

# Heuristics for Dynamic and Stochastic Routing in Industrial Shipping

Gregorio Tirado<sup>1</sup>

Lars Magnus Hvattum<sup>2</sup>  
Jean-François Cordeau<sup>4</sup>

Kjetil Fagerholt<sup>2,3</sup>

<sup>1</sup>Universidad Complutense de Madrid, Spain

<sup>2</sup>Department of Industrial Economics and Technology Management,  
Norwegian University of Science and Technology, Norway

<sup>3</sup>Norwegian Marine Technology Research Institute (MARINTEK), Norway

<sup>4</sup>Canada Research Chair in Logistics and Transportation, HEC Montréal, Canada

July 5, 2011

## Abstract

Maritime transportation plays a central role in international trade, being responsible for the majority of long-distance shipments in terms of volume. One of the key aspects in the planning of maritime transportation systems is the routing of ships. While static and deterministic vehicle routing problems have been extensively studied in the last decades and can now be solved effectively with metaheuristics, many industrial applications are both dynamic and stochastic. In this spirit, this paper addresses a dynamic and stochastic maritime transportation problem arising in industrial shipping. Three heuristics adapted to this problem are considered and their performance in minimizing transportation costs is assessed. Extensive computational experiments show that the use of stochastic information within the proposed solution methods yields average cost savings of 2.5% on a set of realistic test instances.

Keywords: Maritime transportation, simulation, tabu search

## 1 Introduction

Maritime transportation is an important component of international trade, and more than 7 billion tons of goods are carried by ship annually (UNCTAD, 2009). The costs related to ships can be very high, with daily time-charter rates and fuel costs that can amount to tens of thousands of USDs. Proper planning of routes and schedules is crucial for shipping companies to achieve a good fleet utilization and reduce costs. The interest in research on ship routing and

scheduling has therefore quickly increased in the last decades as can be seen from the reviews of Ronen (1983, 1993) and Christiansen et al. (2004, 2007).

One usually distinguishes between three different modes of maritime transportation: liner, tramp, and industrial (Lawrence, 1972). In the latter, the ship operator owns or controls both the cargo to be transported and the ships performing the transportation. The focus of the operator is therefore to minimize the total transportation costs while ensuring that all cargoes are transported. A cargo consists of a specified quantity of a given product that must be picked up at its port of loading, transported, and then delivered to its port of discharge. There are specified time windows during which the loading of the cargoes must start, and there may also be time windows for discharging. The operator controls a heterogeneous fleet of vessels, with different sailing speeds, cost structures, and loading capacities. There are two ways of transporting cargoes: the operator may either transport a cargo using a ship from its own controlled fleet, or may hire the services of other shipping companies. In the latter case the operator will use spot carriers, which implies that a lump sum is paid for the transportation of a given cargo.

The planning of ship routes and schedules in industrial shipping, as well as in the other transportation modes, is to an increasing degree performed with assistance from optimization-based decision support systems, as illustrated by Fagerholt (2004), Fagerholt and Lindstad (2007) and Kang et al. (2011). In practice, this often involves using heuristic solution methods to solve deterministic optimization problems based only on known information. This is done on a continuous basis as new relevant information, such as the occurrence of a new cargo, arrives, or at given time intervals. Most algorithms for ship routing and scheduling thus solve static and deterministic versions of the problem. In land-based transportation, previous work on dynamic and stochastic routing has indicated that the inclusion of stochastic information within a dynamic planning process is valuable (see the work of Bent and Van Hentenryck (2004) and Hvattum et al. (2006, 2007)). The purpose of this paper is to develop and test alternative solution methods for dynamic and stochastic routing problems arising in industrial shipping. Considering that maritime transportation constitutes a major element of global freight transportation and the high costs related to it, there is a strong economic incentive to adapt existing methodologies to the maritime case. We also aim at characterizing the types of industrial ship routing and scheduling problems that can gain the most from the use of stochastic information inside optimization algorithms.

In this paper, a discrete event simulation is used to reproduce a planning environment in which new cargo requests arrive over time. Whenever a replanning action is scheduled, heuristics are run to produce solutions consistent with currently available information. Three different heuristics are considered. In the first heuristic, called the myopic dynamic heuristic (MDH), a deterministic subproblem is solved using an adaptation of the tabu search presented by Korsvik et al. (2010). This method does not utilize any stochastic information. The second heuristic is based on the multiple scenario approach with consensus (MSAC) of Bent and Van Hentenryck (2004) and relies on generating a set of scenarios that consist of the currently known cargoes along with a sampled realization of future cargoes. Each of these scenarios is solved using tabu search, and one of the solutions produced is selected based on a consensus function: the chosen solution is the one that is the most similar to other solutions. The third heuristic that

we propose is an adaption of the branch-and-regret heuristic (BRH) of Hvattum et al. (2007). Again, scenarios are generated and solved using tabu search. Rather than selecting one of the resulting solutions, however, an iterative procedure is followed in which parts of the solution are gradually fixed in each of the scenarios until all scenarios have converged to the same solution. We test the three heuristics on a number of realistic, but randomly generated test instances.

The contribution of this paper consists in studying, for the first time, the dynamic and stochastic aspects of tactical routing in a maritime setting. Although the basic ideas of the solution methods applied are known from literature, some adaptations are required for the maritime transportation context. In particular, for the BRH a completely new branching scheme was implemented. Extensive computational experiments show that including stochastic information when solving the routing problem can lead to large savings for the shipping company. However, the size of the savings depends on the attributes of the problem instance, such as whether partial loads or only full loads are allowed, and on the cost of using spot carriers. We conclude that for full load instances the savings are marginal, yet statistically significant, while for partial load instances the savings are quite large. Both BRH and MSAC produce better solutions than MDH, but the structure of the solutions obtained are very different.

The remainder of the paper is organized as follows. A literature review is presented in Section 2. Section 3 gives a description of the industrial ship routing and scheduling problem, including a mathematical formulation of the static and deterministic version of the problem. In Section 4 we present the three heuristics, as well as the simulation procedure and a tabu search that is used in all three heuristics. A computational study is presented in Section 5, while concluding remarks are given in Section 6.

## 2 Review of the literature

This section presents an overview of related work which is divided into three parts. We first survey work on dynamic routing problems, with a particular focus on problems that have a pick-up and delivery structure similar to that of the industrial ship routing and scheduling problem. Second, we discuss previous work on stochastic and dynamic routing problems. Third, we review the literature on maritime routing problems explicitly incorporating stochasticity.

As explained by Psaraftis (1995), a problem is dynamic if some of the problem inputs are not known beforehand but are revealed as time goes by. In these conditions, the decision maker must solve the problem in a sequential fashion by gradually adapting the solution to the new information. To this end, two main strategies can be used. One is to solve a new problem from scratch every time new information becomes available. Another one is to solve a deterministic problem with the information known at the beginning of the planning horizon and to repeatedly update this solution as the unknown information is revealed.

Dynamic vehicle routing problems have received significant attention in the literature as witnessed by the surveys of Psaraftis (1988, 1995) and Ghiani et al. (2003). An important variant of dynamic vehicle routing is the *dynamic pickup and delivery problem* (DPDP) which arises, for example, in the management of urban courier services. In the DPDP, items must be transported between specific origins and destinations by multiple capacitated vehicles which must also respect time windows on pickup and delivery times. Some requests for transportation

are known at the beginning of the horizon but others arrive dynamically as transportation is being performed. One of the first heuristics for the multi-vehicle DPDP was introduced by Savelsbergh and Sol (1998). The algorithm considers a rolling horizon and decomposes the dynamic problem into a series of static problems which are solved by a branch-and-price based heuristic. More recently, Mitrović-Minić and Laporte (2004) have also developed a heuristic based on a rolling horizon framework for the uncapacitated DPDP with time windows. Again, the problem is solved as a series of static problems, and each problem is optimized with tabu search. The approach was then extended by Mitrović-Minić et al. (2004) to consider a double horizon that concurrently optimizes a short-term objective of minimizing traveled distance and a long term objective of maximizing the slack available to introduce new requests. Another heuristic combining tabu search with an adaptive memory was later described by Gendreau et al. (2006). Dynamic pickup and delivery problems have also been extensively studied in the context of passenger transportation where they are commonly referred to as *dial-a-ride problems* (see for example Madsen et al. (1995)). For a recent survey on dynamic pickup and delivery problems, we refer the reader to Berbeglia et al. (2010).

