A Maritime Inventory Routing Problem with Stochastic Sailing and Port Times

Agra, Agostinho
aagra@ua.pt

Christiansen, Marielle
mc@iot.ntnu.no

Delgado, Alexandrino
alexandrino.delgado@docente.unicv.edu.cv

Hvattum, Lars Magnus
lars.m.hvattum@iot.ntnu.no

Thursday 5th February, 2015

Abstract

This paper describes a stochastic short sea shipping problem where a company is responsible for both the distribution of oil products between islands and the inventory management of those products at consumption storage tanks located at ports. In general, ship routing and scheduling is associated with uncertainty in weather conditions and unpredictable waiting times at ports. In this work, both sailing times and port times are considered to be stochastic parameters. A two-stage stochastic programming model with recourse is presented where the first stage consists of routing, loading and unloading decisions, and the second stage consists of scheduling and inventory decisions. The model is solved using a decomposition approach similar to an L-shaped algorithm where optimality cuts are added dynamically, and this solution process is embedded within the sample average approximation method. A computational study based on real-world instances is presented.

Keywords: Stochastic programming; Maritime transportation; Inventory routing; Uncertainty; L-shaped method; Sample average approximation; Travel time; Service time.

Thursday 5th February, 2015 08:39

1 Introduction

Maritime transportation is characterized by high levels of uncertainty. In practice, operational plans are often adjusted due to factors such as changing weather conditions, port congestions, or mechanical problems at port. A plan that minimizes the transportation and port costs based on expected sailing and port times may not necessarily be good, as it does not account for the consequences resulting from delays. Hence, in most practical situations it will be beneficial to consider the possibility of delays when trying to minimize costs.

In this paper we study a maritime inventory routing problem occurring at the archipelago of Cape Verde. A deterministic variant of this problem was solved to optimality in [3] for short time horizons. Heuristics for the same problem with time horizons up to 6 months were developed in [4]. The deterministic methods assume known and fixed sailing times, but the planner needs to face the uncertainty associated with the ships sailing between ports. This may somehow be circumvented by the inclusion of safety stocks or by artificially increasing the sailing times to compensate for delays.

The ports are used by several independent shipping companies, and limited coordination between the various operators can result in heavy port congestion. This may come from limited capacities in the inner port area, at berths, and of pipes and other important equipment for performing the (un)loading operations. Here, we explicitly consider uncertainty in both sailing times between ports and waiting times at ports.

In the problem considered, a heterogeneous fleet of ships is transporting several oil products between supply and consumption storage tanks. These tanks are located close to or at ports, and a particular port might have both supply and consumption storage tanks for various products. Inventory management is only considered for the consumption storage tanks because these tanks have limited capacity, and it is important that these tanks do not run empty. In contrast, unlimited capacity and supply of products are assumed for the supply storage tanks. By taking the uncertain sailing times and port waiting times into account, the objective of the planning is to design cost efficient ship routes and schedules including determining the number of visits to each port and the (un)loading quantities at each visit.

The purpose of this paper is to describe a stochastic programming model with recourse where the routes and the quantities to load and unload must be fixed *a priori*, that is, before actual values of the uncertain parameters are revealed, while the schedule of the loading and unloading operations can be adjusted according to the observed sailing and port times. In addition, the paper contributes with a solution method that combines the use of the sample average approximation method with a decomposition procedure resembling an L-shaped method [11, 17]. For a given set of scenarios, the corresponding two-stage model is solved to obtain a candidate solution. This is repeated for several different sets of scenarios to obtain several candidate solutions. To choose the best solution, these candidate solutions are evaluated for a larger and independent set of scenarios. To solve the two-stage model for a given set of scenarios, the problem is decomposed into a master problem and one subproblem for each scenario, where the second-stage decisions are considered in the subproblems. We show that feasibility of the second stage is always guaranteed when the solution obtained for the first-stage is feasible. Then we show how to derive optimality cuts from the subproblems that are added dynamically to the master problem. As far as we know, this paper describes the first stochastic programming model and approach for solving a maritime inventory routing problem with uncertain sailing and port waiting times.

The remainder of this paper is organized as follows. In Section 2 we describe the real problem. Relevant literature is reviewed in Section 3. Then, in Section 4 we present a scenario-based mathematical formulation for the problem. The solution approach based on decomposing the problem is discussed in Section 5. In Section 6 we describe how the stochastic sailing and port times have been modeled, and how scenarios have been generated. Section 7 contains a computational study based on real-world instances, and in Section 8 we present the main conclusions of this work.

2 Problem description

In Cape Verde, fuel oil products are imported and delivered to two specific islands and stored in large supply storage tanks, see Figure 1. From these islands, fuel oil products are distributed among all inhabited islands using a small heterogeneous fleet of ships. The two circled islands in Figure 1 are served by a different ship and may therefore be disregarded. Products are stored in separate consumption storage tanks with limited capacity. Some ports have both supply tanks for some products and consumption tanks for other products. As the capacities of the supply tanks are very large compared to the total consumption over the planning horizon, the inventory aspects for these tanks can be ignored. The driving force in the problem is the need for fuel oil products in the consumption storage tanks. If the demand is not satisfied, the backlogged demand will be penalized by a cost.

