Maritime Crude Oil Transportation - a split pickup and split delivery problem

F. Hennig [a,*], B. Nygreen [a], M. Christiansen [a], K. Fagerholt [a],
K. C. Furman [b], J. Song [b], G. R. Kocis [c], P. H. Warrick [c]

[a] *Department of Industrial Economics and Technology Management*
Norwegian University of Science and Technology
Alfred Getz veg 3, N-7491 Trondheim, Norway

[b] *Corporate Strategic Research*
ExxonMobil Research and Engineering Company
1545 Route 22 East, Annandale, NJ 08801, USA

[c] *ExxonMobil Research and Engineering Company*
3225 Gallows Rd, Fairfax, VA 22037, USA

Abstract

The maritime oil tanker routing and scheduling problem is known to the literature since before 1950. In the presented problem, oil tankers transport crude oil from supply points to demand locations around the globe. The objective is to find ship routes, load sizes, as well as port arrival and departure times, in a way that minimizes transportation costs. We introduce a path flow model where paths are ship routes. Continuous variables distribute the cargo between the different routes. Multiple products are transported by a heterogeneous fleet of tankers. Pickup and delivery requirements are not paired to cargos beforehand and arbitrary split of amounts is allowed. Small realistic test instances can be solved with route pregeneration for this model. The results indicate possible simplifications and stimulate further research.

Key words: routing, scheduling, maritime transportation, pickup and delivery, split

* Corresponding author: Phone: +4915124238006, Fax: +493034338562
Email addresses: `Frank.Hennig@iot.ntnu.no` (F. Hennig), `Bjorn.Nygreen@iot.ntnu.no` (B. Nygreen), `Marielle.Christiansen@iot.ntnu.no` (M. Christiansen), `Kjetil.Fagerholt@iot.ntnu.no` (K. Fagerholt), `Kevin.c.furman@exxonmobil.com` (K. C. Furman), `Jinhwa.song@exxonmobil.com` (J. Song), `Gary.r.kocis@exxonmobil.com` (G. R.

1 Introduction

Maritime crude oil transportation began in the end of the nineteenth century. Since then the volume of crude oil transported on seaways has steadily increased. The only significant exceptions have been oil crises in 1973 and 1979 with a subsequent decrease in crude oil consumption and production. Today tanker ships transport more than 1.86 billion tons of crude oil across the seas each year (see (Rodrigue et al., 2006)). The primal driving force for crude oil transportation is refinery requirements. Refineries use crude oil to derive various petroleum products. What type and how much of a petroleum product can be produced depends on refinery capabilities and the types of crude oil, so called grades, available. Refinery operations usually require several different crude oil grades to produce their desired product range. Today's dynamic global market for crude oil and refined products demands versatile refinery operations. Refineries have to adapt to changing crude grade availabilities and varying demand of refined products. This changing environment also affects transportation. If refinery requirements or supply options change, transportation has to be adapted.

The crude oil tanker routing and scheduling problem we study, which is similar to the problem of McKay and Hartley (1974), is potentially applicable to worldwide crude oil transportation. In the problem, a heterogeneous oil tanker fleet transports a number of crude oil grades from several loading ports to several discharging ports. Many loading ports supply a single, location specific crude grade. Some ports however supply several crude grades that also can be found in other loading locations. Refineries usually request several crude grades and hence have to be supplied from several loading ports. Pickups and deliveries are requested in specified time windows. While discharging time windows can be based on refinery production and storage plans, loading time windows usually are the result of negotiations with suppliers. Required pickup and delivery amounts can be split in arbitrary portions and be serviced by several tankers. It can be observed that loading as well as discharging ports often conglomerate in certain geographical regions.

Previous research on maritime crude oil tanker routing and scheduling has treated several aspects of the real world problem. Aspects that have been studied include heterogeneous tanker fleets, multiple products, port restrictions that limit access and cargo onboard, physical ship restrictions and time windows. Typically a cargo is perceived as a quantity of freight to be transported between a loading and a discharging port by a single ship on a single trip. Little attention has been paid to cases where the transportation of single cargoes can be shared between ships. Such a problem is usually referred to as

Kocis), Philip.h.warrick@exxonmobil.com (P. H. Warrick).

split problem. In addition, almost no attention has been paid to cases where the typical cargo definition does not apply. If quantities in pickup locations are not dedicated to certain delivery locations, a pairing of pickup and delivery does not exist and thus is part of a solution. We found this non-paired pickup and delivery in only one crude oil related publication. Often tanker voyages have a rather simple structure or are based on a seemingly rigorous subset of possible ship routes. Where time windows are considered, these seem to be tight. The research in the field of oil tanker routing and scheduling applications has undergone a fairly natural development. We refer to the problem as the oil tanker routing and scheduling problem like for example in Sherali et al. (1999).

The purpose of this paper is to present a model for an oil tanker routing and scheduling problem similar to McKay and Hartley (1974) but more realistic with respect to modern crude oil shipping. The model replicates degrees of freedom present in real operations that are scarcely studied and challenging from an algorithmic point of view. Unlike many others, except McKay and Hartley (1974), we model non-paired supply and demand time windows and arbitrary split of supply and demand amounts. In contrast to McKay and Hartley (1974), we fulfill both pickup and delivery requirements. We also provide details on our solution procedure. Computational results are meant to stimulate further research on the topic that may result in the solving of large scale instances.

The paper is organized as follows: In Section 2 we show previous research on the oil tanker routing and scheduling problem. We also mention research conducted on different kinds of split problems. Section 3 gives a description of the problem and in Section 4 we explain the basics of the path flow model presented in Section 5. In Section 6 we explain how paths can be obtained in a pre-generation phase. Different transportation instances are solved by commercial software and presented in Section 7. In Section 8 discussions and conclusions are made.

2 Previous Research

Oil tanker routing and scheduling is a well known task and, as far as the operations research literature is concerned, goes back to before 1950. It almost seems that the problem has undergone a natural evolution in parallel with increasing computational power and algorithmic advancements. For the purpose of describing oil tanker routing and scheduling problems and their solution approaches it seems justifiable to start in 1954 with the US Navy fuel oil tanker routing problem. In the first part of this section we review publications, which treat the oil tanker routing and scheduling problem, in order of their date of publication. Solution approaches and achievements are discussed. The main characteristics that appear in these papers are listed in Table 1 for the purpose of overview. For a comprehensive review on other maritime routing and scheduling problems see (Christiansen et al., 2004). The second part refers to the scarcity of research on pickup and delivery problems with split. We mention some examples and findings in connection with split problems.