When probabilistic information concerning the unknown problem inputs is available, one faces a stochastic optimization problem. Stochastic routing problems are most often treated as static, rather than dynamic, problems (see the survey of Gendreau et al. (1996)). Nevertheless, the literature on stochastic and dynamic vehicle routing problems is quickly growing. The vehicle routing problem with stochastic demands is often treated as a static problem, but one variant exists where the problem is modeled as a Markov decision process. The most recent contribution is by Secomandi and Margot (2009) who presented partial reoptimization heuristics for the case with a single vehicle. Thomas (2007) also modeled a dynamic and stochastic routing problem as a Markov decision process, considering a single uncapacitated vehicle. The authors analyzed the problem and was able to characterize the optimal policy for the case of a single dynamic customer. Based on this, a heuristic was developed. Swihart and Papastavrou (1999) examined a single-vehicle pickup and delivery problem with the goal of minimizing the expected time in the system for the demands, looking at the problem both from a routing and a queueing perspective. Powell (1996) formulated a stochastic and dynamic model applied to truckload dispatching, where the important decision is which load to take by each truck at the current moment in time.

Ichoua et al. (2006) studied a dynamic and stochastic vehicle routing problem with soft time windows, where the objective function is first to minimize the number of unserved customers and second to minimize the sum of travel time and lateness. They incorporated information from probability distributions by selectively adding dummy customers representing forecast requests to the vehicle routes. Bent and Van Hentenryck (2004) considered a dynamic and stochastic vehicle routing problem with hard time windows, where the only goal is to minimize the number of unserved customers. They proposed to generate multiple scenarios, each one containing a different set of forecast requests. Each scenario is solved separately, and one of them is selected on the basis of a consensus function, yielding a distinguished routing plan. Later, Van Hentenryck et al. (2010) proposed calculating a regret function to determine which customer to visit next, instead of relying on a consensus function to select the solution of a chosen scenario. The idea of representing information about future events using scenarios was also pursued by Hvattum

et al. (2006, 2007), who developed two different heuristics. The problem considered was similar to that of Bent and Van Hentenryck (2004), but the objective function included minimization of the number of vehicles used and travel distance. Hvattum et al. (2007) extended the problem by allowing the demand to be unknown for customers that had placed an order in advance. Schilde et al. (2011) considered variants of the methods of Bent and Van Hentenryck, as well as two variants of variable neighborhood search, for a *dynamic stochastic dial-a-ride problem*. Flatberg et al. (2007) considered issues arising when trying to apply heuristics for stochastic and dynamic vehicle routing problems in practice, including how to model probabilities for future events based on historical data.

Most literature on ship routing and scheduling involves solving static and deterministic problem variants. Two counterexamples that consider operational routing between two points can be mentioned: Lo and McCord (1998) minimized the expected fuel consumption by exploiting uncertain ocean currents and Azaron and Kianfar (2003) represented weather conditions in a stochastic dynamic network. Considering decisions on the tactical level, Cheng and Duran (2004) developed a decision support system for a crude oil transportation and inventory problem that takes into account demand and travel time uncertainty. The crude oil transportation and inventory problem is formulated as a discrete time Markov decision process and solved heuristically. As opposed to the work presented here, the demand uncertainty is with respect to the rate of consumption rather than the appearance of new cargo requests. As only one production port is considered, the problem does not have the same pick-up and delivery structure as in our problem, and only a small set of possible routes are considered from the production port to consumption ports. Hwang et al. (2008) considered the risk associated with fluctuating spot rates and sought to maximize the profit while constraining the variance. The problem is relevant in both tramp and industrial shipping, as optional cargoes are included. Two different formulations were presented, and one of them was solved using a branch-and-price-and-cut method. Only full-load instances were considered, meaning that only one cargo could be onboard a ship at any time, and the problem was considered as a static problem. Uncertainty was associated to future prices and revenues for voyage charter opportunities.

### 3 Problem Description

In the static and deterministic industrial ship routing and scheduling problem, the objective is to minimize the total variable costs for transporting all cargoes by using either the fleet of the company or spot carriers. Each cargo has given origin and destination ports, and the loading and unloading must take place within specified time windows. A ship may carry several cargoes simultaneously, but its total load cannot exceed the ship capacity.

Relying on previous work by Christiansen et al. (2007), the static problem can be formulated on a directed graph  $G = (N, A)$ , where  $N$  is the set of nodes and  $A$  is the set of arcs. Let  $n$  be the total number of cargoes to be transported. Each cargo  $i$  is associated with a loading node  $i$  and an unloading node  $n + i$ . The set of loading nodes is denoted by  $N^P = \{1, \dots, n\}$ , while  $N^D = \{n + 1, \dots, 2n\}$  designates the set of unloading nodes. Let  $V$  denote the set of all ships in the fleet of the shipping company. Each ship  $v \in V$  has a node  $o(v)$  representing its origin and an artificial node  $d(v)$  representing its final destination, which will be determined in the solution

process and corresponds to the last delivery port for ship  $v$ . Let  $N_v \subseteq N^P \cup N^D \cup \{o(v), d(v)\}$  be the subset of nodes that are compatible with ship  $v$  in the sense that the ship can visit the node and carry the cargo associated with the node, if any. Furthermore, define  $N_v^P = N^P \cap N_v$  and  $N_v^D = N^D \cap N_v$ . Let  $A_v \subset N_v \times N_v$  be the set of arcs  $(i, j)$  such that the ship can sail directly from the port associated with node  $i$  to the port associated with node  $j$ . Sailing from node  $i$  to node  $j$  using ship  $v$  incurs a variable cost  $C_{ijv}$  and a total sailing time  $T_{ijv}$ . These values may include, respectively, port fees and handling time at node  $j$ , if any. If cargo  $i$  is transported using a spot carrier, the associated cost is  $C_i^S$ . The capacity of ship  $v$  is given by  $L_v$ , and the size of cargo  $i$  is given by  $Q_i$ . The time window of node  $i \in N$ , if visited by ship  $v$ , is given by  $[B_{iv}, E_{iv}]$ .

To formulate a mathematical program, we define binary flow variables  $x_{ijv}$  equal to 1 iff ship  $v$  sails directly from node  $i$  to node  $j$  and continuous time variables  $t_{iv}$  representing the time at which service begins at node  $i$  if served by ship  $v$ . Similarly, continuous variables  $l_{iv}$  indicate the total load onboard ship  $v$  after completing service at node  $i$ . Finally, binary variables  $y_i$  take value 1 iff cargo  $i$  is performed by a spot carrier. The problem can then be formulated as follows.

$$\min \sum_{v \in V} \sum_{(i,j) \in A_v} C_{ijv} x_{ijv} + \sum_{i \in N^P} C_i^S y_i \quad (1)$$

$$\sum_{v \in V} \sum_{j \in N_v} x_{ijv} + y_i = 1, \quad i \in N^P, \quad (2)$$

$$\sum_{j \in N_v^P \cup \{d(v)\}} x_{o(v)jv} = 1, \quad v \in V, \quad (3)$$

$$\sum_{i \in N_v} x_{ijv} - \sum_{i \in N_v} x_{jiv} = 0, \quad v \in V, j \in N_v^P \cup N_v^D, \quad (4)$$

$$\sum_{i \in N_v^D \cup \{o(v)\}} x_{id(v)v} = 1, \quad v \in V, \quad (5)$$

$$t_{iv} + T_{ijv} - t_{jv} \leq M_{ijv}(1 - x_{ijv}), \quad v \in V, (i, j) \in A_v, \quad (6)$$

$$B_{iv} \leq t_{iv} \leq E_{iv}, \quad v \in V, i \in N_v, \quad (7)$$

$$l_{o(v)v} = 0, \quad v \in V, \quad (8)$$

$$l_{iv} + Q_j - l_{jv} \leq L_v(1 - x_{ijv}), \quad v \in V, j \in N_v^P, (i, j) \in A_v, \quad (9)$$

$$l_{iv} - Q_j - l_{n+j,v} \leq L_v(1 - x_{i,n+j,v}), \quad v \in V, j \in N_v^P, (i, n+j) \in A_v, \quad (10)$$

$$\sum_{j \in N_v} Q_i x_{ijv} \leq l_{iv} \leq \sum_{j \in N_v} L_v x_{ijv}, \quad v \in V, i \in N_v^P, \quad (11)$$

$$\sum_{j \in N_v} (L_v - Q_i) x_{n+i,jv} \geq l_{n+i,v} \geq 0, \quad v \in V, i \in N_v^P, \quad (12)$$

$$\sum_{j \in N_v} x_{ijv} - \sum_{j \in N_v} x_{j,n+i,v} = 0, \quad v \in V, i \in N_v^P, \quad (13)$$