Not all islands consume all products. Consumption rates are assumed to be constant over the time horizon. Each port can receive at most one ship at a time, and in some ports there exists a minimum time interval between the departure of one ship and the arrival of the next ship.

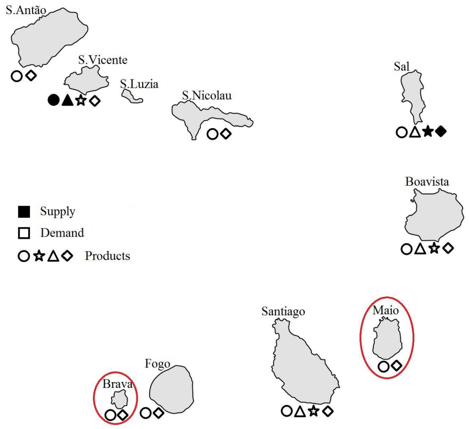

Figure 1: Supply and demand for fuel oil products at several islands in Cape Verde.

Each ship has a specified capacity, fixed speed, and cost structure. The cargo hold of each ship is separated into several cargo tanks. The products cannot be mixed, so we assume that the ships have dedicated tanks for the particular products. The ships are either sailing, waiting outside a port or operating. Here, operating is the common term for loading and unloading.

At port, we consider set-up times for the coupling and decoupling of pipes and operation times which depend on the amount loaded or unloaded. Minimum and maximum unloading quantities can be derived based on properties of the ships and the inventories. The maximum unloading quantity is imposed by the consumption tank capacity and by the ship cargo tank capacity.

The traveling times depend upon the weather conditions and are considered stochastic. The uncertain time parameter at port is mainly related to the time from arrival to start of operation. Hence, a specified waiting time before start of service is defined as stochastic, while the operation times are deterministic.

The inter-island distribution plan consists of routes and schedules for the fleet of ships, and describes the number of visits to each port and the quantity of each product to be loaded or unloaded at each port visit. This plan must satisfy the capacities of the ships and consumption storage tanks while minimizing the sailing and port costs as well as the expected penalty costs of backlogged demand. There is great flexibility in the route pattern of a ship, such that a ship may load in several successive ports as well as performing unloading operations in succession. The quantities loaded or unloaded are variable as well as the number of visits at each port. The problem described here will be referred to as a *stochastic maritime inventory routing problem* (SMIRP), and a scenario-based stochastic programming model for the problem is given in Section 4.

3 Literature review

The amount of literature on maritime transportation optimization has increased steadily over the last decades, as evidenced through the recent survey in [13]. Despite being a transportation mode that is heavily influenced by uncertainty, most of the literature on maritime routing and scheduling involves solving static and deterministic problem variants. However, some contributions exist, and we describe some that are considering problems close to the stochastic maritime inventory routing problem of this paper.

An inventory routing problem with uncertain demands and sailing times was solved heuristically by Cheng and Duran [12]. Rakke et al. [28] and Sherali and Al-Yakoob [30, 31] introduced penalty functions for deviating from the customer contracts and the inventory limits, respectively. Christiansen and Nygreen [14] introduced soft inventory levels to handle uncertainties in sailing and port times, and these levels were transformed into soft time windows.

Agra et al. [5] solved a full-load ship routing and scheduling problem with uncertain travel times using robust optimization. Halvorsen-Weare and Fagerholt [15] considered weather conditions that affect both sailing speeds and the loading and unloading operations for supply vessels servicing offshore installations, and proposed various heuristic strategies to achieve robust weekly voyages and schedules. Heuristic strategies for obtaining robust solutions with uncertain sailing times was also discussed by Halvorsen-Weare et al. [16] for the delivery of liquefied natural gas. None of the aforementioned papers described the use of stochastic programming to model uncertain sailing and port times.

Considering the literature on non-maritime transportation optimization, some researchers have developed models and exact solution methods for vehicle routing problems (VRPs) with uncertain travel times. Considering a VRP with stochastic travel times and service times, Laporte et al. [25] presented a chance-constrained formulation as well as two recourse formulations. The recourse problem was solved to optimality for up to 20 nodes and 5 scenarios using an integer L-shaped method. The VRP with stochastic travel and service times was also studied by Kenyon and Morton [21], considering stochastic programming models that minimized the expected completion time or maximized the probability of completing the routes within a given deadline. An integer L-shaped algorithm was used by Teng et al. [34] to solve a time-constrained *traveling salesman problem* with stochastic travel and service times with up to 35 nodes. Taş et al. [33] studied a VRP with stochastic travel times and soft time windows, solved by branch-and-price. Although these papers present stochastic programming models for routing problems with uncertain travel times and service times, they do not consider heterogeneous fleets, a variable number of visits, and inventory constraints.