Table 1
Main characteristics treated in the reviewed literature

Problem aspects	Characteristics treated in the literature
Fleet types	Homogeneous / heterogeneous
	Sufficiently / insufficiently large
Cargo types	Full / partial shiploads
	Contracted / optional cargoes
	Single / multiple origin(s) and destination(s)
	Splitable cargoes
Cargo carrying	Single / multiple cargo(es) onboard
Ship routing	Single loading to single discharging port
	Single loading port cluster to single discharging port cluster
	Via multiple loading and discharging ports
Restrictions	Bunker fuel consumption
	Port draft restrictions
	Optimal speed selection

2.1 The Oil Tanker Routing and Scheduling Problem

The first problem we present, the US Navy fuel oil tanker routing problem, has received the attention of several researchers. In this problem a homogeneous fleet of tankers is engaged in worldwide fuel oil transportation. Dantzig and Fulkerson (1954), and Flood (1954) treat the problem in a similar manner. They assume a sufficiently large tanker fleet to satisfy the transport demand. The transport demand is given as the number of full shiploads needed between pairs of loading and discharging ports. No scheduling of pickup and delivery dates is necessary. While Dantzig and Fulkerson (1954) are interested in the minimum number of tankers, Flood (1954) minimizes ballast sailing costs. Both problems can be solved by linear programming and as stressed in Dantzig and Fulkerson (1954) as a transportation problem. Later Briskin (1966) points out that the transportation of full shiploads between port pairs is a coarse assumption. He instead proposes discharging port clusters, where the total cargo amount in a cluster is a full shipload. Dynamic programming is used to find routes and indirectly schedules within a discharge cluster. The proposed approach can be combined with the method of Dantzig and Fulkerson (1954) and then allows for a more detailed tanker routing. Finally, an under-sized fleet of tankers is allowed in (Bellmore, 1968). Not all cargoes can be serviced and therefore profit for the transport that can be carried out is maximized. The problem can be formulated as transshipment problem and remains solvable by linear programming.

A shipping problem that is not explicitly linked to oil transportation but in its characteristics probably directly applicable to it is described by Appelgren (1969, 1971). Appelgren (1969) considers a heterogeneous fleet of tankers, where ships have different sizes, speeds and costs. Cargoes are specified by amount, cargo type, loading time window and discharging time window. Each ship carries only one cargo at a time. Whereas different cargo types could in principle be handled in (Dantzig and Fulkerson, 1954), and (Flood, 1954), specific cargo amounts and loading time windows are new. In addition the fleet is allowed to service additional spot cargoes. For solving the problem three solution approaches are discussed: A multi-commodity flow formulation, a path flow formulation with pre-generated routes and a column generation approach. The column generation approach is favored but only the linear relaxation of the master problem is solved to optimality. Feasible solutions were often found. The largest instance that was solved consists of 40 ships, 50 cargoes and a planning horizon of two to three months. In (Appelgren, 1971) the problem of fractional solutions is studied. The paper considers cutting planes and a branch-and-bound method to find feasible, non-fractional solutions. The branch-and-bound method with column generation in the root node proved to be very successful.

Another formulation of the problem is given by Bellmore et al. (1971). The problem is quite similar to the one described by Appelgren (1969) but does not consider spot cargoes. Tankers can be partially loaded and share cargoes. A tanker will only carry one, or part of one, cargo at a time. The authors suggest a column generation approach like Appelgren (1971), but only describe and discuss the branch-and-bound procedure.

The first paper that challenges the assumption of predefined cargoes (or port pairs) is (McKay and Hartley, 1974). The paper assumes independent, non-paired pickup and delivery requirements. Moreover, multiple products can be handled in the model. The authors give an integer programming formulation which they, due to its complexity, reformulate into a model that uses predefined routes. Even though they only use routes with 1-2 loading ports and up to 3 discharging ports, typical problem sizes of their practical applications prove to be too difficult to solve. The authors therefore resort to solving their problem approximately based on a linear relaxation.

Instead of following the challenges shown by McKay and Hartley (1974), Brown et al. (1987) relate to previous problem characteristics, for example full shiploads, and treat spot chartered vessels and optimal speed selection. Spot vessels transport cargoes which cannot be shipped by the controlled fleet. Solutions are obtained after routes are pre-generated and an integer programming formulation is solved.

A further study of similar kind which is also the continuation of Brown et al. (1987) is described by Bausch et al. (1991). They propose a so called elastic set partitioning model. Routes are generated beforehand and the optimal routes are chosen in a set partitioning manner. The specialty here is that set partitioning constraints can be violated at a penalty. The main focus of this article is to show the good applicability of their model in practice.

The next ones who actually extend the tanker routing and scheduling problem are Bremer and Perakis (1992) and Perakis and Bremer (1992). They consider several additional details such as tanker fuel, so called bunker oil, draft restrictions and spot charter costs. The authors consider scheduling explicitly. Their routes however have a rather simple structure as they consider only one loading and one discharging port. Again all possible routes can be pre-generated.

A study where problem size plays a major role is illustrated by Sherali et al. (1999). The described problem goes back on a doctoral thesis, (Al-Yakoob, 1997). Sherali et al. (1999) consider crude oil transportation from Kuwait to North America, Europe and Japan. Also here voyages are simple in structure. In this study the actual assignment of cargoes to compartments in the vessels has been more important. Split delivery and late deliveries are allowed. The problem was finally aggregated and solved based on a rolling horizon approach.

The last article still close to the considered problem is (Chajakis, 2000). In this paper the author mentions a study where routing and scheduling is seen as part of a greater supply chain. Unfortunately no model is presented. The author correctly points out that transport operations cannot be separated completely from refining and storage.

In spite of the growing importance of crude oil in the world economy, research around the oil tanker routing and scheduling problem has not increased during the recent years. With respect to the realistic flexibility in crude oil availability and demand it is unfortunate, that almost no applied research has been conducted on splitting of cargoes. Our research can be seen as an extension of McKay and Hartley (1974), who allow arbitrary split. In addition to them we require a certain amount of crude oil to be transported. Pickup time windows exist and more cargo restrictions are considered. We allow more stops in a route and allow routes to be a combination of laden voyages connected by ballast sailings.

2.2 *Studies on Split Problems*

The pickup and delivery problem (PDP) has been extensively studied in many variants, see for example reviews of Parragh et al. (2008a,b). A commented review can be found in (Berbeglia et al., 2007). Variants, that treat split of transport requirements and non-paired pickup and delivery nodes are scarce. Parragh et al. (2008b) for example names only one study, in which pickup and delivery points are non-paired for the multi-vehicle case. Pickup and delivery problems with split for both pickup and delivery are not mentioned in the review at all. The only PDP with multiple vehicles and allowed split in both pickup and delivery nodes known to the authors is McKay and Hartley (1974).

The problem class that comes closest to the studied problem is the pickup and delivery problem with split loads (PDPSL), which can be found in Nowak et al. (2008). In this problem pre-defined loads, which have a specific origin and destination, can be split between several vehicles. The authors find that load sizes just over one half vehicle capacity have greatest benefit from splitting.

Another problem type that has many similarities is called inventory routing. Here inventories at pickup and/or delivery nodes have to be kept within limits. Usually shipment sizes are not predefined and pickup and delivery nodes might not be paired. This results in a certain form of split of cargo amounts. Examples can be found in Christiansen (1999) and Persson and Göthe-Lundgren (2005).

Most attention with respect to split has been paid to vehicle routing problems (VRP) with either split pickup or split delivery. Some of the most recent publications are for example Archetti et al. (2008), Flisberg et al. (2009) and Chen et al. (2007). Archetti et al. (2008) come to very similar conclusions as Nowak et al. (2008). Requirements of one half to three quarter vehicle capacity are most significant for splitting and a reduction of the number of routes can be found. The actual location of delivery points does not seem to be of importance. A rich practical application for split pickup is given in Flisberg et al. (2009) and for split delivery problems in Chen et al. (2007).