$$t_{iv} + T_{i,n+i,v} - t_{n+i,v} \leq 0, \quad v \in V, i \in N_v^P, \quad (14)$$

$$t_{iv} \geq 0, \quad v \in V, i \in N_v, \quad (15)$$

$$l_{iv} \geq 0, \quad v \in V, i \in N_v, \quad (16)$$

$$y_i \in \{0, 1\}, \quad i \in N^P, \quad (17)$$

$$x_{ijv} \in \{0, 1\}, \quad v \in V, (i, j) \in A_v. \quad (18)$$

The objective function (2) sums all variable costs associated to sailing the company's own fleet plus the cost of using spot carriers. Constraints (2) ensure that each cargo is transported, either by a company-controlled ship or by using a spot carrier. Flow conservation is expressed through constraints (3)–(5). Time variables are controlled through constraints (6) and (7): the former make sure that travel times are added between consecutive visits on a ship route, whereas the latter ensure that all visits are made within the time windows for service. The large constants in constraints (6) can be set equal to  $M_{ijv} = E_{iv} + T_{ijv} - B_{jv}$ . Constraints (8) state that the ship is empty when it leaves the origin node. When visiting a pickup node, constraints (9) make sure that the load is increased at least by the size of the cargo transported. Correspondingly, constraints (10) force the load to be reduced by no more than the size of the cargo delivered. Together, constraints (9) and (10) ensure that  $l_{iv}$  becomes a lower bound on the amount carried by the ship at any time, thus forcing capacity restrictions to hold when coupled with constraints (11) and (12). For each cargo, loading and unloading nodes must either be served by a spot carrier or by the same ship, which is enforced through constraints (13). If served by the same ship, the loading node must be served before the unloading node, in accordance with constraints (14). Finally, the domain of the variables are given by constraints (15)–(18).

The above formulation describes a static and deterministic problem, whereas the problem studied here is in fact dynamic and stochastic. Hence, information about new cargoes, including information about loading and unloading ports, time windows, and sizes, will arrive while the routes for the ships in the fleet are actually being constructed. For an industrial shipping company, the problem is also open-ended, so that new information about cargoes will continue to arrive as long as the company is operating. Finally, the problem is stochastic in the sense that new cargoes will arrive based on stochastic processes which can be exploited in the decision making. It is reasonable to believe, in the context of industrial shipping, that historical data and production forecasts can be used to derive probability distributions on the arrivals of cargoes and their properties.

The model defined through expressions (1)–(18) describes the problem that is encountered by the shipping company when ignoring stochastic information, with the exception that some

of the variables may be fixed due to vessels that already have cargoes onboard. When solving the dynamic problem, cargoes that were previously planned may be reassigned to other vessels, unless the planned vessel has already started to sail towards the port associated with the cargo.

## 4 Solution Methods

This section describes solution methods for the dynamic and stochastic version of the industrial ship routing and scheduling problem of Section 3. To evaluate solution methods, the operations of the shipping company are simulated over a long time horizon. Details of this simulation can be found in Section 4.1. During the simulation, decision points arise where the company is allowed to replan the ship routes and schedules. Three different methods to perform this replanning are considered: the myopic dynamic heuristic is described in Section 4.2, the multiple scenario approach with consensus is described in Section 4.3, and the branch-and-regret heuristic is described in Section 4.4. When performing replanning, each of the three methods mentioned above relies on solving instances of the static and deterministic version of the problem described in Section 3. To do this, they make use of a tabu search which is described in Section 4.5.

### 4.1 Simulation

In a dynamic problem, information is revealed over time. The process of making decisions and receiving new information is represented in a discrete event simulation. The state of the simulation system consists of all available information about cargo requests and vessels, as well as the current plan containing routing and scheduling decisions. Events that are treated in the simulation include the arrival of new cargo requests and vessels arriving at a port. The simulation clock can be updated using either fixed time increments or time increments based on the next event. Whenever the simulation clock is updated, whether with a fixed time increment or an increment such that the next event is reached, the information about vessels and cargoes is updated and one of the three methods described in Sections 4.2 to 4.4 is used to update the current routes. Although one could update the routes only when a new request is received, allowing the algorithm to replan also when a ship arrives at a port will typically yield better solutions because the current plan, which is used as an initial solution, can then be further refined before the decision regarding which cargo to serve next is actually made.

When the simulation starts, we assume that all vessels are empty and waiting at a port. Some cargoes are known, and a first plan is made. Vessels are set to depart to their next destination immediately after finishing a loading or unloading operation at a port, as shipping companies aim to maximize the utilization of their own fleet and most ports operate on a first-come, first-served basis. The simulation continues for 360 days, after which key data, such as total cost incurred, are calculated.

To make a fair comparison of the three solution methods used to make decisions, the end of the time horizon must be handled with care. The solution methods are provided with a set of sample scenarios that include stochastic information, but no stochastic information is included beyond the end of the simulation. Furthermore, when no new events will take place within the simulation, the plan with the remaining cargoes is made by solving a static and deterministic

version of the industrial ship routing and scheduling problem. As long as the event queue of the simulation is not empty, one of the solution methods of Sections 4.2 to 4.4 is used to create the current set of routes.

## 4.2 Myopic Dynamic Heuristic (MDH)

The myopic dynamic heuristic (MDH) is a simple heuristic that can be used on dynamic and stochastic transportation problems. It considers the problem as a purely dynamic problem, using only information that is known with certainty and completely ignoring stochastic information concerning future events. Whenever replanning is allowed, a subproblem is considered that consists of all known cargoes. This subproblem is a static and deterministic version of the problem described in Section 3, and is solved directly using the tabu search procedure described in Section 4.5. The solution of each stage is used as an initial solution for the following stage, after removing all cargoes already served and adding the cargoes that became known during the last stage. The cargoes that were loaded but not unloaded are locked to their corresponding vessel, so that they cannot be reassigned to a different vessel during the search process.

The MDH may be considered as a simple benchmark, which is very similar to current practice among shipping companies that use decision support systems, in that a static and deterministic problem is considered every time a new plan is made. The MDH has only one parameter, given that the tabu search used to solve the subproblem has already been tuned: one simply has to determine how many iterations the tabu search should run for each subproblem encountered. An overview of the MDH is shown as Algorithm 1.

---

### Algorithm 1 Myopic dynamic heuristic

---

- 1: **for** each point in time where replanning is allowed **do**
  - 2:   Formulate an instance of the static and deterministic problem (1)–(18), where variables corresponding to decisions that have already taken place are fixed.
  - 3:   Solve the instance using the tabu search of Section 4.5.
  - 4:   Use the solution as the new plan for the overall problem.
  - 5: **end for**
- 

## 4.3 Multiple Scenario Approach with Consensus Function (MSAC)

The multiple scenario approach with consensus function (MSAC) used in our study is in fact an adaptation of the heuristic introduced by Bent and Van Hentenryck (2004) to solve partially dynamic vehicle routing problems with stochastic customers. MSAC is designed to take advantage of stochastic information that is available when solving dynamic and stochastic routing problems. This is done by creating sample scenarios. The sample scenarios incorporate all information that is known with certainty in addition to a possible set of future information that is sampled from the corresponding probability distributions. Each sample scenario gives rise to a static and deterministic subproblem including all the cargoes that are currently known together with some sampled cargoes, generated from the probability distributions, that may appear later.

The subproblems are solved using the tabu search heuristic described in Section 4.5, and the solutions are used to find a plan for how to serve the set of currently known cargoes.

At each decision point where replanning is allowed, the MSAC proceeds by first generating a set of sample scenarios. The static and deterministic subproblem associated to each scenario is solved using the tabu search, using the current plan as a basis for an initial solution. To complete the initial solution, any new cargo not included in the current plan as well as all sampled cargoes in the scenario are initially served by spot carriers. After solving for each sample scenario, one of the plans associated to a sample scenario will be chosen to generate the new active set of routing and scheduling decisions. The choice is made based on the value of a consensus function, evaluating which scenario solution is most similar to the others.

The consensus function used is defined as follows. Having generated a solution for each sample scenario, we define a two-dimensional matrix  $M \equiv (M[v, c])$  where  $M[v, c]$  represents the number of scenario solutions for which known cargo  $c$  is served first by vessel  $v$ . Then, if the scenario solutions are denoted by  $s_1, \dots, s_k$  and the cargo served first by vessel  $v$  in solution  $s_i$  is denoted by  $c_v^i$ , the considered consensus function  $f$  is defined as  $f(s_i) = \sum_{v \in V} M[v, c_v^i]$ , which represents the score of scenario solution  $s_i$  according to the consensus evaluation. Then, the scenario solution with the highest score is used to construct the final plan for the current stage, with ties broken randomly. To complete the final plan, sampled cargoes are removed from the chosen sample scenario solution. The consensus function is calculated only based on the next cargo handled by each vessel, and the remaining schedule is ignored. This makes sense as long as the method is only being used to select the first destination for each vessel, and it should be applied again when it is time for the vessel to choose its next task.