Other researchers have focused on heuristic solution methods for VRPs with uncertain travel times. A stochastic model for one particular version of the VRP with stochastic travel times was presented by Lambert et al. [24], and a heuristic solution method was proposed. Ando and Taniguchi [7] used a genetic algorithm to solve a VRP with time windows and stochastic travel times. A tabu search was used by Russell and Urban [29] to solve a VRP with soft time windows and stochastic travel times. Li et al. [27] presented both a chance-constrained model and a recourse model for a VRP with time windows and stochastic travel and service times. The authors proposed a tabu search for solving the two versions of the problem. Lei et al. [26] considered a capacitated VRP with stochastic service times, which was solved using a variable neighborhood search heuristic. Taş et al. [32] used tabu search to solve the VRP with stochastic travel times and soft time windows as defined in [33]. Again, neither of these papers considered heterogeneous fleets, a variable number of visits, and inventory constraints.

The literature on non-maritime inventory routing problems (IRPs) also address uncertainty, which is considered as one of the main directions for future research by Andersson et al. [8]. However, the uncertainty is typically found in the demand of the customers. For example, Kleywegt et al. [22, 23] and Adelman [1] created infinite horizon Markov decision process (MDP) formulations where the state was represented by the inventory levels at customers. The authors solved these MDP problems heuristically. Hvattum et al. [19] and Hvattum and Løkketangen [18] also solved the same problems heuristically, but using an approximation of the demand uncertainty based on scenario trees. Bertazzi et al. [10] used a dynamic programming formulation to describe a finite horizon problem with stochastic demands, which was solved using a hybrid roll-out algorithm. Aghezzaf [2] considered an IRP with uncertain demand and uncertain travel times, and coupled Monte Carlo simulation and a non-linear mixed-integer programming to find robust solutions.

4 Mathematical Model

In this section we introduce a two-stage stochastic programming model with recourse for the SMIRP. The routes and the quantities to load and unload are determined in the first stage. However, the schedule of the loading and unloading operations can be adjusted in the second stage. Thus, the inventory level variables are also allowed to change according to the realization of the stochastic parameters. In the following we first describe the variables and constraints related to the first stage in Section 4.1, and then the variables and constraints related to the second stage in Section 4.2.

4.1 First stage

First we model the routing and the loading and unloading constraints.

Routing constraints:

Let N denote the set of ports and V denote the set of ships. Each ship $v \in V$ must depart from its initial position in the beginning of the planning horizon. In order to distinguish between visits to each port we consider an ordering of the visits according to the time of visit. The ship paths are defined on a network where the nodes are represented by a pair (i, m), where i is the port and m is the m^{th} visit to port i. A direct ship movement (arc) from port arrival (i, m) to port arrival (j, n) is represented by (i, m, j, n). Ship paths are illustrated in Figure 2. For instance, ship 2 leaves origin O_2 and sails to port 4 (for the first visit to this port), then sails to port 2 (for the second visit to this port, since the first visit was made by ship 1), and sails to port 1 for its first visit. Finally, the ship sails to port 3 (for the second visit to port 3, since the first visit was made by ship 1) before it ends at its destination.

We define S^A as the set of possible port arrivals (i, m), S^A_v as the set of port arrivals that may be visited by ship v, S^X as the set of all possible ship movements (i, m, j, n), and set S^X_v as the set of all possible movements of ship v. The set of ships that can visit port i is denoted V_i.

For the routing we define the following binary variables: x_{imjnv} that is 1 if ship v sails from port arrival (i, m) directly to port arrival (j, n), and 0 otherwise; x^O_{imv} that indicates whether ship v sails directly from its initial position to port arrival (i, m) or not; w_{imv} is 1 if ship v visits port i at arrival (i, m), and 0 otherwise; z_{imv} is equal to 1 if ship v ends its route at port arrival (i, m),

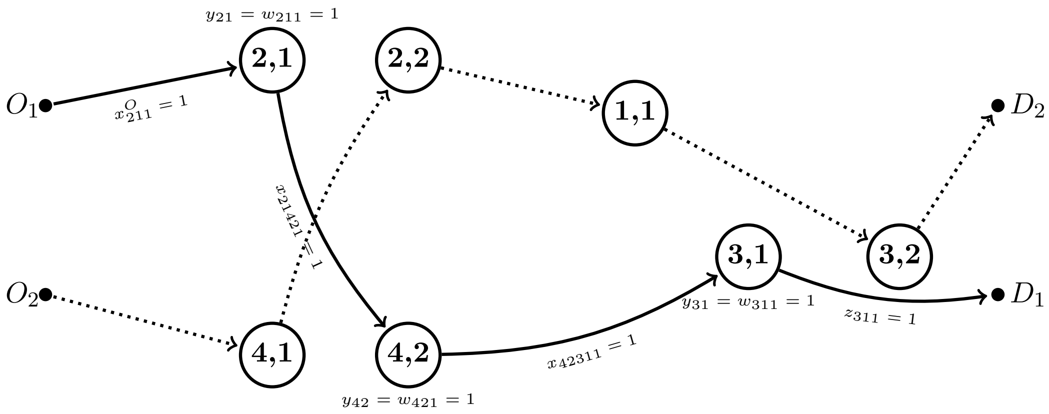

Figure 2: Example of ship routes where each node represents a visit. The first label indicates the port and the second label indicates the visit number. The solid arcs indicate the path of ship 1, while the dotted arcs constitute the path of ship 2.