To our knowledge no problem class has been introduced for the pickup and delivery problem with unpaired pickups and deliveries, and split in all nodes. Without split Parragh et al. (2008b) suggests to name the problem class pickup and delivery VRP, or PDVRP. With pairing of pickup and deliveries Nowak et al. (2008) calls the problem PDPSL. In contrast to the PDVRP, we do not have depots, and in contrast to the PDPSL we do not have paired pickups and deliveries. In addition we have to deal with time windows. Therefore the presented problem could be termed a split pickup split delivery problem with time windows (SPSDPTW).

3 Definition of the Oil Tanker Routing and Scheduling Problem

3.1 Supply and Demand

During a typical planning period of one to several months, refineries (also referred to as discharging ports) place crude oil requests, which often have to be satisfied from different supply locations. A discharging port may request several different crude grades, e.g. grades A and B, at different times (see Figure 1). A request is a specific crude grade and volume in a given time window. Like in Figure 1 time windows can be timely separated or overlapping. Loading ports on the other hand usually supply a unique crude grade. Some ports however can supply several grades, which may also be found in other supply ports. The arcs in Figure 1 suggest a possible pairing of loading and discharging time windows. Note that whenever there are several loading and discharging time windows with the same grade no predefined pairing between the time windows exists. Loading or discharging requests can be split between several tankers.

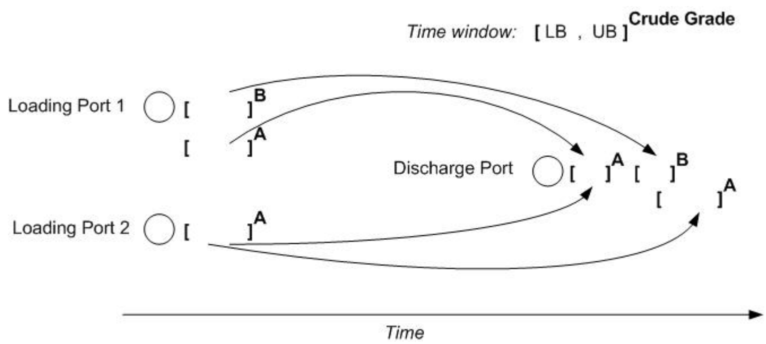

Fig. 1. A possible pairing of supply and demand

The length of a planning period depends on the planning situation. Two situations are common:

- Long-term planning based on rough demand estimates well in advance of plan execution,
- short term planning based on recently updated demand and supply information.

The first case usually involves fewer, wide time windows with larger quantities whereas the second case is based on more and tighter time windows with smaller quantities. Actual planning problems can have characteristics of the long term problem, the short term problem, or some combination of both.

3.2 Tanker Routes

During the planning period tankers are employed on voyages. A voyage is defined as a sequence of ports, more precisely of time window port visits, between which a tanker is laden. Ports are often concentrated in separate geographic loading and discharging regions. Distances between such regions are very long, compared with distances within regions. Hence, all loadings in a voyage are carried out first and are followed by all discharges. Before first loading and after last discharging the vessel is empty and sails in ballast. A voyage can have several loading and several discharging ports. (A single voyage with three loading and two discharging ports is illustrated in Figure 2.) Note that a tanker may load or discharge several grades during one port visit. During the planning period a tanker can possibly carry out several voyages. We call the entire sailing in service of a tanker a route (a detailed specification of routes can be found in Section 4).

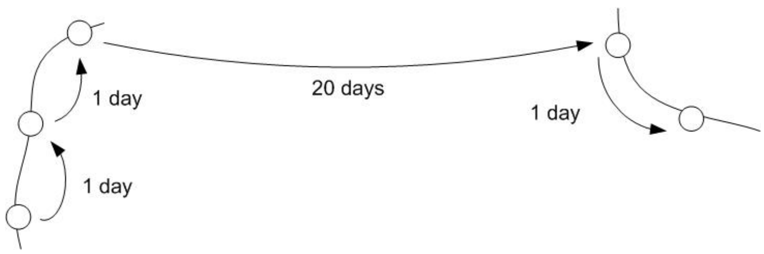

Fig. 2. Voyage with three loading and two discharging ports

3.3 Tankers

A very typical tanker in the considered trade is the so called Very Large Crude Carrier (VLCC). These vessels have an approximate capacity of around 300 000 tons or roughly 2.1 million barrels of crude oil. Another often used tanker, the so called SUEZMAX tanker, has approximately half the size of a VLCC. Smaller tankers are also occasionally used, but less common for long-distance transport. The compartments of a tanker allow it to carry crude oil of different grades simultaneously. For planning purposes it can be assumed that the order of loading different grades and the specific use of compartments is not important. Since a crude oil tanker normally transports only a small number of grades at the same time, it is in practice possible to load almost any mix of crude amounts in their differently sized compartments. Tanker fleets are heterogeneous, since each tanker is different in capacity, speed, dimensions and cost structure. In addition, their initial positions usually differ in location and time of availability.

3.4 Ports and Restrictions

Ports impose several different restrictions on a tanker. Restrictions can originate from several operational and regulatory necessities but can for planning purposes be translated to maximum crude oil weight and volume onboard a vessel when it enters or leaves a port. Restrictions can apply for both incoming and outgoing sailings. Some ports might not at all be suitable for a certain vessel. In addition to port restrictions a tanker has a cargo weight and volume capacity. In different conditions, one or both of these capacities can be limiting. While the real volume capacity only depends on the vessel's tank volume, cargo weight capacity is influenced by the amount of operational supplies onboard. The most dominant variable supply is bunker fuel. Increasing the amount of bunker fuel onboard reduces the amount of cargo that can be transported.

Fuel oil can be bunkered in many locations and for different operational durations. For long voyages, the amount of bunker fuel can be considerable. We address this issue by reducing the cargo weight capacity of a tanker on its sailing leg between a loading region and a discharging region. On this leg a tanker will have its maximum load. The length of an inter-region leg also indicates the amount of bunker fuel needed on the entire voyage. Hence, we base the capacity reduction, or bunker fuel shut-out, on the length of inter-region legs.

Berth constraints that limit the number of simultaneous tanker visits could be an issue. Due to flexibility in practice it seems to be acceptable to exclude this from large-scale crude oil tanker transportation planning.

3.5 Transportation costs

The variable cost of transportation depends on two main components: vessel fuel oil costs and port fees. We do not consider any fixed costs, like manning expenses or charter costs, because we assume a fixed fleet for the transportation task. All fixed costs for the fleet are constant for the planning period and are not subject to optimization. The largest part of a tanker's variable cost on a route is determined by its fuel oil consumption. A tanker burns different fuel amounts per day while sailing, port operations or when waiting. It uses most fuel oil when sailing and least when waiting. The other cost component, port fees, has to be paid whenever a tanker enters a port. While port fees and sailing costs are determined by the actual routing choice, port operation costs and waiting costs are time dependent. The more a tanker loads or discharges in a port, the more costly is the operation. In the same way longer waiting times result in higher cost.

4 Modeling tanker routes

As described in Section 3.2 a tanker route is a sequence of port time windows.