The MSAC has three important parameters: the number of iterations used in the tabu search, the number of sample scenarios used, and finally a parameter controlling the number of sampled cargoes to be included in each scenario. Since the problem is open-ended there is no natural point in time that can be used to stop generating sampled cargoes. It is necessary to include a large enough number of sampled cargoes to properly represent the potential effects of future cargo requests on the routing and scheduling decisions. At the same time, avoiding too many sampled cargoes is necessary so that the subproblems do not become too difficult and time consuming to handle for the tabu search. A parameter is therefore used to specify the number of sampled cargoes to generate as a percentage of the current number of cargo requests known with certainty. The MSAC is summarized in Algorithm 2.

#### 4.4 Branch-and-Regret Heuristic (BRH)

The branch-and-regret heuristic (BRH), introduced by Hvattum et al. (2007) as an improvement of the dynamic stochastic hedging heuristic of Hvattum et al. (2006), is another heuristic designed to take advantage of stochastic information that is available when solving dynamic and stochastic transportation problems. As in the case of MSAC, this is done by using sample scenarios generated from the appropriate probability distributions. Rather than solving a problem for each scenario once and then choosing one of the solutions, the BRH works by iteratively making decisions about the new plan and solving for each scenario to evaluate the decisions. That is, the BRH starts by solving all scenario problems separately, but in these solutions decisions regarding known cargoes will most likely end up being different due to the presence of sampled

---

**Algorithm 2** Multiple scenario approach with consensus function

---

- 1: **for** each point in time where replanning is allowed **do**
  - 2:   Generate  $k$  sample scenarios.
  - 3:   Formulate  $k$  instances of the static and deterministic problem (1)–(18), where variables corresponding to decisions that have already taken place are fixed and where the cargoes included correspond to those present in the sample scenarios.
  - 4:   Solve the instances using the tabu search of Section 4.5.
  - 5:   Calculate  $M[v, c]$ , for each combination of vessel  $v$  and known cargo  $c$ .
  - 6:   Find the scenario  $i$  that maximizes the consensus function  $f(s_i)$ .
  - 7:   Remove all sampled cargoes from the solution, and use as the new plan for the overall problem.
  - 8: **end for**
- 

cargoes. Conceptually, the BRH then identifies structural decisions and branches on them until all scenario problems have solutions where all known cargoes are served according to the same plan.

In detail, at each decision point reached during the simulation, the BRH is divided into two phases as follows. First, all scenario subproblems are solved individually using tabu search. Then, the frequency of assigning cargo  $c$  to vessel  $v$  in the solutions is recorded for every pair of known cargo  $c$  and vessel  $v$ . The pair  $(c, v)$  with the highest frequency is chosen and then two branches are created, one forcing cargo  $c$  to be served by vessel  $v$  and another one forbidding that to happen. Solutions for all scenarios are found for both branches and the one with the smallest overall cost (averaged over all scenarios) is chosen. If the first branch is selected, cargo  $c$  is locked to vessel  $v$  for the remaining iterations. On the contrary, if the second branch is chosen, cargo  $c$  is declared incompatible with vessel  $v$  to ensure that in the remaining iterations cargo  $c$  will not be served by vessel  $v$ . This phase stops when all cargoes are locked to some vessel or there are no unlocked cargoes compatible with more than one vessel.

After finishing the first phase, each cargo is assigned to a specific vessel, but each scenario solution may provide a different sequence for each vessel (containing the same cargoes). That is, for each vessel  $v$  several sequences of known cargoes may be obtained by ignoring sampled cargoes from the solutions of the scenarios. In practice, the number of different sequences will be small due to the presence of time windows. The second phase focuses on determining which sequence (or route) to use for each vessel. This selection is performed as follows. First, the list of vessels is sorted according to the frequency of the most frequent sequence over the set of scenarios, and the vessels are considered following that order. For a given vessel, each available sequence is considered and solutions to all scenarios are found while using that particular sequence, obtaining an evaluation of the sequence. Then the sequence with the best overall evaluation is selected and used for its corresponding vessel. This decision process is performed for each vessel until there are no more vessels remaining. After finishing the second phase a plan is obtained where each cargo is assigned to a vessel and each vessel follows a particular route. At this point, all sample scenarios have solutions wherein the same plan is followed, and the average cost of the sample scenarios gives some indication of what will be the future total cost of performing the routes. An overview of the BRH is shown as Algorithm 3.

---

**Algorithm 3** Branch-and-regret heuristic

---

- 1: **for** each point in time where replanning is allowed **do**
  - 2:   Generate  $k$  sample scenarios.
  - 3:   Formulate  $k$  instances of the static and deterministic problem (1)–(18), where variables corresponding to decisions that have already taken place in reality are fixed and where the cargoes included correspond to those present in the sample scenarios. Also fixed are variables corresponding to decisions selected during the BRH, including that some cargoes are forced to be visited by some vessels, and some cargoes are prohibited from being visited by some vessels.
  - 4:   Solve the instances using the tabu search of Section 4.5.
  - 5:   Count the frequency of known cargo  $c$  being served by vessel  $v$ .
  - 6:   Consider all pairs  $(c, v)$  on which the heuristic has not previously branched, and select the one with highest frequency. If the frequency is 0 go to Step 8.
  - 7:   Solve all scenario problems again, so that solutions are obtained both for the case where cargo  $c$  is served by vessel  $v$  and where cargo  $c$  is served by a different vessel. Calculate the average cost of solutions representing both those decisions and enforce from here on the decision associated to the lowest cost. Go back to Step 3.
  - 8:   For each vessel, make a list of all sequences in which known cargoes are being served. Order the vessels according to the frequency of their most used sequence.
  - 9:   Consider the vessel for which a sequence has not yet been determined that has the most frequently used sequence. Solve all scenario problems again, to obtain a solution for each choice of sequence. Calculate the average cost of scenario solutions for each sequence and enforce from here on the sequence associated to the lowest cost.
  - 10:   If some vessels remain where different sequences are used in the scenario solutions, go to Step 9.
  - 11:   Consider the final solution where all sampled cargoes are disregarded, and use this as the new plan for the overall problem.
  - 12: **end for**
-

The principle behind this implementation of the BRH is the same as in Hvattum et al. (2007), but the details vary considerably. Both implementations consider two phases. Here, the first phase branches on whether a given cargo should be serviced by a given vessel, whereas in Hvattum et al. (2007) the branching was on whether a given cargo should be serviced before or after a given point in time. The second phase here creates branches for a number of different visiting sequences, whereas in Hvattum et al. (2007) the second phase created branches for at most two different combinations of customer and vehicle, thus trying to decide which customer to service first on which vehicle.

## 4.5 Tabu Search

To solve the deterministic subproblems arising every time replanning is to be performed we have adapted the fast tabu search heuristic for ship routing and scheduling of Korsvik et al. (2010), which is itself based on the unified tabu search algorithm of Cordeau et al. (2001) for vehicle routing problems with time windows. Since the subproblems we consider are solved repeatedly, some simplifications are made to speed up the execution of the tabu search. Additionally, our tabu search incorporates some extra features that are necessary to handle the specific ship routing problems arising as subproblems in the BRH.

New instances of the static and deterministic problem arise as the simulation proceeds. In the new instances, some new cargoes may have been added and some previously known cargoes may already be picked up or delivered. The initial solution for the new instance is created by modifying the best solution from the previous planning period, disregarding cargoes that have already been delivered and otherwise fixing parts of the routes corresponding to actions already performed. Thus, deliveries corresponding to cargoes already picked up are not allowed to be re-scheduled. The option of using spot carriers to transport cargoes is handled by a dedicated dummy ship. All cargoes in the new instance that were not part of the solution for the previous planning period are initially allocated to the dummy ship, and in particular, for the very first instance in the simulation all cargoes are initially placed on the spot carrier dummy ship.

At each iteration of the tabu search, the best neighbor of the current solution is selected as the current solution for the next iteration. The neighborhood of a solution  $s$  is formed by all solutions that can be obtained from  $s$  by moving one cargo from one vessel to another. For each cargo to be moved and each candidate vessel, the loading and unloading positions for that cargo must be determined. To reduce the size of the neighborhood, the best loading position is found first, and then the best unloading position is determined subject to the fixed loading position. The data concerning these moves are stored to be used in subsequent iterations and avoid unnecessary recalculations. Once a move is made, moving cargo  $j$  from vessel  $v$ , it will be considered tabu to move the cargo back to the same vessel. The tabu status is upheld for a fixed number of iterations during which moves involving cargo  $j$  onboard vessel  $v$  are prohibited. An aspiration criterion applies, so that the tabu status is relaxed for moves leading to a solution that is better than the best solution found so far.