and 0 otherwise; z_v^O is equal to 1 if ship v is not used and 0 otherwise; y_{im} indicates whether a ship is visiting port arrival (i, m) or not. These variables are included for ship 1 in Figure 2.

$$\sum_{(i,m)\in S_v^A} x_{imv}^O + z_v^O = 1, \qquad v \in V, \tag{1}$$

$$w_{imv} - \sum_{(j,n)\in S_v^A} x_{jnimv} - x_{imv}^O = 0, \qquad v \in V, (i,m) \in S_v^A, \tag{2}$$

$$w_{imv} - \sum_{(j,n)\in S_v^A} x_{imjnv} - z_{imv} = 0, \qquad v \in V, (i,m) \in S_v^A, \tag{3}$$

$$\sum_{v\in V_i} w_{imv} = y_{im}, \qquad (i,m) \in S^A, \tag{4}$$

$$y_{i(m-1)} - y_{im} \geq 0, \qquad (i,m) \in S^A : m > 1, \tag{5}$$

$$x_{imv}^O, w_{imv}, z_{imv} \in \{0,1\}, \qquad v \in V, (i,m) \in S_v^A, \tag{6}$$

$$x_{imjnv} \in \{0,1\}, \qquad v \in V, (i,m,j,n) \in S_v^X, \tag{7}$$

$$z_v^O \in \{0,1\}, \qquad v \in V, \tag{8}$$

$$y_{im} \in \{0,1\}, \qquad (i,m) \in S^A. \tag{9}$$

Equations (1) ensure that each ship either departs from its initial position and sails towards another port or the ship is not used. Equations (2) and (3) are the arc flow conservation constraints, ensuring that a ship arriving at a port also leaves that port or ends its route. Constraints (4) ensure that only one ship visits port (i, m) if y_{im} is equal to one. Constraints (5) state that if port i is visited m times, then it must also have been visited $m - 1$ times. Constraints (6) - (9) define the variables as binary.

Loading and unloading constraints

Let K represent the set of products and K_v represent the set of products that ship v can transport. Not all ports consume all products. Parameter J_{ik} is 1 if port i is a supplier of product k; -1 if port i is a consumer of product k, and 0 if i is neither a consumer nor a supplier of product k. For each port i consuming product k, the demand rate is given by R_{ik}. The quantity of product k on board ship v at the beginning of the planning horizon is given by Q^O_{vk} and C_{vk} is the capacity of the compartment of ship v dedicated to product k. Minimum and the maximum unloading quantities of product k at port i are given by $\underline{Q}_{ik}$ and $\overline{Q}_{ik}$, respectively. Parameter T is the length of the time horizon.

To model the loading and unloading constraints, we define the following binary variables: o_{imvk} is equal to 1 if product k is loaded onto or unloaded from ship v at port visit (i,m), and 0 otherwise. In addition, we define the following continuous variables: q_{imvk} is the amount of product k loaded onto or unloaded from ship v at port visit (i,m); f_{imjnvk} denotes the amount of product k that ship v transports from port visit (i,m) to port visit (j,n), and f^O_{imvk} gives the amount of product k that ship v transports from its initial position to port visit (i,m).

The loading and unloading constraints are given by:

$$f^O_{imvk} + \sum_{(j,n)\in S^A_v} f_{jnimvk} + J_{jk} q_{imvk} = \sum_{(j,n)\in S^A_v} f_{imjnvk}, \qquad v \in V, (i,m) \in S^A_v, k \in K_v, \tag{10}$$

$$f^O_{imvk} = Q^O_{vk} x^O_{imv}, \qquad v \in V, (i,m) \in S^A_v, k \in K_v, \tag{11}$$

$$f_{imjnvk} \le C_{vk} x_{imjnv}, \qquad v \in V, (i,m,j,n) \in S^X_v, k \in K_v, \tag{12}$$

$$0 \le q_{imvk} \le C_{vk} o_{imvk}, \qquad v \in V, (i,m) \in S^A_v, k \in K_v : J_{ik} = 1, \tag{13}$$

$$\underline{Q}_{ik} o_{imvk} \le q_{imvk} \le \overline{Q}_{ik} o_{imvk}, \qquad v \in V, (i,m) \in S^A_v, k \in K_v : J_{ik} = -1, \tag{14}$$

$$\sum_{k\in K_v} o_{imvk} \ge w_{imv}, \qquad v \in V, (i,m) \in S^A_v, \tag{15}$$

$$\sum_{(i,m)\in S^A_v} \sum_{v\in V} \sum_{k\in K_v : J_{ik}=-1} q_{imvk} \ge \sum_{i\in N} \sum_{k\in K : J_{ik}=-1} R_{ik} T, \tag{16}$$

$$o_{imvk} \le w_{imv}, \qquad v \in V, (i,m) \in S^A_v, k \in K_v, \tag{17}$$

$$f_{imjnvk} \ge 0, \qquad v \in V, (i,m,j,n) \in S^A_v, k \in K_v, \tag{18}$$

$$f^O_{imvk} \ge 0, \qquad v \in V, (i,m) \in S^A_v, k \in K_v, \tag{19}$$

$$q_{imvk} \ge 0, \qquad v \in V, (i,m) \in S^A_v, k \in K_v : J_{ik} \ne 0, \tag{20}$$

$$o_{imvk} \in \{0,1\}, \qquad v \in V, (i,m) \in S^A_v, k \in K_v : J_{ik} \ne 0. \tag{21}$$

