The Liner Shipping Fleet Repositioning Problem with Cargo Flows

Kevin Tierney and Rune Møller Jensen

IT University of Copenhagen, Copenhagen, Denmark
{kevt,rmj}@itu.dk

Abstract. We solve an important problem for the liner shipping industry called the Liner Shipping Fleet Repositioning Problem (LSFRP). The LSFRP poses a large financial burden on liner shipping firms. During repositioning, vessels are moved between services in a liner shipping network. Shippers wish to reposition vessels as cheaply as possible without disrupting the cargo flows of the network. The LSFRP is characterized by chains of interacting activities with a multi-commodity flow over paths defined by the activities chosen. Despite its great industrial importance, the LSFRP has received little attention in the literature. We introduce a novel mathematical model of the LSFRP with cargo flows based on a carefully constructed graph and evaluate it on real world data from our industrial collaborator.

1 Introduction

Responsible for transporting over 1.3 billion tons of cargo in 2011 [10], liner shipping networks reliably and cheaply connect the world's markets. Vessels are regularly repositioned between services in liner shipping networks to adjust the networks to the world economy and to stay competitive. Since repositioning a single vessel can cost hundreds of thousands of US dollars, optimizing the repositioning activities of vessels is an important problem for the liner shipping industry.

The Liner Shipping Fleet Repositioning Problem (LSFRP) consists of finding sequences of activities that move vessels between services in a liner shipping network while respecting the cargo flows of the network. The LSFRP maximizes the profit earned on the subset of the network affected by the repositioning, balancing sailing costs and port fees against cargo and equipment revenues, while respecting important liner shipping specific constraints dictating the creation of services and movement of cargo. A unique feature of the LSFRP is the state-based nature of the activities in the problem. Many LSFRP activities span multiple physical locations and depend on where vessels are located at particular times in order to be performed. Automated planning techniques were used to represent a high-level version of the LSFRP that ignored cargo flows in [9]. Cargo flows, however, are an important aspect of the LSFRP that drive decisions on how vessels should be repositioned.

To this end, we present a novel model of the LSFRP with cargo flows using a mixed-integer programming (MIP) model on top of a detailed graph that embeds many LSFRP constraints. We solve our model using CPLEX and study the performance of our model on real world data from our industrial collaborator. Our instances contain two actual repositioning scenarios as well as several constructed scenarios to investigate the scaling performance of our model.

2 Liner Shipping Fleet Repositioning

Liner shipping networks consist of a set of cyclical routes, called services, that visit ports on a regular, usually weekly, schedule. Liner shipping networks are designed to serve customer's cargo demands, but over time the economy changes and liner shippers must adjust their networks in order to stay competitive. Liner shippers make adjustments by adding, removing and modifying the services in their network. Whenever a new service is created, or an existing service is expanded, vessels must be *repositioned* from their current service to the service being added or expanded. Vessel repositioning is expensive due to the cost of fuel (in the region of hundreds of thousands of dollars) and the revenue lost due to cargo flow disruptions. Given that liner shippers around the world reposition hundreds of vessels per year, optimizing vessel movements can significantly reduce the economic and environmental burdens of containerized shipping, and allow shippers to better utilize repositioning vessels to transport cargo.

The aim of the LSFRP is to maximize the profit earned when repositioning a number of vessels from their initial services to a service being added or expanded, called the goal service. We focus on the case where a new service is being added to the network because expanding a service can be seen as a special case of adding a new service, in which vessels are repositioned from the service being expanded to itself along with extra vessels from elsewhere in the network.

Liner shipping services are composed of multiple *slots*, each of which represents a cycle that is assigned a particular vessel. Each slot is composed of a number of *visitations*, which represent a specific time when a vessel is scheduled to call a port. A vessel that is assigned to a particular slot sequentially sails to each visitation in the slot. Figure 1 shows a schedule of an example service that contains three slots and visits five ports. The service requires three weeks to complete a cycle, and therefore needs three vessels in order to maintain weekly frequency. Each line (black, dark gray, light gray) represents a slot, and each dot is a visitation at a port at a particular time.

Vessel sailing speeds can be adjusted throughout repositioning to balance cost savings with punctuality. The bunker fuel consumption of vessels increases cubically based on the speed of the vessel. *Slow steaming*, in which vessels sail near or at their minimum speed, therefore, allows vessels to sail cheaper between two ports than at higher speeds, albeit with a longer duration. We linearize the bunker consumption of each repositioning vessel in order to more easily model the LSFRP.

Phase-out & Phase-in The repositioning period for each vessel starts at a specific time when the vessel may cease normal operations, that is, it may stop sailing to scheduled visitations and go somewhere else. Each vessel is assigned a different time when it may begin its repositioning, or *phase-out* time. After this time, the vessel may undertake a number of different activities to reach its goal service at low cost. In order to complete the repositioning, each vessel must *phase in* to a slot on the goal service before a time set by the repositioning coordinator. After this time, normal operations on the goal service are set to begin, and all scheduled visitations on the service are to be undertaken. Figure 1 shows a phase-in service with a phase-in deadline at port c in week 2. The solid lines connect all the visitations that must be undertaken, whereas the dashed lines connect visitations that will only be carried out if they are profitable during the repositioning.

Within a *tradezone*, which is a contiguous geographical area, vessels may sail freely from their initial service to goal service, as well as back from the goal service to the initial service. However, if two ports lie in different tradezones, vessels may only sail between them when going from the initial service to the goal service. Tradezone restrictions ensure cargo is not brought to places that would violate the law, as well as help keep vessels from experiencing unexpected delays.

When a port is visited that is not in the initial or goal service, or is visited out of order, it is called an *inducement*. If a port on the initial or goal service is left off of the repositioning vessel's schedule, it is called an *omission*. Figure 2 shows a vessel's repositioning (solid line) from its initial service (dashed) to its goal service (dotted) within a tradezone. Although FXT is on both the goal and initial services, it is omitted from the repositioning. Note also that the ports RTM and BRV are induced into the repositioning path. This is only possible because the induced ports are in the same tradezone as LEH and AAR.