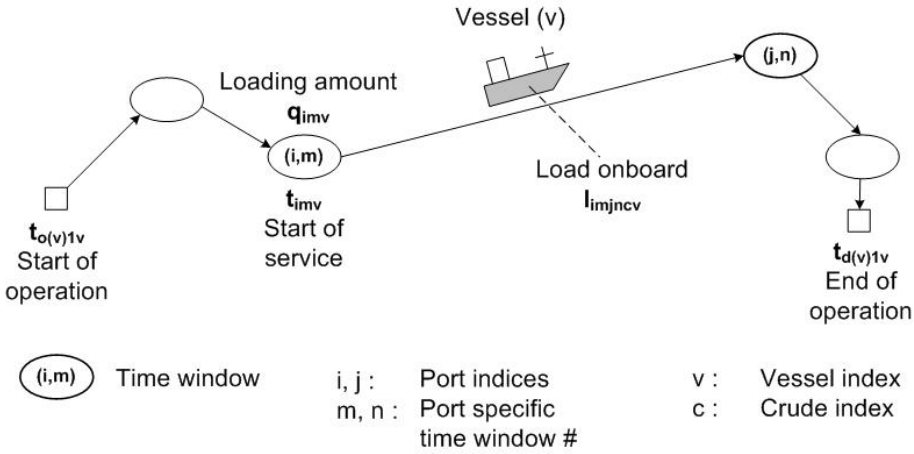

Fig. 3. Single-voyage tanker route

Figure 3 shows a tanker route where nodes are visited time windows. The illustrated route has two loading followed by two discharging time windows. Squared nodes represent vessel origin and destination positions. The route shown consists of one voyage only. A time window in a route is specified by port name, indexed i or j, and time window number, indexed m or n. In the following we will refer to port time window pairs (i, m) and (j, n) as time windows. Two or more consecutive time windows can belong to the same port.

In each time window a vessel, v, loads or discharges a certain amount of crude oil, q_{imv}. Technical circumstances can suggest minimum loading amounts if loading first takes place. The crude grade c, which is loaded or discharged in a time window, is time window specific and therefore not in the index subscript of q_{imv}. Different loading time windows in a voyage may supply different crude grades. Therefore, the load l_{imjncv} onboard a vessel v between two time windows (i, m) and (j, n) has to be tracked for each crude grade. Time window lower and upper bounds limit the time for start of service t_{imv} for vessel v in time window (i, m). Arrival at the port is allowed to be earlier than start of service. An early arrival results in idle/waiting time. In addition, waiting time may be constrained. Time windows in vessel origin and destination locations can limit the time a vessel is available for operation.

The most common models applied in large ship routing and scheduling applications are path flow models, where paths represent ship routes for which visited ports and transported cargo amounts are known. The optimization then has to select one route per vessel, so that all constraints are fulfilled. Examples for that can be found in Appelgren (1969), Bausch et al. (1991) and Perakis and Bremer (1992). Paths can also be mere sequences of time windows like in McKay and Hartley (1974). Tanker loads and schedules then have to be decided by the optimization model.

5 A Path Flow Formulation with continuous cargo quantities

In this paper we present a path flow model, in which paths are ship routes consisting of sequences of time windows. No information about loading or discharging amounts are related to a route. As shown in Figure 3, the model needs continuous variables, q_{imv}, t_{imv} and l_{imjncv}, to distribute cargo between the several ships and to ensure that time and cargo constraints are fulfilled. Binary variables λ_{vr} take the value one, if ship v uses route r and zero otherwise. Each ship sails one route only.

Actually used (port-time-window to port-time-window) sailing legs can be retrieved from the formulation by means of the following formula, which right-hand side appears several times in the path flow formulation:

$$x_{imjnv} = \sum_{r \in \mathcal{R}_v} A_{imjnvr} \cdot \lambda_{vr}.$$

For a given sailing leg (i, m, j, n) from time window (i, m) to time window (j, n), A_{imjnvr} equals one, if vessel v uses sailing leg (i, m, j, n) on route r and zero otherwise. With $\mathcal{R}_v$ as the set of all routes for vessel v, binary sailing leg variable x_{imjnv} equals one, if sailing leg (i, m, j, n) is included in a route actually sailed by vessel v. Otherwise x_{imjnv} is zero.

In the following section we give the mathematical description of the path flow model. The generation of paths is explained in Section 6.

5.1 Model

The model combines tanker routes - one route per tanker - and decides on loading and discharging quantities to find a cost minimal routing plan and sailing schedule. Each part of the model is explained separately. We introduce the needed nomenclature for each model part at the beginning of each subsection.

5.1.1 Objective Function

The objective of the model is to minimize total transportation cost, which has two components: Bunker fuel costs and port fees.

Indices:

i, j	Ports
m, n	Time window numbers
v	Vessel
r	Route
$o(v)$	Origin position of vessel v
$d(v)$	Destination position of vessel v

Sets:

$\mathcal{V}$	Vessels
$\mathcal{R}_v$	Routes for vessel v
$\mathcal{N}_v$	Ports that vessel v can visit
$\mathcal{T}_i$	Time window numbers for port i

Data:

C_{vr}	Fixed part of cost for sailing route r by vessel v
T_{im}^{Q}	Loading/discharging time needed per weight unit crude oil in time window (i, m)
F_v^P	Reduced fuel cost per time unit of port operation for vessel v
F_v^I	Fuel cost per time unit of idle time for vessel v

Variables

λ_{vr}	Binary routing variable; takes value 1, if vessel v sails route r and 0 otherwise
q_{imv}	Cargo weight loaded or discharged in time window (i, m) by vessel v
t_{imv}	Time vessel v starts service in time window (i, m)

$$
\begin{aligned}
min \sum_{v \in \mathcal{V}} \sum_{r \in \mathcal{R}_v} & C_{vr} \cdot \lambda_{vr} \\
+ \sum_{v \in \mathcal{V}} \sum_{i \in \mathcal{N}_v} \sum_{m \in \mathcal{T}_i} & F_v^P \cdot T_{im}^Q \cdot q_{imv} \\
+ \sum_{v \in \mathcal{V}} & F_v^I \cdot (t_{d(v)1v} - t_{o(v)1v}).
\end{aligned} \tag{1}
$$

The objective of the model is to minimize fuel costs and port fees. Port fees apply whenever a vessel visits a port. Fuel costs arise per day of vessel operation. The first term incurs the cost of sailing, C_{vr}, for vessel v on an entire route r. C_{vr} also includes port fees for the entire route. The second term covers the variable part of the costs in port. The fuel consumption in port depends on the amount of handled cargo. The last term accounts for the waiting time, or idle, fuel consumption. Instead of tracking how much waiting time a vessel spends on a route, we reduce sailing and port fuel consumption by the amount of idle fuel consumption and charge idle fuel consumption for the entire time a vessel is in service.

5.1.2 Convexity Constraints

$$\sum_{r\in\mathcal{R}_v} \lambda_{vr} = 1 \qquad \forall v \in \mathcal{V}. \tag{2}$$

Each vessel is allowed to sail one route only. If a vessel is not used in the optimal solution it sails a dummy route from its origin directly to its destination at no cost.