Equations (10) are the load flow conservation constraints. Equations (11) determine the quantity on board when ship v sails from its initial position to port arrival (i,m). Constraints (12) guarantee that the ships' tank capacities are not exceeded. Upper bounds on the quantities loaded at the ports are imposed by constraints (13), while constraints (14) impose lower and upper limits on the unloaded quantities. Constraints (15) ensure that if ship v visits port arrival (i,m), then at least one product must be (un)loaded. To prevent extreme situations at the end of the planning horizon with full storages at the supply side and empty consumption storage tanks, thereby avoiding transportation costs, constraints (16) ensure that the sum of delivered goods should not be less than the sum of the consumption over the entire horizon T. Constraints (17) ensure that if ship

v (un)loads one product at visit (i,m), then w_{imv} must be one. Constraints (18)-(21) are the non-negativity and integrality requirements.

4.2 Second stage

Now we present the second stage model where the variables can be adjusted to the scenario. The set of scenarios Ω will be indexed by c.

Time constraints

To keep track of the inventory level it is necessary to determine the start and end times of operation at each port arrival. We define the following parameters: T^Q_{ik} is the time required to (un)load one unit of product k at port i; T^S_{ik} is the set up time required to operate product k at port i. T_{ijvc} is the sailing time between port i and j by ship v for scenario c; T^O_{ivc} indicates the sailing time required by ship v to travel from its initial position to port i for scenario c; T^B_i is the minimum interval between the departure of one ship and the next arrival at port i; T^W_{imc} is the waiting time at port arrival (i,m) for scenario c. The parameter μ_i denotes the maximum number of visits at port i. For each scenario c we define start time t_{imc} and end time t^E_{imc} variables for port arrival (i,m). Variables t^+_{ic} give the remaining time from the end of the last visit at port i until time T for scenario c, when this visit occurs before time T. Notice that due to the stochastic time parameters, the start time of a visit might appear after the end of the planning horizon T.

Assuming that a ship v travels from (i,m) to (j,n) under scenario c and loads product k using vessel v, Figure 3 shows the parameters involved when calculating the time variables for node (j,n).

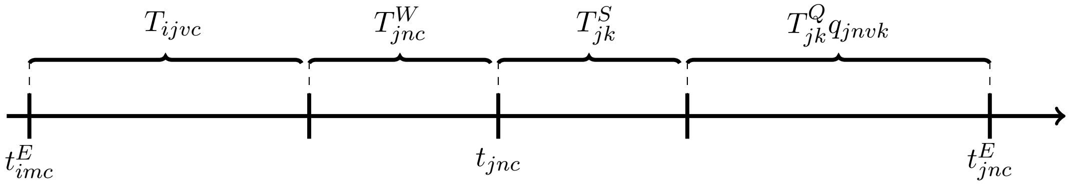

Figure 3: Illustration of the parameters involved when calculating start and end times for node (j,n).

The set of time constraints is as follow:

$$t^E_{imc} \geq t_{imc} + \sum_{v\in V}\sum_{k\in K_v} T^S_{ik} o_{imvk} + \sum_{v\in V}\sum_{k\in K_v} T^Q_{ik} q_{imvk}, \qquad (i,m)\in S^A, c\in\Omega, \tag{22}$$

$$t_{imc} - t^E_{i(m-1)c} - T^B_i y_{im} \geq 0, \qquad (i,m)\in S^A : m>1, c\in\Omega, \tag{23}$$

$$t^E_{imc} + \sum_{v\in V_i\cap V_j} T_{ijvc} x_{imjnv} + T^W_{jnc} y_{jn} - t_{jnc} \leq M(1-\sum_{v\in V_i\cap V_j} x_{imjnv}), (i,m,j,n)\in S^X, c\in\Omega, \tag{24}$$

$$\sum_{v\in V} T^O_{ivc} x^O_{imv} + T^W_{imc} y_{im} \leq t_{imc}, \qquad (i,m)\in S^A, c\in\Omega, \tag{25}$$

$$t^+_{ic} \geq T - t^E_{i\mu_i c}, \qquad i\in N, c\in\Omega, \tag{26}$$

$$t_{imc}, t^E_{imc} \geq 0, \qquad (i,m)\in S^A, c\in\Omega, \tag{27}$$

$$t^+_{ic} \geq 0, \qquad i\in N, c\in\Omega. \tag{28}$$

Constraints (22) define the end time of service at port visit (i, m). Constraints (23) impose a minimum interval between two consecutive visits at port i. Constraints (24) relate the end time of port visit (i, m) to the start time of port visit (j, n) when ship v sails directly from port (i, m) to (j, n). The big-M constant, denoted by M was set to $2T$, since the start time of a visit can occur after time T. These constraints are a stronger version of the usual family of constraints $t^E_{imc} + T_{ijvc} + T^W_{jnc} - t_{jnc} \leq M(1 - x_{imjnv})$ defined for each $v \in V$. Constraints (25) ensure that if ship v travels from its initial position directly to port visit (i, m), then the start time is at least the sailing time between the two positions plus the waiting time at port visit (i, m). Constraints (26) together with (28) determine the time gap between the last visit to port i and time T. The continuous time variables are declared as non-negative in (27) and (28).