The tabu search includes a relaxation mechanism to allow the consideration of intermediate infeasible solutions violating capacity or time window constraints. These violations are penalized in the move evaluation function, and the weights associated to these penalties are updated dynamically. If we denote the capacity and time window violations of a solution  $s$  as  $q(s)$  and

$w(s)$ , respectively, a solution  $s$  is evaluated according to the function  $f(s) = c(s) + \alpha q(s) + \beta w(s) + p(s)$ , where  $c(s)$  is the cost of the solution (transportation and spot carrier costs) and  $\alpha$  and  $\beta$  are the weights (real positive values) of the violations. The term  $p(s)$  refers to a dynamic, solution-dependent penalty that is added to diversify the search. This penalty is proportional to the number of times that the move leading to a new solution has been performed and a scaling factor to control the intensity of the diversification. More details on the tabu search heuristic can be found in Korsvik et al. (2010).

Since the static and deterministic version of the ship routing and scheduling problem must be solved frequently, it is important that the execution of the tabu search is as fast as possible. Thus, we do not consider the periodic route reoptimization process and the final intensification phase used by Korsvik et al. (2010). We employ the same parameter settings as Korsvik et al. (2010), except that the update factor of  $\alpha$  and  $\beta$  and the diversification intensity are kept fixed during the search. A problem feature that is distinct for problem instances encountered in the BRH is that certain cargoes may be either fixed to a specific vessel or prohibited from being served by a specific vessel. To allow these conditions to be imposed, a compatibility matrix is introduced, stating which vessels can be used for each cargo, and only moves satisfying the conditions of the compatibility matrix are considered in the search.

## 5 Computational Study

Computational experiments were carried out in two parts. First, we have calibrated the solution methods on a small set of instances as explained in Section 5.1. Second, we have run the methods on a set of realistic instances with varying characteristics as explained in Section 5.2. Each instance was simulated over a one-year horizon, using the same realized demands irrespective of the solution method used.

Separate sets of instances were used for the calibration and main tests, respectively. In the following we describe the main test set. The instances vary based on the fleet size (small size fleet with 6 vessels, medium size fleet with 12 vessels, or large size fleet with 24 vessels), the fleet composition (heterogeneous, with vessels with different cost and capacity, or homogeneous), the size of cargoes (full load or partial load), the number of cargoes (high or low demand), and the number of ports (either 4 production ports and 12 consumption ports or 8 production ports and 24 consumption ports). Varying these aspects simultaneously and generating two distinct instances per setting yields a total of 96 different instances. In addition we varied the cost of using spot carriers (normal fees, medium-high fees, and very high fees), corresponding to market situations where demand is either closely matched by supply or, for the latter, significantly higher than the supply.

When generating instances, a distance matrix based on actual distances between real ports is used. Sailing times between ports vary from one to 36 days. Disjoint sets of production ports and consumption ports are selected randomly from the full set of ports. Each pair of production and consumption port is assigned a daily probability of generating new cargo requests, and this probability is set to zero if the consumption port will never require delivery from the production port. In addition, there is a minimum number of days between cargo requests between two given ports. Uniform probability distributions are generated for the cargo size and the number of days

prior to the pickup time window that the request will be known. The size of the time windows for pickup and delivery is fixed, but varies between different pairs of ports: time windows for pickups span between 10 and 20 days, and time windows for deliveries span up to three months depending on the distance between ports. The vehicle fleet may consist of up to three different types of ships, where the capacity of the largest ship is slightly less than twice that of the smallest ship. When generating full load instances the size of each cargo is set equal to the size of the smallest ship, thus ensuring that at most one cargo can be onboard a ship at any time. When generating partial load instances, the size of each cargo is random but scaled so that the largest ship can carry about four average size cargoes simultaneously. Note that no special treatment of the full load instances is given by the solution methods.

## 5.1 Calibration

Calibration tests were performed on a set of five selected instances, generated separately from the instances used in the main test. These instances all had heterogeneous fleets and partial loads, but varied with respect to fleet size, level of demand, and the number of ports. In a first set of experiments illustrated in Figure 1 we calibrated the number of iterations to use per subproblem for the MDH. At the same time we tested different strategies for updating the time during the simulation: either based on events such as arrival of new cargoes and arrivals at ports, or using a fixed time increment of two, four, eight, or sixteen days. The numbers on the secondary axis have the following meaning: each data point consists of the average normalized quality over the five instances used in the calibration, where the quality is a number in the  $[0, 1]$  interval for each instance. For each instance, the best result for any method in the test is assigned a value of 1, and the worst result is assigned a value of 0. This also applies to Figures 2–5. For the MDH, using a fixed interval time update is not worse than using event-based updates, although large steps seem worse than small steps. To facilitate comparisons with other solution methods, where MSAC is built around using event-based updates, all three methods will use the event-based updates in the following. It is also clear that increasing the number of tabu search iterations per subproblem does not necessarily improve the total cost: better solutions to the static and deterministic subproblems do not translate into better solutions for the underlying dynamic and stochastic problem.

For the MSAC there are three parameters to consider: the number of scenarios, the number of sampled cargoes contained in each scenario, and the number of tabu search iterations used to solve each scenario. Figure 2 shows results where the number of scenarios is fixed to 15, and where the number of iterations is varied over the primary axis. The number of sampled cargoes in each scenario is measured relative to the number of known cargoes, so that fewer sampled cargoes are added when few cargoes are known. Results are reported where the number of sampled cargoes in each scenario varies between 10 and 150%. It is a plausible conjecture that the number of iterations per subproblem required to obtain good solutions increases with the number of sampled cargoes. It then becomes a matter of balancing the computational effort required with the quality of results obtained, to select suitable parameters for the method. With a reasonable number of iterations per subproblem, using 100% sampled cargoes looks best. Similar results were obtained with the number of scenarios fixed to 30.

Figure 3 shows the effect of varying the number of scenarios used for each subproblem by

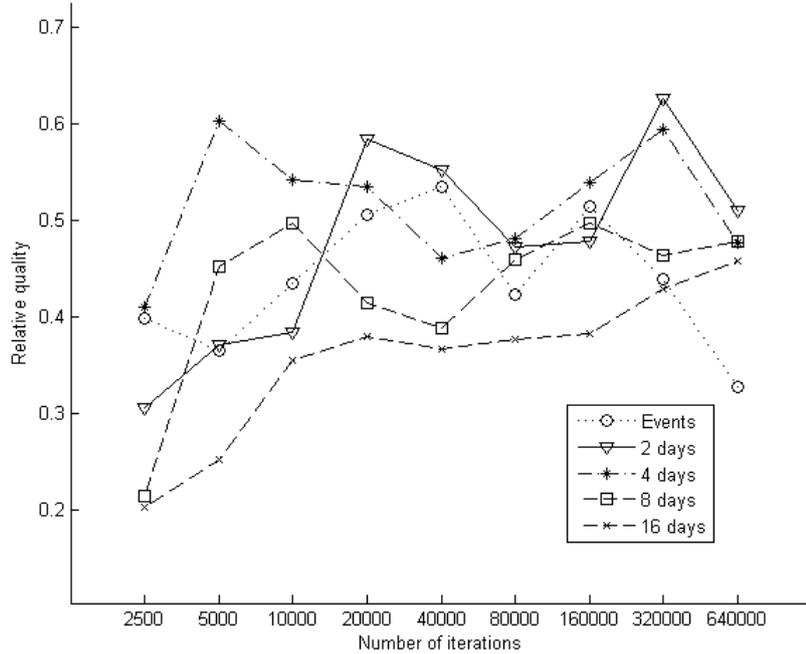


Figure 1: Testing different time update mechanisms, with varying number of iterations per subproblem, for the MDH.

the MSAC, using 100% sampled cargoes. It appears that stable performance is obtained once the number of scenarios reaches about 15. Also, if there is a potentially significant improvement to be achieved by increasing the number of iterations, this improvement is well hidden by the randomness of the solution instances.