Cargo and Equipment Revenue is earned through delivering cargo and equipment (empty containers). We use a detailed view of cargo flows. Cargo is represented as a set of port to port demands with a cargo type, a latest delivery time, an amount of TEU[1] available, and a revenue per TEU delivered. We subtract the cost of loading and unloading each TEU from the revenue to determine the profit per TEU of a particular cargo demand. In contrast to cargo, which can be seen as a multi-commodity flow where each demand is a commodity with a start and end port, equipment can be sent from any port where it is in surplus to any port where it is in demand. Ports which have an equipment surplus or deficit can be considered to have an infinite amount of supply

[1] TEU stands for *twenty-foot equivalent unit* and represents a single twenty-foot container.

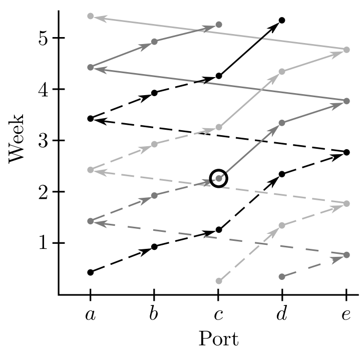

Fig. 1: A service time-space graph.

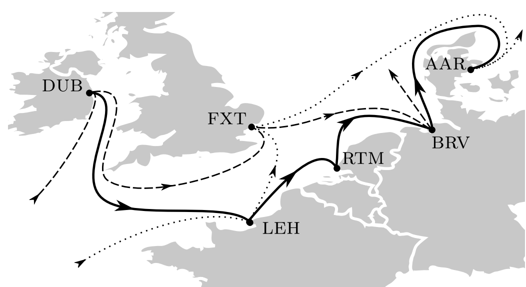

Fig. 2: An example repositioning.

or demand for a particular type of equipment. This is reasonable since the amount of extra containers on-hand or that are required tends to be much greater than the size of a vessel. Each piece of equipment brought from a port where it is in excess to a port where it is needed earns a small revenue. The revenue earned is an estimation of how much money was saved by bringing the equipment on a repositioning vessel instead of moving the equipment through other, more expensive means.

We consider both *dry* and *reefer* (refrigerated) cargo. Dry containers are standard containers with no specific handling requirements. Reefer containers, in contrast, must be stowed on a vessel in a location with a plug in order to keep the refrigeration unit running. Some ports have equipment, but are not on any service visited by repositioning vessels. These ports are called *flexible* ports, and are associated with flexible visitations. The repositioning coordinator may choose the time a vessel arrives at such visitations, if at all. All other visitations are called *inflexible*, because the time a vessel arrives is fixed.

Sail-on-service (SOS) Opportunities While repositioning, vessels may use certain services to cheaply sail between two parts of the network. These are called *SOS opportunities*. There are two vessels involved in SOS opportunities, referred to as the *repositioning vessel*, which is the vessel under the control of a repositioning coordinator, and the *on-service vessel*, which is the vessel assigned to a slot on the service being offered as an SOS opportunity. Repositioning vessels can use SOS opportunities by replacing the on-service vessel and sailing in its place for a portion of the service. SOS opportunities save significant amounts of money on bunker fuel, since one vessel is sailing where there would have otherwise been two. Using an SOS can even earn money from the *time-charter bonus*, which is money earned by the liner shipper if the on-service vessel is leased.

Consider Figure 3, in which the AC3 service is offered as a SOS opportunity to the vessel repositioning from CHX to Intra-WCSA. The repositioning vessel can leave CHX at TPP, and sail to HKG where it picks up the AC3, replacing the on-service vessel. The repositioning vessel then sails along the AC3 until it gets to BLB where it can join the Intra-WCSA. Note that no vessel sails on the backhaul of the AC3, and this is allowed because very little cargo travels on the AC3 towards Asia.

When a repositioning vessel uses an SOS opportunity, the on-service vessel is either laid-up or leased out, freeing a slot on the service. The repositioning vessel may join the freed slot in any of the *starting visitations* and may leave the slot in one of the

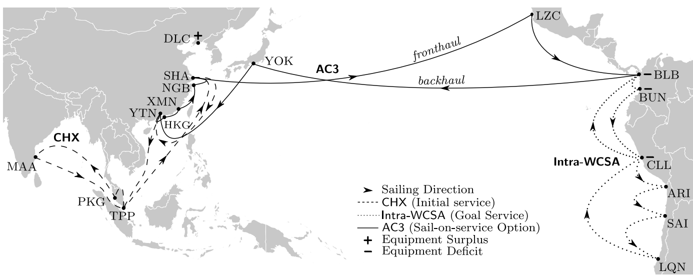

Fig. 3: A subset of the case study we performed with our industrial collaborator.

ending visitations. There are two ways for repositioning vessels to start an SOS: *transshipment* and *parallel sailing*. When starting an SOS by transshipment, all cargo loaded on the on-service vessel is transshipped (moved) to the repositioning vessel. Each TEU transshipped has a fee roughly equal the cost of loading a TEU. Transshipment is not always possible, due to *cabotage restrictions*, which are laws preventing foreign flagged vessels from offering domestic cargo services in certain markets. Cabotage rules can be legally worked around using a *parallel sailing*, in which both the repositioning vessel and the on-service vessel visit ports in tandem. The repositioning vessel only loads cargo, and the on-service vessel only discharges cargo. Although parallel sailing is expensive, since fuel consumption is doubled, it is sometimes cheaper than performing a transshipment when the transshipment fees at a port are high.

2.1 Literature Review

The LSFRP has received little attention in the literature and was not mentioned in either of the most influential surveys of work in the liner shipping domain [3,4]. Although there has been significant work on problems such as the Fleet Deployment Problem (FDP) [8] and the Network Design Problem (NDP) [1], neither captures the movement of vessels through the network inherent in the LSFRP.