Inventory constraints

The inventory constraints are considered for each port i consuming product k ($J_{ik} = -1$). They ensure that the inventory levels are kept within the corresponding bounds, and link the inventory levels to the unloaded quantities.

For each consumption storage tank for product k at port i, the minimum $\underline{S}_{ik}$, the maximum $\overline{S}_{ik}$, and the initial S^O_{ik} inventory levels are given.

We define the nonnegative continuous variables s_{imkc} and s^E_{imkc} indicating the inventory levels of product k at the start and at the end of port visit (i, m) for scenario c, respectively; s^T_{ikc} is the inventory level of product k, above the minimum stock level for port i at the end of time T or at the end of the last visit if this occurs after T, for scenario c; r_{imkc} and r^E_{imkc} indicate the backlog of product k at the start and at the end of port visit (i, m) for scenario c, respectively; and r^T_{ikc} is the amount of product k below the minimum level for port i at the end of time T or at the end of the last visit if this visit occurs after T, for scenario c.

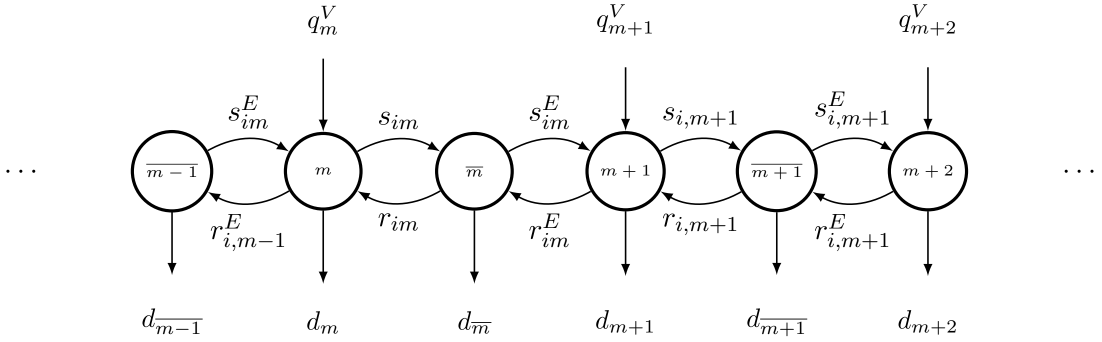

Figure 4: Inventory balance flow for each event represented by a node (start of a visit m, represented by a node with label m, and end of a visit, represented by $\overline{m}$) for a given port i and a given product k (omitted from the nodes and from the variables for simplicity). For the start of visit m, the inflow from the ships is given by $q^V_m = \sum_{v \in V} q_{imvk}$. The demand at the start of visit m, represented in the figure by d_m, is given by the consumption between the end-time of the last visit and the start-time of visit m, $R_{ik}(t_{imc} - t^E_{i(m-1)c})$. The demand at the end of visit m, denoted by $d_{\overline{m}}$ is the consumption between the start-time of visit m and the end-time of visit m, $R_{ik}(t^E_{imc} - t_{imc})$.

The inventory constraints for the consumption storage tanks are as follows (see Figure 4):

$$s_{i1kc} = S^O_{ik} - R_{ik}t_{i1c} + r_{i1kc}, \qquad i \in N, k \in K : J_{ik} = -1, c \in \Omega, \tag{29}$$

$$\begin{aligned} s^E_{imkc} + r_{imkc} = s_{imkc} + r^E_{imkc} + \sum_{v \in V} q_{imvk} - R_{ik}(t^E_{imc} - t_{imc}), \qquad (i,m) \in S^A, \\ k \in K : J_{ik} = -1, c \in \Omega, \end{aligned} \tag{30}$$

$$\begin{aligned} s_{imkc} + r^E_{i(m-1)kc} = s^E_{i(m-1)kc} + r_{imkc} - R_{ik}(t_{imc} - t^E_{i(m-1)c}), \qquad (i,m) \in S^A : m > 1, \\ k \in K : J_{ik} = -1, c \in \Omega, \end{aligned} \tag{31}$$

$$s^E_{i\mu_i kc} + r^T_{ikc} = r^E_{i\mu_i kc} + s^T_{ikc} + R_{ik}t^+_{ic} + \underline{S}_{ik}, \qquad i \in N, k \in K : J_{ik} = -1, c \in \Omega, \tag{32}$$

$$s_{imkc}, s^E_{imkc} \leq \overline{S}_{ik}, \qquad (i,m) \in S^A, k \in K : J_{ik} = -1, c \in \Omega, \tag{33}$$

$$s_{imkc}, s^E_{imkc}, r_{imkc}, r^E_{imkc} \geq 0, \qquad (i,m) \in S^A, k \in K : J_{ik} = -1, c \in \Omega. \tag{34}$$

$$s^T_{ikc}, r^T_{ikc} \geq 0, \qquad i \in N, k \in K : J_{ik} = -1, c \in \Omega. \tag{35}$$