5.1.3 Scheduling Constraints

Scheduling constraints are necessary, because the exact time needed in port is unknown in the route generation phase. Pre-generated routes might be feasible with respect to sailing time, but can become infeasible for certain port stay durations. The time spent in port is first known, when handled cargo amounts are decided. The scheduling of a route depends therefore on handled cargo amounts. In the same way, scheduling constraints can limit handling amounts.

Sets:

$\mathcal{A}_v$	Arcs (i, m, j, n) vessel v can sail (including arcs from origin time window and to destination time window)

Data:

T^S_{ijv}	Sailing time between ports i and j for vessel v
A_{imjnvr}	=1, if vessel v sails leg (i, m, j, n) on route r; =0 otherwise
$\underline{T}_{im}$	Earliest time for start of service in time window (i, m)
$\overline{T}_{im}$	Latest time for start of service in time window (i, m)
U_{imjnv}	Sailing leg and vessel specific big-M constant for unused sailing legs

$$t_{imv} + T^Q_{im} \cdot q_{imv} + T^S_{ijv} - t_{jnv} - U_{imjnv} \cdot (1 - \sum_{r \in \mathcal{R}_v} A_{imjnvr} \cdot \lambda_{vr}) \leq 0$$
$$\forall v \in \mathcal{V}, (i, m, j, n) \in \mathcal{A}_v, \tag{3}$$

$$\underline{T}_{im} \leq t_{imv} \leq \overline{T}_{im} \qquad \forall v \in \mathcal{V}, i \in \mathcal{N}_v \cup \{o(v), d(v)\}, m \in \mathcal{T}_i, \tag{4}$$

A lower bound on the start of service in time window (j, n) is calculated in (3) by adding to the time for start of service, t_{imv}, in time window (i, m) the cargo handling time, $T^Q_{im} \cdot q_{imv}$, in (i, m) and sailing time, T^S_{ijv}, from i to j. In addition, we have to make sure in constraint (4) that the times for start of service in each port lie within specified time windows. The values of the t_{imv} variables in non-visited time windows can be chosen freely by the optimization and do not have any meaning. We define one time variable per vessel for each time window. As a consequence, multiple visits by a single vessel in a time window are not feasible. Christiansen (1999) presents a shipping model, where repeated time window visits are feasible. The same approach could be adopted here, too. However, we deem this as unneccessary, since we assume that a repeated time window visit is not meaningful in the context of the considered application. (Further discussion of this issue can be found in Section 6.1.)

5.1.4 Cargo Constraints

The cargo constraints make sure that supply and demand requirements are met. They also keep track of the crude grade specific cargo amounts onboard the vessels. Note that cargo restrictions have to be obeyed on each sailing leg and that supply and demand time windows are not explicitly linked.

Sets:

$\mathcal{N}^P$	Loading ports
$\mathcal{N}^D$	Discharging ports
$\mathcal{V}^N_i$	Vessels that can visit port i
$\mathcal{A}^W_v$	Arcs for vessel v that possess a possibly binding cargo weight restriction
$\mathcal{A}^V_v$	Arcs for vessel v that possess a possibly binding cargo volume restriction
$\mathcal{C}$	Crude grades in the problem

Data:

Q_{im}	Cargo amount to be loaded/discharged in total by all vessels in time window (i, m)
C_{im}	Crude grade supplied or demanded in port i and time window m
D_c	Density of crude grade c
I_i	Sign modifier; =1, if port i is a loading port and =-1, if port i is a discharging port
$\delta_{c,C_{im}}$	Kronecker delta; =1 for $c = C_{im}$ and 0 otherwise
$\overline{W}_{ijv}$	Maximum allowed cargo weight for sailings from port i to j by vessel v
$\overline{V}_{ijv}$	Maximum allowed cargo volume for sailings from port i to j by vessel v

Variable:

l_{imjncv} Load of crude grade c onboard vessel v on leg (i, m, j, n)

$$\sum_{c \in \mathcal{C}} l_{imjncv} - \overline{W}_{ijv} \cdot \sum_{r \in \mathcal{R}_v} A_{imjnvr} \cdot \lambda_{vr} \leq 0$$
$$\forall v \in \mathcal{V}, (i, m, j, n) \in \mathcal{A}_v^W, \quad (5)$$

$$\sum_{c \in \mathcal{C}} \frac{l_{imjncv}}{D_c} - \overline{V}_{ijv} \cdot \sum_{r \in \mathcal{R}_v} A_{imjnvr} \cdot \lambda_{vr} \leq 0$$
$$\forall v \in \mathcal{V}, (i, m, j, n) \in \mathcal{A}_v^V, \quad (6)$$

$$\sum_{(j,n)} l_{jnimcv} + \delta_{c,C_{im}} \cdot I_i \cdot q_{imv} - \sum_{(j,n)} l_{imjncv} = 0$$
$$\forall v \in \mathcal{V}, i \in \mathcal{N}_v, m \in \mathcal{T}_i, c \in \mathcal{C}. \quad (7)$$

Constraints (5) and (6) only allow cargo onboard a vessel on used sailing legs. At least one of these constraints will exist for each sailing leg. The total cargo amount onboard a vessel on a specific leg is the sum of the amounts for each crude grade onboard. This load amount has to be less than or equal to a weight or volume limit, $\overline{W}_{ijv}$ or $\overline{V}_{ijv}$. The cargo volume can easily be calculated by dividing each grade's load amount, l_{imjncv}, by its density D_c. We can check in advance if a leg (i, m, j, n) for a vessel v might be weight restrictive, volume restrictive or both. Constraint (7) is an inventory balance constraint for cargo amounts onboard the vessels. Each time window supplies or demands a unique crude grade C_{im}. Load amounts onboard need to be changed for only this grade. For all other grades load amounts remain unchanged. For a particular vessel, there will be only one used incoming and one used outgoing sailing leg for each visited time window. Only on these legs are load variables, l_{imjncv}, allowed to be positive. Hence, the correct load variables will be updated. The constraints are needed for each vessel and each grade that could be onboard, when the vessel visits the time window.

$$\sum_{v \in \mathcal{V}_i^N} q_{imv} = Q_{im} \qquad \forall i \in \mathcal{N}^P \cup \mathcal{N}^D, m \in \mathcal{T}_i, \tag{8}$$

In each time window (i, m) loading or discharging requirements Q_{im} have to be met. This amount can be larger or smaller than a single ship's capacity. Since we allow an arbitrary split of all requirements, the total requirement amount Q_{im} equals the sum of all loadings or dischargings over all vessels in constraint (8).

5.1.5 Variable Type Constraints

The variable type constraints complete the model. The only specialty here are the semi-continuous loading variables in (10).