Although tramp shipping problems, such as [7], maximize cargo profit in the face of sailing costs and port fees as in the LSFRP, they lack liner shipping specific constraints, such as phase-in requirements and strict visitation times. Airline disruption management [5,6], while also relying on time-based graphs, differs from the LSFRP in two key ways. First, airline disruption management requires an exact cover of all flight legs over a planning horizon. The LSFRP has no such requirement over visitations or sailing legs. Second, there are no flexible visitations in airline disruption management.

Martin W. Andersen's PhD thesis [2] discusses a problem similar to the LSFRP, called the Network Transition Problem (NTP). No mathematical model or formal problem description is provided, so it is difficult to exactly ascertain what the NTP solves in comparison to the LSFRP. However, it is clear that the NTP lacks cost saving activities like SOS opportunities, empty equipment flows and slow steaming.

The primary previous work on the LSFRP in the literature is [9], by the authors, which introduced Linear Temporal Optimization Planning (LTOP), a hybrid of automated planning and linear programming that performs a branch-and-bound search for repositioning solutions. The authors focus on an abstraction of the LSFRP without cargo/equipment flows and SOS parallel sailings. The hybrid method works well on a version of the LSFRP solved due to the state dependent nature of the activities in the LSFRP, and outperforms a MIP model. The LSFRP seems to reside on the border between automated planning and mixed-integer programming. On the one hand, vessels have important state information that must be taken into account. On the other hand, there is not so much state information such that embedding that information in a graph for the MIP explodes as in classical automated planning problems. Nevertheless, LTOP is not capable of handling cargo flows, as the automated planning actions used must either model only costs or only revenues. LTOP is, therefore, unable to handle our model of the LSFRP.

3 Mathematical Model

We model the LSFRP with cargo flows on a graph $G = (V, A)$, where V is the set of nodes and A the set of directed arcs between nodes. Each node in V represents a *visitation* of a vessel at a particular port[2], and each arc in A represents an allowed sailing between two visitations. The graph encompasses all of the activities each vessel may undertake during a fixed *repositioning period*, which is the period from the time the vessel is first allowed to leave its phase-out service until the time when normal operations must begin on the phase-in service. The path of each vessel through the graph represents the activities to be undertaken by that vessel, and we therefore require the paths to be node disjoint to prevent multiple vessels from performing the same activity. This is an important constraint because there is only space for a single vessel at the quay during a visitation. Note that flexible visitations, i.e. visitations without a prior fixed schedule, can be undertaken by multiple vessels, even simultaneously. For ease of modeling, we therefore replicate flexible visitations for each vessel and consider them as node disjoint. We give more details about this process (and justifications) later. We embed a number of problem constraints and objectives directly in the graph, including sailing costs, sail-on-service opportunities, cabotage restrictions, phase-in/out requirements, and canal fees, which are described in detail in the next section, followed by our MIP model over the graph.

3.1 Graph Description

The visitations in the graph are split into two sets, thus $V = V^i \cup V^f$, where V^i is the set of *inflexible* visitations, i.e. visitations associated with a specific port call time, and V^f is the set of *flexible* visitations, which are assigned a time only if a vessel performs the visitation. The set V^f contains visitations in which a vessel can pick up/deliver equipment or incremental cargo that are not on any phase-out, phase-in, or SOS service.

Figure 4 shows a high-level view of the structure of the graph. The graph's visitations are classified into the sets PO, PI, F, and SOS, which contain the phase-out, phase-in, flexible and sail-on-service visitations, respectively. The node τ represents the graph sink. Vessels begin their repositioning in the PO visitations with the goal of

[2] We use the terms visitation and node interchangeably.

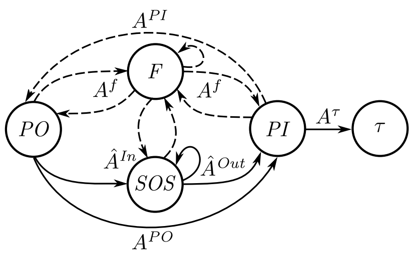

Fig. 4: A high level overview of the graph structure.

reaching τ, which can only be reached through the PI nodes in order to enforce constraints on how services are started. At any time during repositioning, each vessel is either performing a visitation, or sailing between two visitations. The vessel may sail to a visitation in a different set if an arc connects that set from the vessel's current visitation, with dashed lines indicating that the two visitations must be in the same tradezone, i.e. they must be geographically close to each other. We now formally define the graph of the LSFRP, using the following constants.

S	Set of ships, indexed by s.
L	Set of phase-in slots, where $\lvert L\rvert = \lvert S\rvert$, indexed by ℓ.
SOS	The set of SOS slots, indexed by o.
R_ℓ^{PI}	Set of visitations of phase-in slot ℓ.
R_s^{PO}	Set of phase-out visitations of vessel s.
$O_o^{\{P,TS,T,E\}}$	Sets of parallel, transshipment, transit, and end visitations.
V^R	Set of non-SOS inflexible visitations, $V^R = V^i \backslash (O_o^P \cup O_o^{TS} \cup O_o^T \cup O_o^E)$.
TZ	Set of trade zones.
$z_i \in TZ$	Trade zone of visitation $i \in V$.
$enter(i) \in \mathbb{R}^+$	Time a vessel begins inflexible visitation $i \in V^i$.
$exit(i) \in \mathbb{R}^+$	Time a vessel ends inflexible visitation $i \in V^i$.
$\tau \in V$	Graph sink, which is not an actual visitation.
$V' = V \setminus \tau$	Set of nodes without the graph sink.
$\Delta_{i,j}$	Minimum time required for any ship to sail from visitation i to j.
$A^{SD}(R)$	Set of arcs connecting subsequent visitations in the visitation set R.
A^{PO}	Set of arcs connecting phase-out slots to phase-in slots.
A^{PI}	Set of arcs from phase-in visitations to same tradezone phase-out visitations.
A^τ	Set of arcs from the phase-in to the graph sink.
A^f	Set of arcs connecting flexible visitations to other visitations.
$\hat{A}_o^{In}$	Set of arcs connecting to the start nodes of o.
$\hat{A}_o^{Out}$	Set of arcs extending from the end nodes of o.
$\hat{A}_o^{PTS}$	Set of arcs connecting the parallel nodes to transhipment nodes of o.
$\hat{A}_o^{TST}$	Set of arcs connecting transshipment nodes to transit nodes of o.
$\hat{A}_o^{TT}$	Set of arcs between transit nodes of o.
$\hat{A}_o^{EE}$	Set of arcs between sequential end nodes of o.
$\hat{a}_o^{TE}$	Arc from the latest transit node in o to its earliest end node.