Figures 4 and 5 illustrate the calibration of the BRH method. In Figure 4 the number of sampled cargoes is fixed to 50%, and in Figure 5 the number of scenarios is fixed to 30. Again, it appears that a number of scenarios approaching 15 to 30 is sufficient to obtain good results. For the BRH the number of iterations take a different meaning since one cannot tell in advance how many times each scenario is solved. However, using around 200 iterations per recalculation looks sufficient. Contrary to the MSAC, the BRH may benefit from a lower number of sampled cargoes per scenario. Although there is a trend that increasing the number of iterations per recalculation allows using a larger number of sampled cargoes, the best results are obtained using relatively few sampled cargoes.

In summary, we use for all methods the event-based time update. We simulate for one year, and to have a fair treatment of end-of-simulation effects, we solve the last part as a deterministic problem using 10,000 iterations of the tabu search heuristic. We use the solution from the previous subproblem as the initial solution for the new subproblem in all methods. For MDH we use 40,000 iterations of tabu search per subproblem. For BRH we use 30 scenarios with 25%

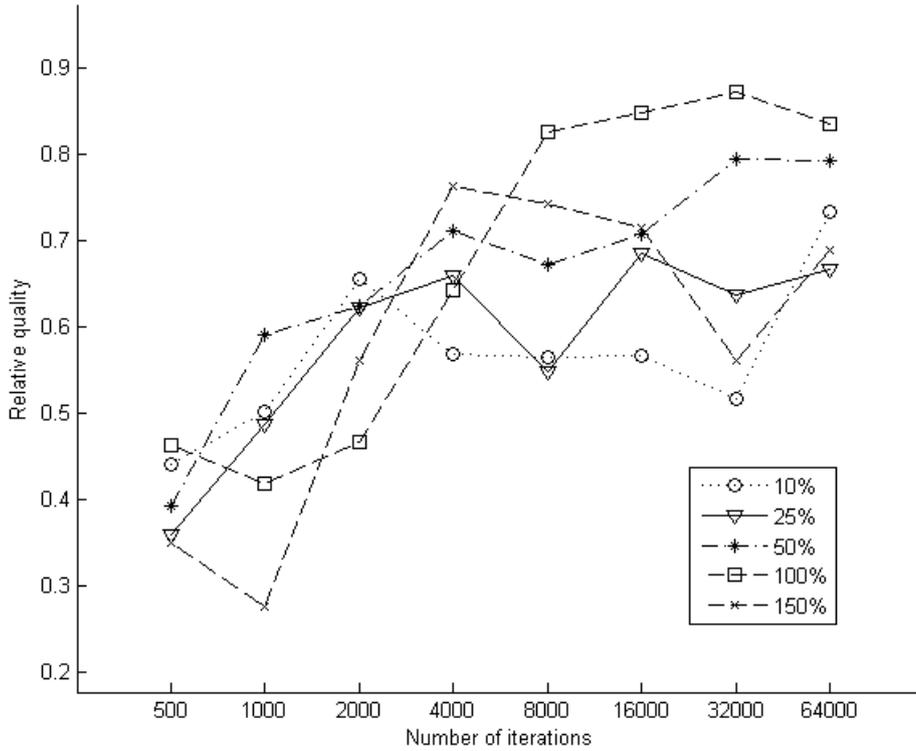


Figure 2: Testing the effect of varying the number of iterations per subproblem and the percentage of sampled cargoes per scenario in MSAC when using 15 scenarios per subproblem.

extra sampled cargoes and 200 iterations of tabu search to evaluate each decision. For MSAC we use 30 scenarios with 100% extra sampled cargoes and 16,000 iterations per subproblem.

## 5.2 Main results

The main results for instances with normal spot carrier costs are presented in Table 1 for instances with partial load, and in Table 2 for full load instances. The running times per decision point are lowest for the MDH and highest for the MSAC. Over all instances, the most computationally demanding instances require slightly less than half an hour of CPU time per decision point when using the MSAC. This is acceptable, given that planning is typically performed at most a few times on a given day. Both MSAC and BRH are improving on MDH in terms of the total cost, by about 1–2% on the partial load instances and only marginally on the full load instances. The cost reduction is broken down on different subsets of instances, and numbers in bold in the tables indicate that the cost saving is statistically significant on a 5% level. So even if the improvement on full load instances is marginal, it is significant. Moreover,

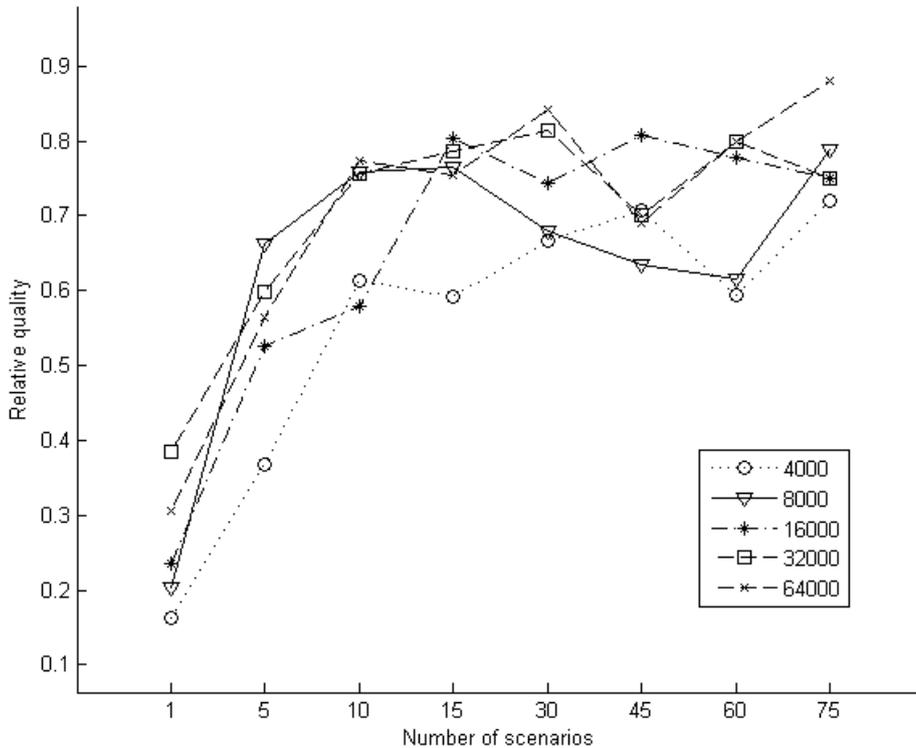


Figure 3: Illustrating the effect of varying the number of scenarios and the number of tabu search iterations per subproblem in MSAC, using 100% sampled cargoes in each scenario.

the cost saving is larger for instances with high demand, and for instances with fewer ports.

For instances with normal spot carrier costs, the improvements by MSAC are accompanied by a reduction of the number of spot carriers used, whereas for BRH it is accompanied by an increase in the number of spot carriers used. This indicates that the three methods produce solutions that are structurally different. Moreover, the BRH is able to produce solutions with reduced total cost despite using more spot carriers than the other methods. This observation led us to perform tests on the same instances but with increased cost of using spot carriers. The results of these tests are presented in Table 3 for medium high spot carrier costs and in Table 4 for high spot carrier costs. Only results for partial load instances are presented, as results for full load instances follow the same pattern but with smaller effects on the cost savings. For the MSAC, having larger spot carrier costs leads to a larger reduction in the number of spot carriers used compared to the MDH. Also, the cost saving becomes larger as the spot carrier costs increase, and on instances with high demand and high spot carrier costs the MSAC can reduce total cost by almost 10%. This indicates that incorporating stochastic information in this dynamic setting is more important when the total demand is high compared to the total supply. The BRH only slightly reduces the number of spot carriers used, even for problems with high

