and emerging freight transportation problems, such as same day delivery and emissions reduction for routing. She serves as the PhD coordinator for the Department of Management Sciences at Iowa and is the author of many papers published in *Transportation Science*, *EJOR*, *Networks*, among other journals.

Daniel Müller is a PhD candidate in the Department of Business Information Systems at the University of Paderborn. His thesis focuses on various topics of liner shipping optimization.