We define the set of inflexible nodes as $V^i = \bigcup_{\ell \in L} R_\ell^{PI} \bigcup_{s \in S} R_s^{PO} \bigcup_{o \in SOS} (O_o^P \cup O_o^T \cup O_o^{TS} \cup O_o^E)$. The set of flexible visitations, V^f, contains all visitations that have equipment surpluses/deficits such that $V^f \cap V^i = \emptyset$. In order to formally define the set of arcs contained in the graph, let $follows(i,j) \in \mathbb{B}$ return *true* if and only if visitation j is scheduled on any service to immediately follow visitation i, with $i,j \in V^i$. In addition, we let $can\text{-}sail(i,j) \in \mathbb{B}$ be *true* if and only if $enter(j) \geq exit(i) + \Delta_{i,j}$, where $i,j \in V'$. This indicates whether or not it is possible to sail between two visitations at the fastest speed of the fastest vessel in the model. Note that all of the arc sets are disjoint. We now formally define all of the previously mentioned sets of arcs.

$$
\begin{aligned}
A^{SD}(R) &= \{(i,j) \mid i,j \in R \wedge \mathit{follows}(i,j)\}, R \in \bigcup_{s\in S} R_s^{PO} \bigcup_{\ell\in L} R_\ell^{PI} \\
A^{PO} &= \{(i,j) \mid i \in \bigcup_{s\in S} R_s^{PO} \wedge j \in \bigcup_{\ell\in L} R_\ell^{PI} \wedge \mathit{can\text{-}sail}(i,j)\} \\
A^{PI} &= \{(i,j) \mid i \in \bigcup_{\ell\in L} R_\ell^{PI} \wedge j \in \bigcup_{s\in S} R_s^{PO} \wedge z_i = z_j \wedge \mathit{can\text{-}sail}(i,j)\} \\
A^{\tau} &= \{(i,j) \mid i \in \bigcup_{\ell\in L} \operatorname*{argmax}_{i'\in R_\ell^{PI}}\{\mathit{exit}(i')\} \wedge j = \tau\} \\
A^{f} &= \{(i,j) \mid \big((i \in V^f \wedge j \in V^R) \vee (i \in V^R \wedge j \in V^f)\big) \wedge z_i = z_j\} \\
\hat{A}_o^{In} &= \{(i,j) \mid i \in \bigcup_{s\in S} R_s^{PO} \wedge j \in (O_o^P \cup O_o^{TS}) \wedge \mathit{can\text{-}sail}(i,j)\} \\
&\quad \bigcup \{(i,j) \mid i \in V^f \wedge j \in (O_o^P \cup O_o^{TS}) \wedge z_i = z_j \wedge \mathit{can\text{-}sail}(i,j)\} \\
\hat{A}_o^{Out} &= \{(i,j) \mid i \in O_o^E \wedge j \in \left(\bigcup_{\ell\in L} R_\ell^{PI} \bigcup_{o'\in\{SOS\setminus o\}} (O_{o'}^P \cup O_{o'}^{TS})\right) \wedge \mathit{can\text{-}sail}(i,j)\} \\
&\quad \bigcup \{(i,j) \mid i \in O_o^E \wedge j \in V^f \wedge z_i = z_j \wedge \mathit{can\text{-}sail}(i,j)\} \\
\hat{A}_o^{PTS} &= \{(i,j) \mid i \in O_o^P \wedge j \in O_o^{TS} \wedge \mathit{follows}(i,j)\} \\
\hat{A}_o^{TST} &= \{(i,j) \mid i \in O_o^{TS} \wedge j \in O_o^{T} \wedge \mathit{follows}(i,j)\} \\
\hat{A}_o^{TT} &= \{(i,j) \mid i,j \in O_o^{T} \wedge \mathit{follows}(i,j)\} \\
\hat{A}_o^{EE} &= \{(i,j) \mid i,j \in O_o^{E} \wedge \mathit{follows}(i,j)\} \\
\hat{a}_o^{TE} &= (\operatorname*{argmax}_{i\in O_o^T}\{\mathit{exit}(i)\}, \operatorname*{argmin}_{j\in O_o^E}\{\mathit{enter}(j)\})
\end{aligned}
$$

The set of all arcs in the graph, A, is therefore defined by

$$
\begin{aligned}
A = &\bigcup_{s\in S} \big(A^{SD}(R_s^{PO})\big) \bigcup_{\ell\in L} \big(A^{SD}(R_\ell^{PI})\big) \cup A^{PI} \cup A^f \\
&\cup \hat{A}_o^{In} \cup \hat{A}_o^{Out} \cup A^{\tau} \cup A_o^{ST} \cup \hat{A}_o^{TT} \cup \hat{A}_o^{EE} \cup \hat{a}_o^{TE}.
\end{aligned}
$$

Phase-out and Phase-in Each vessel begins its repositioning at a pre-specified earliest phase-out visitation on its initial slot. The graph structure of the phase-out slots is fully described by $A^{SD}(R), R \in \cup_{s\in S} R_s^{PO}$; A^{PO}; and $\hat{A}_o^{In}, o \in SOS$ since the only activities a vessel may undertake on its initial service slot are sailing between scheduled visitations in the slot, or leaving the slot for a flexible node, SOS opportunity, or the phase-in service.