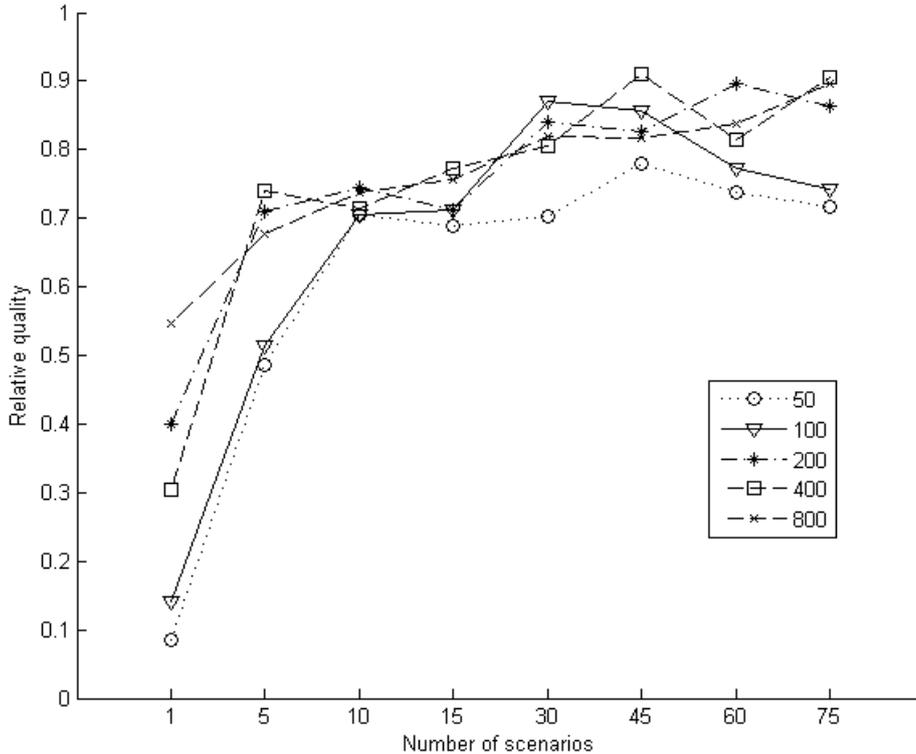


Figure 4: Calibrating the number of scenarios per subproblem and the number of tabu search iterations per recalculation in BRH, using 50% sampled cargoes per scenario.

spot carrier costs. However, the total cost savings increase with the spot carrier costs, although not as much as for MSAC. The conclusion is that the methods still produce solutions that are structurally very different.

## 6 Concluding Remarks

This paper has addressed a dynamic and stochastic routing problem in industrial shipping. Three heuristics based on different strategies were presented and described in detail. MDH is a deterministic algorithm based only on known information, and BRH and MSAC are heuristics that take into account the stochastic information available concerning future cargo requests to design the current plan. In all three methods a tabu search procedure is used to solve the deterministic subproblems arising when replanning is performed. To test the solution methods, a large set of test instances were generated. Results show that the MDH is effective up to a certain point when increasing the number of iterations performed in the tabu search. However, when increasing the number of iterations further, the solutions of the deterministic replanning

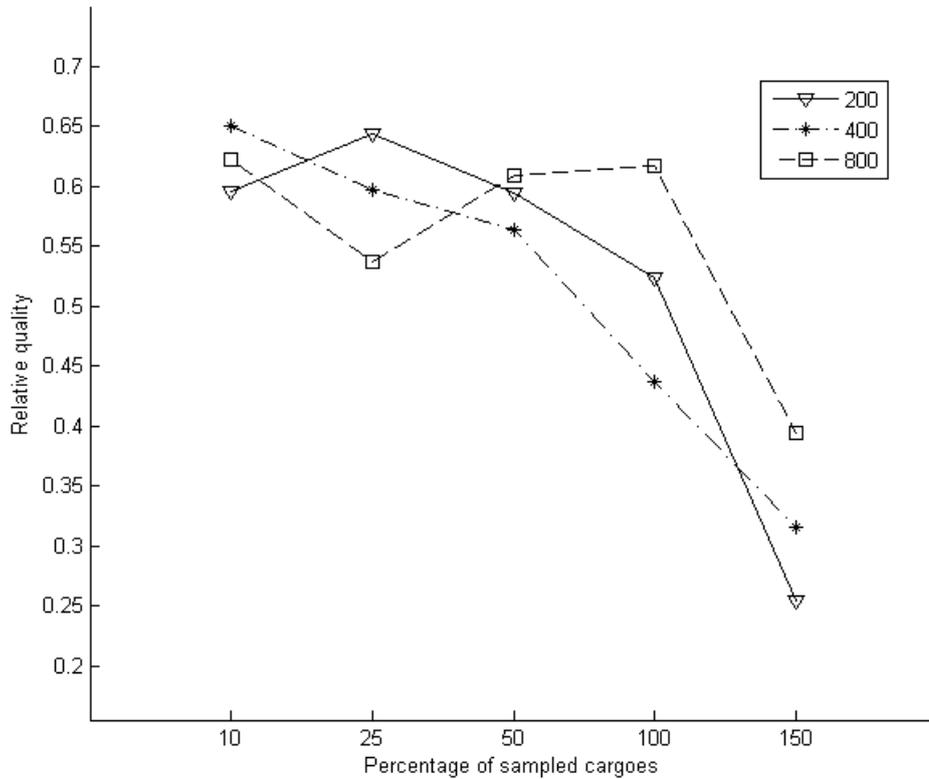


Figure 5: Calibrating the number of tabu search iterations and the percentage of sampled cargoes in BRH, using 30 scenarios per subproblem.

problems may improve without reducing the final costs observed by the simulation. Both the MSAC and the BRH are able to make better use of increased computational effort, as increasing the number of scenarios and increasing the number of tabu search iterations used to solve the scenarios both improve the quality of the implemented solutions. Overall, the costs obtained using BRH and MSAC are roughly between zero and five percent lower than when using MDH, but this varies according to the structure of the problem instance.

The results obtained have implications with respect to which methods should be implemented in decision support systems for maritime transportation. It is quite clear that improved solutions for the static and deterministic problems may not correspond to improved results for the shipping company, as seen in Figure 1. Hence, attempting to solve this static and deterministic problem to optimality may be of little value in practice. However, making use of alternative methods that incorporate stochastic information requires that this information is readily available, which may be an obstacle for efficient use. The cost savings obtained in the test results reported here may at first glance appear small, but actually correspond to large amounts of money considering the high costs associated with maritime shipping.

Subset of instances	#I	Seconds per decision point			Reduction in spot carriers		Cost saving	
		MDH	MSAC	BRH	MSAC	BRH	MSAC	BRH
Large fleet	16	26	1134	518	18.31	-5.31	0.67%	<b>2.64%</b>
Medium fleet	16	6	241	45	7.75	-2.13	<b>2.41%</b>	<b>1.88%</b>
Small fleet	16	1	35	3	1.81	0.63	<b>0.97%</b>	<b>1.13%</b>
High demand	24	17	780	340	16.58	-1.67	<b>2.77%</b>	<b>2.80%</b>
Low demand	24	5	160	37	2.00	-2.88	-0.07%	<b>0.96%</b>
16 ports	24	12	535	234	11.33	-1.42	<b>1.92%</b>	<b>2.31%</b>
32 ports	24	10	405	143	7.25	-3.13	0.77%	<b>1.46%</b>
All	48	11	470	188	9.29	-2.27	<b>1.35%</b>	<b>1.88%</b>

Table 1: Summary of results for 48 instances with partial load and normal spot carrier costs.

Subset of instances	#I	Seconds per decision point			Reduction in spot carriers		Cost saving	
		MDH	MSAC	BRH	MSAC	BRH	MSAC	BRH
Large fleet	16	2	53	9	0.88	-1.13	<b>0.25%</b>	<b>0.23%</b>
Medium fleet	16	1	13	1	0.25	0.06	0.19%	0.08%
Small fleet	16	0	3	0	0.00	0.06	-0.17%	0.24%
High demand	24	2	39	6	0.71	-0.75	0.11%	<b>0.26%</b>
Low demand	24	0	6	0	0.04	0.08	-0.07%	0.11%
16 ports	24	1	21	3	0.25	-0.42	0.00%	<b>0.22%</b>
32 ports	24	1	24	4	0.50	-0.25	<b>0.18%</b>	0.15%
All	48	1	23	3	0.38	-0.33	<b>0.09%</b>	<b>0.19%</b>

Table 2: Summary of results for 48 instances with full load and normal spot carrier costs.

Although the cost savings obtained with the proposed methods are quite significant, several lines of research are worth exploring further. One is to consider more sophisticated solution methods, such as introducing local search procedures to better exploit the stochastic information available based on the performance on all scenarios. Another relevant opportunity for future work is to extend this research by developing specialized tools for other modes of maritime transportation, such as tramp shipping, that can also largely benefit from stochastic information.

## Acknowledgements

This research was carried out with financial support from the DESIMAL and MARRISK projects, partly funded by the Research Council of Norway, from the NILS Mobility Project, EEA grant UCM-EEA-ABEL-02-2009, from the Government of Spain, grant TIN2009-07901, and from the Canadian Natural Sciences and Engineering Research Council. This support is gratefully acknowledged.