Data:

$\underline{P}_{imv}$ Minimum loading amount in time window (i, m) for vessel v

$\overline{P}_{imv}$ Maximum loading amount in time window (i, m) for vessel v

$$\begin{aligned}
\lambda_{vr} &\in \{0, 1\} && \forall v \in \mathcal{V}, r \in \mathcal{R}_v, && (9)\\
q_{imv} &\in \{0, [\underline{P}_{imv}, \overline{P}_{imv}]\} && \forall i \in \mathcal{N}^P, v \in \mathcal{V}_i^N, m \in \mathcal{T}_i, && (10)\\
q_{imv} &\geq 0 && \forall i \in \mathcal{N}^D, v \in \mathcal{V}_i^N, m \in \mathcal{T}_i && (11)\\
t_{imv} &\geq 0 && \forall v \in \mathcal{V}, i \in \mathcal{N}_v \cup \{o(v), d(v)\}, m \in \mathcal{T}_i && (12)\\
l_{imjncv} &\geq 0 && \forall v \in \mathcal{V}, (i, m, j, n) \in \mathcal{A}_v, c \in \mathcal{C}. && (13)
\end{aligned}$$

If a vessel v does not visit time window (i, m), variable q_{imv} has to be zero. If however loading takes place, the vessel loads at least a minimum amount $\underline{P}_{imv}$. This amount is time window and vessel specific and is based on technical and business issues. A vessel can at maximum load its own capacity, possibly reduced by other restrictions, or the maximum available amount. Discharge amounts have no lower bound. Constraint (12) is in principle unnecessary due to constraint (4).

6 Route generation

For the given model, we pre-generate tanker routes. Only those routes need to be generated, which have relevance in reality. The following sections describe, which assumptions are used and how routes are pre-generated.

6.1 Problem specific assumptions

Routes are based on an incomplete network of sailing legs. Sailing legs considered unrealistic by the problem owner are excluded from the network. In addition, unrealistic port sequences can be excluded during the route generation. The following assumptions limit the number of possible voyages from which routes are built.

A realistic voyage has a limited number of port visits and will in practice not include long waiting times. The risk of interruptions is usually kept on a reasonable level if the number of loading and discharging ports in a voyage does not exceed three each for planning purposes. Since a tanker might service several time windows in a port, the number of time windows in a voyage can exceed the number of ports. We allow at maximum four time windows each for loading and discharging to keep the number of different grades onboard on a reasonable level and to allow several time window servicings per port visit. The total waiting time in a voyage, which cannot be avoided due to given time window bounds, should be limited and is assumed to be approximately one quarter of the entire sailing time at maximum. Allowed waiting time for a ship between the ship being ready to service a time window and actual service is shorter. The reason for that is that many days of planned waiting at a port are not perceived good solutions in practice. That is valid even if, from a theoretical planning point of view, such waiting would be beneficial.

The question arises, if a vessel is allowed to visit a time window several times during a voyage or a route. In realistic operations a vessel would not interrupt loading/discharging of a grade to load/discharge another grade in the same or even a different port. Even if a data specification is imaginable, for which such an interruption could be part of an optimal solution, we assume that these occurrences are rare and exclude them from the optimization. At the same time, this provides solutions more acceptable in practice. In different consecutive voyages of the same route, a repeated visit of a time window would only be feasible for extremely wide time windows, present in rough long-term planning (see Section 3.1). We turn our attention to short-term planning where extremely wide time windows are not an issue. (One way to approach the long-term planning with the proposed model would be to split wide time windows in several (for example two), for which revisit is infeasible.) At the present state this seems to be acceptable.

In this study we assume a heterogeneous fleet of tankers, which is sufficiently large to carry out the entire transportation. We only consider tankers of VLCC size.

6.2 Route Generator

The route generator pre-generates all feasible routes for a given problem instance. It can also make a selection, if all feasible routes lead to a prohibitive large model. In principle four steps are carried out for each vessel:

(1) For each subset of discharging time windows, generate all time feasible delivery sequences (delivery routings).
(2) For each delivery routing: Consider all subsets of grade matching loading time windows and, for each subset, generate all feasible pickup sequences (pickup routings). Each matching pair of pickup and delivery routings constitutes a voyage.
(3) Select a subset of promising voyages.
(4) Generate single-voyage routes and combine voyages into routes with multiple, consecutive voyages.

Steps (1) and (2) constitute the voyage generation phase. Voyages consist of a pickup route followed by a delivery route. Since there may be many routing options for the pickup and delivery routings, several different voyages with identical time windows exist. To keep all voyages in the model is impractical for problem solving as the results in Section 7.2 indicate. Many voyages are very similar and some voyages may be dominated by others. Dominated voyages can never be part of an optimal solution and as such could be omitted. However, dominance is not trivial to check in the problem at hand. We therefore choose a voyage selection strategy which is derived from a simplified and approximate dominance consideration. The idea is as follows: Generally we are interested in cheap voyages. However, the cheapest, which typically means shortest, route to connect ports may not allow us to transport a maximum of cargo due to draft restrictions. If, for example, the first discharging port in a route can only be visited by a partially laden ship, we miss the chance to utilize our ships fully. Therefore we must be interested in finding relatively cheap voyages that have a high potential for capacity utilization. How much capacity utilization in fact will be required on a voyage is unknown a priori. In step (3) we have chosen to select a number of good voyages per group of voyages with identical time windows. The description and analysis of the voyage selection can be found in Section 7.2.

In step (4) voyages are converted to routes. Each voyage has a vessel specific time interval for feasible start of voyage. The start interval is the time frame in which a ship has to start the voyage in order to arrive at each time window before it closes. At the same time it must not arrive so early that it would have to wait longer than the set waiting time limit. At ship origin position we assume that no waiting takes place. Therefore the start time window of a ship can be interpreted as end of voyage zero in a route. Now we can combine voyages into routes. If it is not possible for a ship to leave its origin time window and arrive within the start intervals of the following voyages, the route consisting of these voyages is infeasible. A feasible single-voyage route consists of origin, a single voyage and destination. A multiple-voyage route comprises origin, two or more consecutive voyages and destination. We generate all possible combinations of voyages into single or multiple-voyage routes in the following way:

Definitions:

- k-base route: origin follows by k single voyages,
- k-voyage route: k-base route followed by ship destination.

For each ship:

(1) Set k = 1.
(2) Generate all time feasible 1-base routes.
(3) Generate all k-voyage routes: Extend all k-base routes with the ship's destination, if feasible.
(4) Consider all k-base routes: Build all time feasible pairs of k-base routes and single voyages to conceive all k+1 base routes.
(5) If there is at least one k+1-base route set k = k+1 and continue with step 3.

Step (3) is the actual route finalization step.

7 Cases and Computational Results

In this section we first present six realistic test instances. Then we describe the voyage selection strategy mentioned in Section 6.2 step (3) and finally report computational results. All computations are carried out on a HP DL140 G3 with two 64-bit dual core processors (1.6 GHz, 8 GB RAM) and Linux operating system. The optimization software Xpress-MP 2008A is used.

7.1 Test instances

We consider two types of test instances. The planning horizon of instances 1 to 3 only allows single-voyage routes. In instances 4 to 6 routes can consist of up to two voyages. A test instance can be characterized by the number of ships in the fleet, number of loading and discharging time windows and the total crude oil amount (in thousand barrels) to be transported. Table 2 gives on overview about the instance sizes.

Table 2
Test instances

Instances	1	2	3	4	5	6
# ships	2	3	5	2	4	6
# loading time windows	2	3	6	5	6	6
# discharging time windows	5	8	14	8	12	16
total crude amount (kbbl)	3770	6270	10270	6950	10550	12120

7.2 Selection of voyages

To choose reasonably good voyages for the route generation, we focus on sailing cost and potential for capacity utilization. Voyages can visit the same time windows but in different order. The goal is to select a number of voyages for each group of voyages with identical time windows. In general we prefer cheap voyages, but we have to ensure that a ship is able to carry a sufficiently large load on the heaviest loaded sailing leg. In other words, we want to ensure that the maximum possible capacity utilization of a ship on the heaviest loaded leg is greater or equal to a reasonable value.