The phase-in graph structure ensures that the goal service has a vessel in each of its slots. An example phase-in graph structure is given in Figure 5, which shows the graph for the service in Figure 1. Each sequence of visitations (colored light gray, dark gray, and black) represents a slot on the goal service. Each visitation is labeled with the

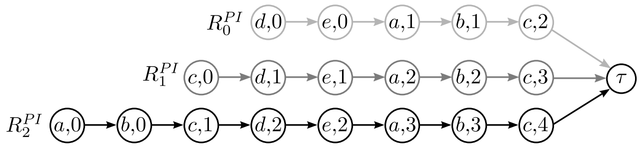

Fig. 5: The phase-in graph structure for the service in Figure 1.

port and week that it is visited. The last node in each sequence corresponds to the on-time requirement (node $(c, 2)$) extended to each slot. After each of these visitations, the service begins normal operations, and is no longer under the control of the repositioning coordinator. This graph structure ensures that all vessels perform a legal phase-in, namely that each slot is assigned a single vessel. Each phase-in slot is guaranteed to be assigned a single vessel since there are as many slots as there are vessels (three), the graph sink τ only has a single incoming node from each slot, and the paths of vessels are node disjoint (except for τ).

Flexible visitations Flexible visitations are modeled by replicating the flexible visitation for each vessel in the model. If we did not replicate flexible visitations, only one vessel would be able to visit a flexible visitation, even though it is possible (and even desirable) for multiple vessels to undertake a single flexible visitation. This is because when a vessel visits a flexible visitation, the visitation must be assigned a time when it can take place. Therefore, we replicate the visitation and the node-disjoint paths of each vessel ensure two vessels do not use the same flexible visitation node. Since our instances generally do not contain many flexible visitations, this duplication does not significantly hinder the solvability of the instances. This opens the possibility that two vessels may visit the same flexible visitation at the same time. We do not consider this to be a problem since flexible visitations are at ports that will probably have the capacity to deal with multiple ships. Since flexible visitations do not have fixed entry and exit times, the time required for a vessel to visit them must be taken into account. The *piloting time* is the time required to maneuver the vessel in to, and out of, a port. The amount of cargo loaded or unloaded from the vessel at the port further extends the amount of time necessary for a vessel to stay at a flexible port, depending on the efficiency of the port at moving containers on and off the vessel.

SOS Figure 6 shows the graph structure of an SOS opportunity. Vessels may enter the SOS either through parallel sailing nodes (O_o^P) or transshipment nodes (O_o^{TS}). Parallel sailings end in a transshipment, shown with graph arcs $\hat{A}_o^{PTS}$, in which cargo is moved to the repositioning vessel. Port p_1 does not allow any transshipments due to a cabotage restriction, necessitating a parallel sailing to carry cargo from the port.

The transshipment nodes then connect to the transit nodes (O_o^T), which represent the "normal" port calls of the SOS. The vessel calls each port of the SOS in sequence until it reaches an end port (O_o^E), where it may then leave the SOS. A single SOS opportunity may only be used by a single vessel, and since the paths of the vessels through the graph are node disjoint, the sequential nature of the end nodes ensures that only one vessel may traverse a particular slot in an SOS.

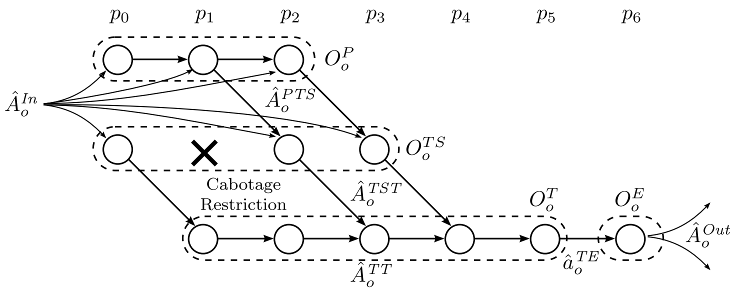

Fig. 6: The graph structure of an SOS opportunity.

3.2 MIP Model

We now define the MIP model that guides the vessels through the graph, and controls the flow of cargo and equipment, using the following constants and variables to supplement the constants used to define the graph.

Constants

T	Set of equipment types. $T = \{dc, rf\}$.
$\sigma_s \in V'$	Starting visitation of vessel $s \in S$.
$V^{t+}, (V^{t-})$	Set of visitations with an equipment surplus (deficit) of type t.
V^{t^*}	Set of visitations with an equipment surplus or deficit of type t.
V^{Goal}	Set of visitations corresponding to visitations on the goal service.
$In(i) \subseteq V'$	Set of visitations with an arc connecting to visitation $i \in V$.
$Out(i) \subseteq V'$	Set of visitations receiving an arc from $i \in V$.
$c_i^{Mv} \in \mathbb{R}^+$	Cost of a TEU move at visitation $i \in V'$.
$f_{s,i}^{Port} \in \mathbb{R}$	Port fee associated with vessel s at visitation $i \in V'$.
$d_{s,i}^{Mv} \in \mathbb{R}$	Move time per TEU for vessel s at visitation $i \in V'$.
$r_t^{Eqp} \in \mathbb{R}^+$	Revenue for each TEU of equipment of type $t \in T$ delivered.
$u_s^t \in \mathbb{R}^+$	Capacity of vessel s for cargo type $t \in T$.
A'	The set of arcs $(i, j) \in A$, where $i, j \in V'$.
$c_{i,j}^s$	Fixed cost of vessel s utilizing arc $(i, j) \in A'$.
$\alpha_{i,j}^s$	Variable hourly cost of vessel $s \in S$ utilizing arc $(i, j) \in A'$.
$\Delta_{i,j,s}^{Min}$	Minimum duration for vessel s to sail on flexible arc (i, j).
$\Delta_{i,j,s}^{Max}$	Maximum duration for vessel s to sail on flexible arc (i, j).
$(o, d, t) \in \Theta$	A demand triplet, where $o \in V', d \subseteq V'$ and $t \in T$ are the origin visitation, destination visitations and the cargo type, respectively.
$a^{(o,d,t)} \in \mathbb{R}^+$	Amount of demand available for the demand triplet.
$r^{(o,d,t)} \in \mathbb{R}^+$	Amount of revenue gained per TEU for the demand triplet.
$\Theta_i^O, (\Theta_i^D) \subseteq \Theta$	Set of demands with an origin (destination) visitation $i \in V$.
Θ_o^{SOS}	Set of demands corresponding to $o \in SOS$.