Equations (29) calculate the inventory level of each product at the first visit. Equations (30) calculate the inventory level of each product when the service ends at port visit (i, m). Equations (31) relate the inventory level at the start of port visit (i, m) to the inventory level at the end of port visit $(i, m - 1)$. Constraints (32) are the inventory balance constraints at time T, or at the end of the last visit if this visit occurs after time T, for each product. When the last visit occurs before T, the inventory level at the last visit needs to be adjusted by the consumption until time T. Additionally the quantity below the lower bound $\underline{S}_{ik}$ is penalized as backlogged demand. Hence, these constraints penalize situations where the storages are below the lower inventory limit at the end of the planning. Constraints (33) ensure that the capacities of the depots are not exceeded. Finally, non-negativity requirements (34), and (35) are imposed on the inventory and backlog variables.

4.3 Objective function

The objective is to minimize the sailing, setup and operating costs plus the penalty for backlogged demand. The storages at both the supply and consumption tanks belong to the same company, so there exists no inventory holding costs in any storage tank. In order to avoid extreme inventory levels in the end of the planning horizon, we included constraints (16).

The sailing cost of ship v from port i to port j is denoted by C^T_{ijv}, while C^{TO}_{oiv} represents the sailing cost of ship v from its initial port position to port i. The operating cost of product k at port i is denoted by C^O_{ik}. The penalty cost for backlogging of product k at port i is denoted C^P_{ik}. The objective function is as follow:

$$\begin{aligned} z = \min \sum_{v \in V} \sum_{(i,m,j,n) \in S^X_v} C^T_{ijv} x_{imjnv} + \sum_{v \in V} \sum_{(i,m) \in S^A_v} C^{TO}_{oiv} x^O_{imv} + \sum_{v \in V} \sum_{(i,m) \in S^A_v} \sum_{k \in K_v | J_{ik} \neq 0} C^O_{ik} o_{imvk} + \\ \sum_{c \in \Omega} \frac{1}{|\Omega|} \left(\sum_{(i,m) \in S^A} \sum_{k \in K_v | J_{ik} = -1} C^P_{ik}(r_{imkc} + r^E_{imkc}) + \sum_{i \in N} \sum_{k \in K_v | J_{ik} = -1} C^P_{ik} r^T_{ikc} \right). \end{aligned} \tag{36}$$

The objective function (36) consists of five main terms. The three first terms minimize the sailing and port operation costs. Then the backlogged demand for each port visit is penalized. Finally we penalize, for each product unloaded in a port, the difference between the lower inventory limit and the actual inventory level for the last visit or time T if the last visit is earlier.

5 Solution approach

Since the complete model is too large to be solved efficiently, it is decomposed into a master problem and one subproblem for each scenario, following the idea of the L-shaped algorithm [11]. Let the problem (1) - (36) be re-written as:

$$
\begin{aligned}
z = \min\ & C(X) + \sum_{c \in \Omega} \frac{1}{|\Omega|} H(X, c) \\
s.t.\ & (1) - (21)
\end{aligned}
$$

where

$$
\begin{aligned}
C(X) = & \sum_{v \in V} \sum_{(i,m,j,n) \in S_v^X} C_{ijv}^T x_{imjnv} + \sum_{v \in V} \sum_{(i,m) \in S_v^A} C_{oiv}^{TO} x_{imv}^O \\
& + \sum_{v \in V} \sum_{(i,m) \in S_v^A} \sum_{k \in K_v | J_{ik} \neq 0} C_{ik}^O o_{imvk}
\end{aligned}
$$

and

$$
\begin{aligned}
H(X, c) = \min\ & \sum_{(i,m) \in S^A} \sum_{k \in K | J_{ik} = -1} C_{ik}^P (r_{imkc} + r_{imkc}^E) + \sum_{i \in N} \sum_{k \in K_v | J_{ik} = -1} C_{ik}^P r_{ikc}^T \\
s.t.\ & (22) - (35), \text{with } \Omega = \{c\}.
\end{aligned}
$$

The master problem consists of the first stage, but with iteratively added variables and constraints to reflect the recourse costs. The subproblems consider fixed first stage decisions, and are solved for each scenario to supply optimality cuts to the master problem.

The problem (1) - (36) has relatively complete recourse, since feasibility in the second stage is guaranteed if the inventory levels do not exceed the capacities of the inventories. Hence, for each feasible solution to the first stage, the second stage has always a feasible solution (it suffices to delay the unloading when necessary). The details for solving the problem are given in Section 5.2. To solve problems with a large number of scenarios, the sample average approximation method is used as described in Section 5.1.