Voyages are selected according to the strategy outlines below. A numerical example with six possible voyages is shown in Table **??**:

(1) Choose one or several ship capacity utilizations and call them $U_1, U_2, ...$ for example 75% (U_1) and 90% (U_2) utilization. ($U_i < U_{i+1} < U_{i+2}...$)
(2) Sort the voyages that visit the same time windows after increasing sailing cost. Consider only voyages during sorting that allow ship utilization at least equal to U_1. (Disregard all voyages that have both a higher sailing cost and a lower maximum possible utilization than any other voyage.)
(3) Select the first (cheapest) voyage into the list of voyages to be used in the optimization. Consider the in (1) chosen utilization values and call the one nearest above the maximum possible ship utilization of the just selected voyage U_k.

(4) Delete those of the sorted voyages that do not allow ship utilization at least equal to U_k.
(5) If there is at least one more sorted voyage, then continue from 3 until there are no more voyages left.

Table 3
Voyage selection example

1) $\mathbf{U_1 = 75\%, U_2 = 90\%, U_3 = 97\%}$														
2)			**3)**			**4)**			**5-3)**			**5-4)**		
No.	% cost	% max. util.	No.	% cost	% max. util.	No.	% cost	% max. util.	No.	% cost	% max. util.	No.	% cost	% max. util.
1	80	80	1	80	80									
2	85	82	2	85	82	~~2~~	~~85~~	~~82~~						
3	88	88	3	88	88	~~3~~	~~88~~	~~88~~						
4	93	92	4	93	92	4	93	92	4	93	92			
5	98	95	5	98	95	5	98	95	5	98	95	~~5~~	~~98~~	~~95~~
6	100	98	6	100	98	6	100	98	6	100	98	6	100	98
		$\geq$ **75**	$\mathbf{U_k = U_2}$					$\geq$ **90**	$\mathbf{U_k = U_3}$					$\geq$ **97**

If the utilizations $U_1, U_2, \ldots$ are set sensibly, voyages are contained in the selection that allow large cargo amounts and at the same time are cheapest possible. In case of a cargo situation that does not require only heavily loaded vessels cheaper voyages with less cargo potential are also available. To find a good choice of utilizations we have tested four settings on two instances, instances 2 and 5. The settings and their utilizations are listed in Table 4.

Table 4
Different utilization settings

Setting	No. of utilizations	Utilizations U
0	–	0
1	1	50
2	1	75
3	1	90
4	2	$75, 90$

Setting 0 is a reference calculation and contains all possible voyages. No selection is made here. No voyage among all possible voyages allows a ship utilization of less than 50%. That means, in setting 1 the cheapest voyage is chosen for each group of voyages with identical time windows. Setting 4 is the only setting where we use two utilization levels to enrich the route pool as compared to setting 2 and 3.

Table 5 shows calculation results for test instance 2 and different ship capacity utilizations U. In setting 1 to 3 cheapest voyages are chosen that allow at least

Table 5
Voyage selection results for instance 2

Setting	0	1	2	3	4
ship capacity utilizations U	–	50	75	90	$75, 90$
# routes	7267	1037	1037	882	1100
LP solution at root node	3494	3593	3593	3578	3570
Best solution found	4078	4096	4096	4078	4078
Total running time (s)	43200	4120	4120	1419	4749
% gap at running time end	10.1	0	0	0	0
# explored nodes	744900	503390	503390	201747	563986

50%, 75%, or 90% ship utilization. For settings 1 to 4 optimal solutions can be found for the particular selections of voyages. Setting 1 and 2 are in fact identical, since in this test instance no voyage limits ship utilization to less than 75%. In terms of running time and solution value setting 3 is clearly to be favored. Setting 0 does not find a better solution within twelve hours. It can be however dangerous to only allow voyages with a fairly high maximum utilization. For different test instances cheaper solutions might be lost, if not all voyages require high utilization. This is the case for instance 5, for which results are shown in Table 6.

Table 6
Voyage selection results for instance 5

Setting	0	1	2	3	4
ship capacity utilizations U	–	50	75	90	$75, 90$
# routes	19391	2220	1769	811	1943
LP solution at root node	5760	–	5989	6054	5982
Best solution found	7205	inf	7399	7814	7428
Total running time (s)	43200	–	43200	17838	43200
% gap at running time end	18.24	–	14.27	0	16.19
# explored nodes	205500	–	767100	1463768	598600

Table 6 illustrates the danger of relying on high utilization or cheapest voyages only. Setting 3 can be solved, but the optimal solution is about 9% more costly than the best found solution in instance 0. Setting 4 does not lead to a proven optimum for instance 5, but its best solution is much closer to the one in instance 0. If only the cheapest voyages are chosen, the problem becomes infeasible because not all cargo can be transported. In order not to sacrifice too much solution quality, but ensure high possible vessel utilization, we use setting 4 for further calculations.

7.3 Calculation results

We report optimal results or best known results within twenty-four hours running time for all cases based on voyage selection setting 4.

Table 7
Test results for setting 4

Instances	1	2	3	4	5	6
# routes	40	1100	12635	520	1943	22843
Route pre-generation time (s)	0	4	1422	0	20	1038
# of variables before presolve	138	1445	16526	928	3623	26257
# of variables after presolve	93	1324	15312	779	2972	25182
LP solution in root node	2620	3570	5792	3967	5982	6299
Best solution found	2638	4078	6777	4810	7399	7687
# ships used	2	3	5	2	4	4
% gap at best solution found	11.19	6.80	14.36	3.59	15.67	17.97
Time to best solution (s)	0	1173	40715	44	52108	68558
% gap at run time end	0	0	14.26	0	15.10	17.95
Total run time (s)	0	4762	86400	45	86400	86400
# explored nodes	1	563986	236300	4679	1085200	140200
Solution at 6 hours	–	–	6860	–	7428	8530
% gap at 6 hours	–	–	15.46	–	17.13	26.11

Only the smallest instances can be solved to optimality within twenty-four hours. Instance 3, 5 and 6 finish with rather large gaps. Only the solutions to instance 5 and 6 still improve after twelve hours. Out of twenty-six voyages in all final solutions twenty-one voyages have only one or two loading time windows. Four loading time windows do not occur. The majority of discharging sequences have three or four time windows. For the instances solved to optimality, four discharging time windows occur only once.

A closer look at the vessel itineraries found by the routing and scheduling optimization reveals different reasons for quantity splitting:

- Loading time window quantity exceeds ship capacity,
- Loading time window quantity has to be delivered on several voyages,
- Splitting takes place without technical necessity.