Variables

$y^s_{i,j} \in \{0,1\}$	Indicates whether vessel s is sailing on arc $(i,j) \in A$.
$t^E_i \in \mathbb{R}^+_0$	Defines the enter time of a vessel at visitation i.
$t^X_i \in \mathbb{R}^+_0$	Defines the exit time of a vessel at visitation i.
$t^V_i \in \mathbb{R}^+_0$	Defines the time spent at flexible visitation $i \in V^f$ loading and unloading cargo/equipment.
$w^s_{i,j} \in \mathbb{R}^+_0$	The duration that vessel $s \in S$ sails on flexible arc $(i,j) \in A^f$.
$x^{(o,d,t)}_{i,j} \in \mathbb{R}^+_0$	Amount of flow of demand triplet $(o,d,t) \in \Theta$ on $(i,j) \in A'$.
$x^t_{i,j} \in \mathbb{R}^+_0$	Amount of equipment of type $t \in T$ flowing on $(i,j) \in A'$.

$$\max \sum_{s \in S} \left(\sum_{(i,j) \in A'} c^s_{i,j} y^s_{i,j} + \sum_{(i,j) \in A^f} \alpha^s_{i,j} w^s_{i,j} \right) \tag{1}$$

$$+ \sum_{(o,d,t) \in \Theta} \left(\sum_{j \in d} \sum_{i \in In(j)} \left(r^{(o,d,t)} - c^{Mv}_o - c^{Mv}_j \right) x^{(o,d,t)}_{i,j} \right) \tag{2}$$

$$+ \sum_{t \in T} \left(\sum_{i \in V^{t+}} \sum_{j \in Out(i)} \left(r^{Eqp}_t - c^{Mv}_i \right) x^t_{i,j} - \sum_{i \in V^{t-}} \sum_{j \in In(i)} c^{Mv}_i x^t_{j,i} \right) \tag{3}$$

$$+ \sum_{j \in V'} \sum_{i \in In(j)} \sum_{s \in S} f^{Port}_{s,j} y^s_{i,j} \tag{4}$$

$$\text{s. t.} \sum_{s \in S} \sum_{i \in In(j)} y^s_{i,j} \le 1 \qquad \forall j \in V' \tag{5}$$

$$\sum_{i \in Out(\sigma_s)} y^s_{\sigma_s,i} = 1 \qquad \forall s \in S \tag{6}$$

$$\sum_{i \in In(\tau)} \sum_{s \in S} y^s_{i,\tau} = |S| \tag{7}$$

$$\sum_{i \in In(j)} y^s_{i,j} - \sum_{i \in Out(j)} y^s_{j,i} = 0 \qquad \forall j \in \{V' \setminus \bigcup_{s \in S} \sigma_s\}, s \in S \tag{8}$$

$$\sum_{(o,d,rf) \in \Theta} x^{(o,d,rf)}_{i,j} - \sum_{s \in S} u^{rf}_k y^s_{i,j} \le 0 \qquad \forall (i,j) \in A' \tag{9}$$

$$\sum_{(o,d,t) \in \Theta} x^{(o,d,t)}_{i,j} + \sum_{t' \in T} x^{t'}_{i,j} - \sum_{s \in S} \left(\sum_{t' \in T} u^{t'}_s \right) y^s_{i,j} \le 0 \quad \forall (i,j) \in A' \tag{10}$$

$$\sum_{i \in Out(o)} x^{(o,d,t)}_{o,i} \le a^{(o,d,t)} \sum_{i \in Out(o)} \sum_{s \in S} y^s_{o,i} \qquad \forall (o,d,t) \in \Theta \tag{11}$$

$$\sum_{i \in In(j)} x^{(o,d,t)}_{i,j} - \sum_{k \in Out(j)} x^{(o,d,t)}_{j,k} = 0 \qquad \forall (o,d,t) \in \Theta, j \in V' \setminus (o \cup d) \tag{12}$$

$$\sum_{i \in In(j)} x^t_{i,j} - \sum_{k \in Out(j)} x^t_{j,k} = 0 \qquad \forall t \in T, j \in V' \setminus V^{t^*} \tag{13}$$

$$\Delta^{Min}_{i,j,s} y^s_{i,j} \le w^s_{i,j} \le \Delta^{Max}_{i,j,s} y^s_{i,j} \qquad \forall (i,j) \in A^f, s \in S \tag{14}$$

$$t_i^E = enter(i) \sum_{s \in S} \sum_{j \in In(i)} y_{i,j}^s \qquad \forall i \in V^i \tag{15}$$

$$t_i^X = exit(i) \sum_{s \in S} \sum_{j \in Out(i)} y_{i,j}^s \qquad \forall i \in V^i \tag{16}$$

$$t_i^X + \sum_{s \in S} w_{i,j}^s - t_j^E \leq 0 \qquad \forall (i,j) \in A^f \tag{17}$$

$$t_i^X - t_i^E - t_i^V \leq -pilot(i) \qquad \forall i \in V^f \tag{18}$$

$$\begin{aligned} &\sum_{(o,d,t) \in \Theta_i^O} \sum_{j \in Out(o)} d_{s,o}^{Mv} x_{o,j}^{(o,d,t)} + \sum_{(o,d,t) \in \Theta_i^D} \sum_{d' \in d} \sum_{j \in In(d')} d_{s,d}^{Mv} x_{j,d'}^{(o,d,t)} \\ &+ \sum_{t \in T} \left(\sum_{i' \in \{V^{t+} \cap i\}} \sum_{j \in Out(i')} d_{s,i'}^{Mv} x_{i',j}^t + \sum_{i' \in \{V^{t-} \cap i\}} \sum_{j \in In(i')} d_{s,j}^{Mv} x_{j,i'}^t \right) \\ &- \ t_i^V + M_i^s \sum_{j \in Out(i)} y_{i,j}^s \leq M_i^s \qquad \forall i \in V^f, s \in S \end{aligned} \tag{19}$$

The objective consists of several components. The sailing cost (1) takes into account the precomputed sailing costs on arcs between inflexible visitations, as well as the variable cost for sailings to and from flexible visitations. Note that the fixed sailing cost on an arc does not only include fuel costs, but can also include canal fees or the time-charter bonus for entering an SOS. The profit from delivering cargo (2) is computed based on the revenue from delivering cargo minus the cost to load and unload the cargo from the vessel. Note that the model can choose how much of a demand to deliver, even choosing to deliver a fractional amount. We can allow this since each demand is an aggregation of cargo between two ports, meaning at most one container between two ports will be fractional. Equipment profit is taken into account in (3). Equipment is handled similar to cargo, except that equipment can flow from any port where it is in supply to any port where it is in demand. Finally, port fees are deducted in (4).