Subset of instances	#I	Seconds per decision point			Reduction in spot carriers		Cost saving	
		MDH	MSAC	BRH	MSAC	BRH	MSAC	BRH
Large fleet	16	28	1250	592	22.06	-4.19	<b>2.70%</b>	<b>4.13%</b>
Medium fleet	16	6	268	54	9.13	0.06	<b>3.22%</b>	<b>2.32%</b>
Small fleet	16	1	38	3	1.25	0.31	<b>1.91%</b>	1.33%
High demand	24	19	871	394	21.17	-1.42	<b>5.35%</b>	<b>4.81%</b>
Low demand	24	5	166	39	0.46	-1.13	-0.12%	<b>0.37%</b>
16 ports	24	13	604	272	13.88	-1.21	<b>2.84%</b>	<b>2.77%</b>
32 ports	24	10	433	161	7.75	-1.33	<b>2.39%</b>	<b>2.42%</b>
All	48	12	519	216	10.81	-1.27	<b>2.61%</b>	<b>2.59%</b>

Table 3: Summary of results for 48 instances with partial load and medium high spot carrier costs.

Subset of instances	#I	Seconds per decision point			Reduction in spot carriers		Cost saving	
		MDH	MSAC	BRH	MSAC	BRH	MSAC	BRH
Large fleet	16	28	1262	608	26.94	0.44	<b>5.76%</b>	<b>6.42%</b>
Medium fleet	16	6	275	54	9.19	0.56	<b>6.07%</b>	<b>4.20%</b>
Small fleet	16	1	40	3	1.75	0.75	<b>3.55%</b>	<b>2.62%</b>
High demand	24	19	882	404	24.75	1.21	<b>9.77%</b>	<b>7.69%</b>
Low demand	24	5	169	39	0.50	-0.04	0.49%	<b>1.14%</b>
16 ports	24	14	613	278	15.17	-0.63	<b>4.66%</b>	<b>4.07%</b>
32 ports	24	10	439	165	10.08	1.79	<b>5.59%</b>	<b>4.75%</b>
All	48	12	526	222	12.63	0.58	<b>5.13%</b>	<b>4.41%</b>

Table 4: Summary of results for 48 instances with partial load and high spot carrier costs.

## References

- A. Azaron and F. Kianfar. Dynamic shortest path in stochastic dynamic networks: Ship routing problem. *European Journal of Operational Research*, 144:138–156, 2003.
- R.W. Bent and P. Van Hentenryck. Scenario-based planning for partially dynamic vehicle routing with stochastic customers. *Operations Research*, 52:977–987, 2004.
- G. Berbeglia, J.-F. Cordeau, and G. Laporte. Dynamic pickup and delivery problems. *European Journal of Operational Research*, 202:8–15, 2010.
- L. Cheng and M.A. Duran. Logistics for world-wide crude oil transportation using discrete event simulation and optimal control. *Computers and Chemical Engineering*, 28:897–911, 2004.
- M. Christiansen, K. Fagerholt, and D. Ronen. Ship routing and scheduling: Status and perspectives. *Transportation Science*, 38:1–18, 2004.

- M. Christiansen, K. Fagerholt, B. Nygreen, and D. Ronen. Maritime transportation. In C. Barnhart and G. Laporte, editors, *Handbooks in Operations Research and Management Science*, pages 189–284. Elsevier, Amsterdam, 2007.
- J.-F. Cordeau, G. Laporte, and A. Mercier. A unified tabu search heuristic for vehicle routing problems with time windows. *Journal of the Operational Research Society*, 52:928–936, 2001.
- K. Fagerholt. A computer-based decision support system for vessel fleet scheduling - experience and future research. *Decision Support Systems*, 37:35–47, 2004.
- K. Fagerholt and H. Lindstad. TurboRouter: An interactive optimisation-based decision support system for ship routing and scheduling. *Maritime Economics and Logistics*, 9:214–233, 2007.
- T. Flatberg, G. Hasle, O. Kloster, E.J. Nilssen, and A. Riise. Dynamic and stochastic vehicle routing in practice. In V.S. Zeimpekis, C.D. Tarantilis, G.M. Giaglis, and I.E. Minis, editors, *Dynamic Fleet Management: Concepts, Systems, Algorithms & Case Studies*, volume 38 of *Operations Research Computer Science Interfaces Series*, chapter 3, pages 41–64. Springer, New York, 2007.
- M. Gendreau, G. Laporte, and R. Séguin. Stochastic vehicle routing. *European Journal of Operational Research*, 88:3–12, 1996.
- M. Gendreau, F. Guertin, J.-Y. Potvin, and R. Séguin. Neighborhood search heuristics for a dynamic vehicle dispatching problem with pick-ups and deliveries. *Transportation Research Part C*, 14:157–174, 2006.
- G. Ghiani, F. Guerriero, G. Laporte, and R. Musmanno. Real-time vehicle routing: Solution concepts, algorithms and parallel computing strategies. *European Journal of Operational Research*, 151:1–11, 2003.
- L.M. Hvattum, A. Løkketangen, and G. Laporte. Solving a dynamic and stochastic vehicle routing problem with a sample scenario hedging heuristic. *Transportation Science*, 40:421–438, 2006.
- L.M. Hvattum, A. Løkketangen, and G. Laporte. A branch-and-regret heuristic for stochastic and dynamic vehicle routing problems. *Networks*, 49:330–340, 2007.
- H.-S. Hwang, S. Visoldilokpun, and J.M. Rosenberger. A branch-and-price-and-cut method for ship scheduling with limited risk. *Transportation Science*, 42:336–351, 2008.
- S. Ichoua, M. Gendreau, and J.-Y. Potvin. Exploiting knowledge about future demands for real-time vehicle dispatching. *Transportation Science*, 40:211–225, 2006.
- M.H. Kang, H.R. Choi, H.S. Kim, and B.J. Park. Development of a maritime transportation planning support system for car carriers based on genetic algorithm. *Applied Intelligence*, 2011. DOI:10.1007/s10489-011-0278-z.
- J.E. Korsvik, K. Fagerholt, and G. Laporte. A tabu search heuristic for ship routing and scheduling. *Journal of the Operational Research Society*, 61:594–603, 2010.

- S.A. Lawrence. *International Sea Transport: The Years Ahead*. Lexington Books, Lexington, MA, 1972.
- H.K. Lo and M.R. McCord. Adaptive ship routing through stochastic ocean currents: General formulations and empirical results. *Transportation Research, Part A*, 32:547–561, 1998.
- O.B.G. Madsen, H.F. Ravn, and J.M. Rygaard. A heuristic algorithm for a dial-a-ride problem with time windows, multiple capacities and multiple objectives. *Annals of Operations Research*, 60:193–208, 1995.
- S. Mitrović-Minić and G. Laporte. Waiting strategies for the dynamic pickup and delivery problem with time windows. *Transportation Research Part B*, 38:635–655, 2004.
- S. Mitrović-Minić, R. Krishnamurti, and G. Laporte. Double horizon based heuristics for the dynamic pickup and delivery problem with time windows. *Transportation Research Part B*, 38:669–685, 2004.
- W.B. Powell. A stochastic formulation of the dynamic assignment problem, with an application to truckload motor carriers. *Transportation Science*, 30:195–219, 1996.
- H.N. Psaraftis. Dynamic vehicle routing problems. In B.L. Golden and A.A. Assad, editors, *Vehicle Routing: Methods and Studies*, pages 223–248. North-Holland, Amsterdam, 1988.
- H.N. Psaraftis. Dynamic vehicle routing: Status and prospects. *Annals of Operations Research*, 61:143–164, 1995.
- D. Ronen. Cargo ships routing and scheduling: Survey of models and problems. *European Journal of Operational Research*, 12:119–126, 1983.
- D. Ronen. Ship scheduling: The last decade. *European Journal of Operational Research*, 71:325–333, 1993.
- M.W.P. Savelsbergh and M. Sol. DRIVE: Dynamic routing of independent vehicles. *Operations Research*, 46:474–490, 1998.
- M. Schilde, K.F. Doerner, and R.F. Hartl. Metaheuristics for the dynamic stochastic dial-a-ride problem with expected return transports. *Computers and Operations Research*, 38:1719–1730, 2011.
- N. Secomandi and F. Margot. Reoptimization approaches for the vehicle-routing problem with stochastic demands. *Operations Research*, 57:214–230, 2009.
- M.R. Swihart and J.D. Papastavrou. A stochastic and dynamic model for the single-vehicle pick-up and delivery problem. *European Journal of Operational Research*, 114:447–464, 1999.
- B.W. Thomas. Waiting strategies for anticipating service requests from known customer locations. *Transportation Science*, 41:319–331, 2007.
- UNCTAD. Review of maritime transportation. Technical report, UNCTAD, 2009.

P. Van Hentenryck, R.W. Bent, and E. Upfal. Online stochastic optimization under time constraints. *Annals of Operations Research*, 177:151–183, 2010.