The first two types represent splitting that must be expected. But whenever there is as choice between loading time windows that supply the same grade, how the splitting takes place is unknown. The third form of splitting, a splitting of amounts to find a better solution, can be observed, too. Here two types can be identified. One, which we could call the split load case, is present, if a single pickup time window supplies a complete delivery, but does that on more than one voyage. The other one, the mixed split case, means that a delivery time window quantity is supplied from different pickup time windows. The mixed split case is a result of the non-paired pickup and delivery time windows. Table 8 shows the occurrences of the split cases in the instances. Note that in a mixed split case the delivery time window quantity needs not to be split.

Table 8
Occurrences of split in the solutions

Occurrences	1	2	3	4	5	6
# of split pickup time windows	1	2	3	1	2	5
# of split delivery time windows	1	1	4	1	2	4
# of split load cases	1	0	1	0	1	2
# of mixed split cases	0	1	4	1	2	2

Each instance has a single pair of time windows, with a unique grade. (In Instance 3 there are two pairs). For these pairs split does not take place. For those discharging time windows, for which an implicit pairing exists, i.e. it is obvious which loading time window is going to supply the demanded crude, split in a discharging time window occurs only three times. Where no implicit pairing is given, discharging time window split happens ten times. The maximum number of loading time windows that supply a single discharging time window is three.

8 Discussion and Conclusion

In this paper we have described an oil tanker routing and scheduling problem, which is based on realistic transport operations. We have formulated the problem as a path flow model with continuous loading and discharging variables. Paths represent ship routes and are pre-generated before the optimization.

During route generation many practically reasonable assumptions can be incorporated and need not be modeled explicitly. The optimization can split cargo amounts specified in time windows on different vessels in an arbitrary fashion. This model is an extension of the model proposed by McKay and Hartley (1974) and has the additional advantage compared to (McKay and Hartley, 1974) that the number of constraints does not grow with the number of generated routes.

The structure of the used test instances naturally influences the obtained results. But since these test instances are based on realistic data interesting observations can be made. In all used test instances, loading time windows can be assigned to discharging time windows in a way that all discharging time windows are supplied by only one loading time window. In other words, no mixed split is necessary. If such a pairing would be given a priori the problem would reduce to a PDPSL with time windows. As the results for the optimally solved instances show, other forms of split are beneficial, too. Extensive mixed split takes place for the explicitly non-paired time windows. Here a greater flexibility exists and obviously is exploited. These results point to the benefits of the proposed model compared to a PDPSL with time windows where all time windows are paired beforehand. A general conclusion on the benefit of mixed split should not drawn based on the results. The proposed model finds feasible solutions relatively quickly but only small instances can be solved to proven optimality.

Further studies on the presented problem can go along two lines. It could be tested, if an a priori pairing for implicitly paired time windows help the solution process. In addition split of these ”cargoes” could be prohibited. The final solution might not worsen significantly. Another line of research can focus on solving larger instances by applying a branch-and-price framework. In connection with that it can also be studied if it is beneficial to replace the continuous loading and discharging variables in the model with discretized cargo amounts. These amounts can be incorporated into the routes. The model then gets a simpler structure at the expense of more, and more complicated routes.

References

Al-Yakoob, S. M., 1997. Mixed-integer mathematical programming optimization models and algorithms for an oil tanker routing and scheduling problem. Ph.D. thesis, Virginia Polytechnic Institute.

Appelgren, L. H., 1969. A column generation algorithm for a ship scheduling problem. Transportation Science 3, 53–68.

Appelgren, L. H., 1971. Integer programming methods for a vessel scheduling problem. Transportation Science 5, 64–78.

Archetti, C., Savelsbergh, M. W. P., Speranza, M. G., 2008. To split or not to split: That is the question. Transportation Research Part E 44, 114–123.

Bausch, D. O., Brown, G. G., Ronen, D., 1991. Elastic set partitioning: A powerful tool for scheduling transportation of oil and gas. In: Breton, M., Zaccour, G. (Eds.), Advances in Operations Research in the Oil and Gas Industry. Paris, Editions Technip.

Bellmore, M., 1968. A maximum utility solution to a vehicle constrained tanker scheduling problem. Naval Research Logistics Quarterly 15(3), 403–411.

Bellmore, M., Bennington, G., Lubore, S., 1971. A multivehicle tanker scheduling problem. Transportation Science 5(1), 36–47.

Berbeglia, G., Cordeau, J., Gribkovskaia, I., Laporte, G., 2007. Static pickup and delivery problems: a classification scheme and survey. TOP 15, 1–31.

Bremer, W. M., Perakis, A. N., 1992. An operational tanker scheduling optimization system: model implementation, results and possible extensions. Maritime Policy and Management 19(3), 189–199.

Briskin, L. E., 1966. Selecting delivery dates in the tanker scheduling problem. Management Science 12(6), 224–235.

Brown, G. G., Graves, G. W., Ronen, D., 1987. Scheduling ocean transportation of crude oil. Management Science 33(3), 335–346.

Chajakis, E. D., 2000. Management science for marine petroleum logistics. Decision Making: Recent Developments and Applications, 279–292.

Chen, S., Golden, B., Wasil, E., 2007. The split delivery vehicle routing problem: Applications, algorithms, test problems, and computational results. Networks 49 (4), 318–329.

Christiansen, M., 1999. Decomposition of a combined inventory and time constrained ship routing problem. Transportation Science 33(1), 3–16.

Christiansen, M., Fagerholt, K., Ronen, D., 2004. Ship routing and scheduling: status and perspectives. Transportation Science 38(1), 1–18.

Dantzig, G. B., Fulkerson, D. R., 1954. Minimizing the number of tankers to meet a fixed schedule. Naval Research Logistics Quarterly 1, 217–222.

Flisberg, P., Lidén, B., Rönnqvist, M., 2009. A hybrid method based on linear programming and tabu search for routing of logging trucks. Computers & Operations Research 36(4), 1122–1144.

Flood, M. M., 1954. Application of transportation theory to scheduling a military tanker fleet. Journal of the Operations Research Society of America 2(2), 150–162.

McKay, M. D., Hartley, H. O., 1974. Computerized scheduling of seagoing tankers. Naval Research Logistics 21(2), 255–264.

Nowak, M., Ergun, Ö., White, C. C., 2008. Pickup and delivery with split loads. Transportation Science 42 (1), 32–43.

Parragh, S., Doerner, K., Hartl, R., 2008a. A survey on pickup and delivery problems - part I: Transportation between customers and depot. Journal für Betriebswirtschaft 58 (1), 21–51.

Parragh, S., Doerner, K., Hartl, R., 2008b. A survey on pickup and delivery problems - part II: Transportation between pickup and delivery locations.

Journal für Betriebswirtschaft 58 (2), 81–117.
Perakis, A. N., Bremer, W. M., 1992. An operational tanker scheduling optimization system: background, current practice and model formulation. Maritime Policy and Management 19(3), 177–187.
Persson, J. A., Göthe-Lundgren, M., 2005. Shipment planning at oil refineries using column generation and valid inequalities. European Journal of Operational Research 163, 631–652.
Rodrigue, J.-P., Comtois, C., Slack, B., 2006. The Geography of Transport Systems. Hofstra University, Department of Economics & Geography, http://people.hofstra.edu/geotrans.
Sherali, H. D., Al-Yakoob, S. M., Hassan, M. M., 1999. Fleet management models and algorithms for an oil-tanker routing and scheduling problem. IIE Transactions 31(5), 395–406.