Multiple vessels are prevented from visiting the same visitation in (5). The flow of each vessel from its source node to the graph sink is handled by (6), (7) and (8), with (7) ensuring that all vessels arrive at the sink.

Arcs are assigned capacities if a vessel utilizes the arc in (9), which assigns the reefer container capacity, and in (10), which assigns the total container capacity, respectively. Note that constraints (9) do not take into account empty reefer equipment, since empty containers do not need to be turned on, and can therefore be placed anywhere on the vessel. Cargo is only allowed to flow on arcs with a vessel in (11). The flow of cargo from its source to its destination, through intermediate nodes, is handled by (12). Constraints (13) balance the flow of equipment in to and out of nodes. In contrast to the way cargo is handled, equipment can flow from any port where it is in supply to any port where it is in demand. Since the amount of equipment carried is limited only by the capacity of the vessel, no flow source/sink constraints are required.

Flexible arcs have a duration constrained by the minimum and maximum sailing time of the vessel on the arc in (14). The enter and exit time of a vessel at inflexible ports is handled by (15) and (16), and we note that in practice these constraints are only necessary if one of the outgoing arcs from an inflexible visitation ends at a flexible visitation. Constraints (17) sets the enter time of a visitation to be the duration of a vessel on a flexible arc plus the exit time of the vessel at the start of the arc, and (18) forces the enter and exit time at a flexible node to take into account the time the vessel spends in port, including piloting time. The time to load and unload cargo and equipment is reflected in (19), which sets the t^V variable to the cargo movement time.

The model forms a disjoint path problem in which a fractional multicommodity flow is allowed to flow over arcs in the paths, along with a small scheduling component in the flexible nodes. Flexible arcs could be alternatively represented using a discretized approach, however we forego a discretization because of the vast differences in timescales between port activities and sailing activities, which are on the order of hours and days, respectively. In order to achieve such a fine grained view of flexible arc activities, we would require numerous extra arcs and nodes for each flexible node.

3.3 Complexity

We reduce the knapsack problem to the LSFRP in order to show that the LSFRP is NP-complete. Given n items, each with a profit p_i and a size s_i, and a knapsack with a capacity C, the knapsack problem maximizes the objective $\sum_{i=0}^{n} p_i x_i$ where x_i is a binary variable indicating whether or not item i is in the knapsack, subject to the capacity constraint $\sum_{i=0}^{n} s_i x_i \leq C$.

Theorem 1. *The LSFRP with flexible visitations is NP-complete.*

We first note that the LSFRP is clearly in NP, as the total profit can be easily computed from the paths of the vessels through the graph. We initialize an LSFRP with a single vessel and no cargo or equipment demands. The problem instance contains a single phase-out visitation, ω, and a single phase-in visitation, λ. The port fees at both ω and λ are 0, and we let $enter(\omega) = exit(\omega) = 0$ and $enter(\lambda) = exit(\lambda) = C$. In other words, the timespan in which the repositioning must take place is limited to the capacity of the knapsack. For each knapsack item, we create a flexible visitation, f_i, which has a duration of exactly s_i, i.e. $pilot(f_i) = s_i$. The port fee for visiting f_i is $-p_i$, since the LSFRP maximizes profit (i.e. minimizes fees). All flexible nodes, as well as ω and λ, are in a single tradezone. Therefore, the specification of the LSFRP graph ensures that the phase-out node, ω, connects to all flexible nodes, all flexible nodes connect to each other, and all flexible nodes connect to the phase-in node, λ. The sailing time of the vessel between all nodes in the graph is set to 0.

Item i is included in the knapsack solution if and only if the vessel visits f_i during its repositioning. Since the vessel can only visit a single flexible visitation at a time, the duration of each flexible visitation is fixed to the size of the item it represents, and the phase-in visitation is fixed in time to the size of the knapsack, the capacity constraint of the knapsack must be satisfied. Additionally, according to the objective of the LSFRP, only the flexible visitations corresponding to the maximum profit knapsack items will be chosen. Therefore, the LSFRP with flexible visitations is NP-complete. □

4 Computational Study

We created a benchmark set of instances containing two real world repositioning scenarios, with three and eleven vessels each. The rest of our benchmark set consists of hypothetical repositionings crafted using real liner shipping data to examine the scaling behavior of the MIP. Our instances have graphs varying from as few as 49 nodes and 214 arcs up to 392 nodes and 12,105 arcs. Table 3 shows information about each instance in our dataset and the time it takes to solve each instance to optimality in seconds. The table gives the instance ID, number of nodes, number of inflexible and flexible arcs, number of demands, number of ports with equipment and number of SOS opportunities, along with the CPU time of each instance. Instances were alloted an hour of CPU time on an AMD Opteron 2425 HE processor and allowed 10 GB of RAM each. Instances that exceeded the maximum CPU time of one hour are marked with "CPU" and instances that ran out of memory are indicated with "Mem".

Instances 1 – 9 represent variations of the full case study shown in Figure 3. They are relatively easy to solve, even under the presence of flexible arcs for equipment. However, as the number of vessels grows, so does the size of the graph and the difficulty of the problems. Instances 42 and 43 correspond to another real world scenario, but neither can be solved in a reasonable amount of time using this model alone. In order to solve these problems to optimality, a more advanced technique, such as branch-and-price, will have to be employed.

There is no obvious way to characterize the hardness of instances. For example, instance 28 requires more than seven times the amount of time than instance 31, even though they have comparable graph sizes and almost the same number of demands. Additionally, instances 12 and 13 have similar graph structures, but instance 13 is solved in half the time of instance 12, perhaps because of the extra two SOS opportunities which may allow for better pruning. Overall, instances that do not timeout are solved relatively quickly, the slowest taking only 22 minutes. In fact, one third of our instances can be solved faster than a minute, including instance 30 with 8 vessels and 1220 arcs. On the instances where the MIP times out, not a single one is able to reach a gap of less than 10% with the linear relaxation, indicating the difficulty of such instances.

5 Conclusion

We presented a novel model of an important real-world problem, the LSFRP, using a MIP model and a constraint embedded graph. Our model takes into account all of the key aspects of the LSFRP, including liner shipping service construction constraints, cargo flows, empty equipment repositioning, cabotage restrictions, and sail-on-service opportunities, and maximizes the profit earned during repositioning. We proved the NP-completeness of the LSFRP, studied the performance of our model on two repositioning scenarios from our industrial collaborator, as well as on a number of instances crafted from real-world data.

For future work, we intend to use a branch and price framework to overcome memory issues and solve instances to optimality faster. In addition, we will investigate using heuristics to provide good solutions to the LSFRP in the low amount of time that would be required for a decision support system.

ID	$\|S\|$	$\|V\|$	$\|A^i\|$	$\|A^f\|$	$\|\Theta\|$	$\sum_{t\in T}\|V^{t^*}\|$	$\|SOS\|$	CPU Time (s)
1	3	49	214	0	72	0	1	0.12
2	3	60	227	0	115	0	2	0.2
3	3	61	200	0	81	0	3	0.1
4	3	68	201	0	108	0	2	0.17
5	3	70	351	0	81	0	3	0.18
6	3	84	218	0	146	0	4	0.24
7	3	86	612	0	99	0	3	0.54
8	3	113	864	132	138	6	3	42.05
9	3	113	864	132	138	10	3	50.52
10	4	61	542	0	99	0	0	37.51
11	4	65	542	42	99	6	0	CPU
12	4	77	630	0	134	0	2	311.87
13	4	83	651	0	146	0	4	115.72
14	4	83	651	0	146	25	4	139.34
15	5	74	381	0	251	0	0	0.74
16	5	109	428	0	398	0	5	1.32
17	6	112	1371	0	118	0	0	63.25
18	6	147	1913	0	128	0	4	CPU
19	6	147	1913	0	128	13	4	CPU
20	6	147	1913	0	128	37	4	CPU
21	6	160	1518	0	153	0	9	68.12
22	6	160	1518	0	153	42	9	45.01
23	6	172	1645	162	153	76	9	CPU
24	7	81	533	0	215	0	0	2.18
25	7	81	533	0	215	16	0	2.06
26	7	83	660	0	289	0	0	713.12
27	7	115	592	0	289	0	3	2.86
28	7	118	674	0	372	0	4	960.28
29	7	118	674	0	372	23	4	1345.34
30	8	119	1220	0	117	0	0	29.12
31	8	126	1621	0	337	0	0	127.43
32	8	161	1473	0	166	0	3	60.79
33	8	196	1495	429	168	50	3	CPU
34	9	317	11817	0	407	0	0	CPU
35	9	370	12004	40	777	124	4	Mem
36	9	377	12038	0	929	0	4	CPU
37	9	384	12099	0	1001	120	7	Mem
38	9	386	12038	40	929	132	4	Mem
39	9	392	12105	0	1057	0	7	Mem
40	9	392	12105	0	1057	124	7	Mem
41	10	257	8521	0	568	0	0	CPU
42	11	317	12438	0	781	0	0	CPU
43	11	325	12931	0	884	0	4	Mem

Table 3: Instance information and the time to solve each instance to optimality with CPLEX 12.4 and a timeout of one hour.

6 Acknowledgements

We would like to thank our industrial collaborators Mikkel Muhldorff Sigurd and Shaun Long at Maersk Line for their support and detailed description of the fleet repositioning problem. We would also like to thank David Pisinger and Björg Áskelsdóttir for their insightful comments and suggestions on this work. This research is sponsored in part by the Danish Council for Strategic Research as part of the ENERPLAN research project.

References

1. J.F. Álvarez. Joint routing and deployment of a fleet of container vessels. *Maritime Economics and Logistics*, 11(2):186–208, June 2009.
2. M.W. Andersen. *Service Network Design and Management in Liner Container Shipping Applications*. PhD thesis, Technical University of Denmark, Department of Transport, 2010.
3. M. Christiansen, K. Fagerholt, B. Nygreen, and D. Ronen. Maritime transportation. *Transportation*, 14:189–284, 2007.
4. M. Christiansen, K. Fagerholt, and D. Ronen. Ship routing and scheduling: Status and perspectives. *Transportation Science*, 38(1):1–18, 2004.
5. J. Clausen, A. Larsen, J. Larsen, and N.J. Rezanova. Disruption management in the airline industry–concepts, models and methods. *Computers & Operations Research*, 37(5):809–821, 2010.
6. N. Kohl, A. Larsen, J. Larsen, A. Ross, and S. Tiourine. Airline disruption management–perspectives, experiences and outlook. *Journal of Air Transport Management*, 13(3):149–162, 2007.
7. J.E. Korsvik, K. Fagerholt, and G. Laporte. A large neighbourhood search heuristic for ship routing and scheduling with split loads. *Computers & Operations Research*, 38(2):474 – 483, 2011.
8. B.J. Powell and A.N. Perakis. Fleet deployment optimization for liner shipping: An integer programming model. *Maritime Policy and Management*, 24(2):183–192, Spring 1997.
9. K. Tierney, A.J. Coles, A.I. Coles, C. Kroer, A.M Britt, and R.M. Jensen. Automated planning for liner shipping fleet repositioning. In *Proceedings of the 22nd International Conference on Automated Planning and Scheduling*, 2012. To Appear.
10. United Nations Conference on Trade and Development. Review of maritime transport, 2011.