5.1 Sample average approximation method

To solve the SMIRP with many scenarios, we apply the sample average approximation method as in [9]. First we consider M separate sets of scenarios. Each set of scenarios, $i \in \{1, \ldots M\}$ contains a small number of m scenarios, $\{c^{i1}, \ldots, c^{im}\}$. The model (1) - (36) is solved for each set of scenarios i using a decomposition approach. Let X^i denote the obtained first stage solution. The M candidate solutions $X^1, \ldots, X^M$, are then compared using a different, and much larger, set of n scenarios $\{\hat{c}^1, \ldots, \hat{c}^n\}$. The best solution is given by $X^* = argmin\{z_n(X^i) : i \in \{1, \ldots, M\}\}$ where $z_n(X^i) = C(X^i) + \frac{1}{n} \sum_{j=1}^{n} H(X^i, \hat{c}^j)$.

With the first stage solutions $X^1, \dots, X^M$ being obtained, the optimal values are denoted $z_m^i = z_m(X^i) = C(X^i) + \frac{1}{m}\sum_{j=1}^{m} H(X^i, c^{ij})$. The average value over all sets of scenarios, $\bar{z}_m = \frac{1}{M}\sum_{i=1}^{M} z_m^i$ is a statistical estimate for a lower bound on the optimal value of the true problem.

For the larger set of n scenarios, which can be regarded as a benchmark scenario set representing the true distribution (see [20]), the cost $z_n(X^i)$ of each solution $X^i, i \in \{1, \dots, M\}$ is computed as well as $X^* = argmin\{z_n(X^i) : i \in \{1, \dots, M\}\}$. The best value, $z_n(X^*)$, is a statistical estimate for an upper bound on the optimal value. The estimated optimality gap (GAP) is given by $GAP(\%) = \frac{z_n(X^*) - \bar{z}_m}{\bar{z}_m} \times 100(\%)$.

When employing a scenario generation method it is desirable that no matter which set of scenarios is used, by solving the two-stage model, one obtains approximately the same value for the optimal solution. This is named as stability requirement conditions in [20]. Here we evaluate stability, following [9], by computing the variances:

$$\hat{\sigma}^2_{z_n(X^*)} = \frac{1}{(n-1)n} \sum_{j=1}^{n} \left(C(X^*) + H(X^*, \hat{c}^j) - z_n(X^*) \right)^2, \tag{37}$$

$$\hat{\sigma}^2_{\bar{z}_m} = \frac{1}{(M-1)M} \sum_{i=1}^{M} (z_m^i - \bar{z}_m)^2, \tag{38}$$

where $\hat{\sigma}^2_{\bar{z}_m}$ is the variance between samples and $\hat{\sigma}^2_{z_n(X^*)}$ is the variance within the larger sample. The estimated variance of the estimated optimality gap is

$$\hat{\sigma}^2_G = \hat{\sigma}^2_{z_n(X^*)} + \hat{\sigma}^2_{\bar{z}_m}.$$

5.2 Optimization process

To solve the model (1) - (36) for a set of scenarios Ω, we first solve to optimality a master problem including only one scenario. Since the problem has relatively complete recourse, the resulting values for the first stage decision variables are feasible for the complete problem with all scenarios. However, we need to check whether the solution is optimal for the complete model. To do that we check, for each scenario, whether there is backlogged demand when the deliveries are made as early as possible. If such a scenario with backlogged demand is found, we add to the master problem additional variables and constraints (which are implied by the time constraints and inventory constraints) enforcing the backlogged demand to be counted in the objective function. Then the revised master problem is solved again, and the process is repeated until all the optimality constraints are satisfied. Hence, as in the L-shaped method, the master problem initially disregards the recourse cost, and an improved estimation of the recourse cost is gradually added to the master problem by solving optimality subproblems and adding the corresponding cuts. The algorithm may also be terminated if the additional recourse cost added in an iteration is less than a given small amount ϵ. A formal description of this process is given below.

Algorithm 1 Optimization procedure for an input set of scenarios Ω.

```
Choose a scenario $c \in \Omega$
Solve the master problem for scenario $c$
while There are new violated optimality cuts and a change in the objective function greater than $\epsilon$ do
    Add all the violated optimality cuts
    Solve again the master problem with the new cuts
end while
```

Next we explain how separation of constraints imposing backlog for each scenario (optimality cuts) is done in each iteration.

The backlog variables are bounded as follows:

$$
\begin{aligned}
r_{imkc} &\geq R_{ik}t_{imc} - S^O_{ik} - \sum_{n\leq m-1}\sum_{v\in V} q_{invk}, (i,m) \in S^A, k \in K : J_{ik} = -1, c \in \Omega, && (39)\\
r^E_{imkc} &\geq R_{ik}t^E_{imc} - S^O_{ik} - \sum_{n\leq m}\sum_{v\in V} q_{invk}, (i,m) \in S^A, k \in K : J_{ik} = -1, c \in \Omega, && (40)\\
r^T_{ikc} &\geq R_{ik}t^E_{i\mu_i c} + R_{ik}t^+_{ic} + \underline{S}_{ik} - S^O_{ik} - \sum_{n\leq \mu_i}\sum_{v\in V} q_{invk}, i \in N, k \in K : J_{ik} = -1, c \in \Omega. && (41)
\end{aligned}
$$

Constraints (39) - (40) are implied by (29) - (31) (adding alternately (31) and (30) from (i,m) to $(i,1)$ and then (29)), and from the non-negativity requirements on the inventory variables (34). Constraints (41) are implied by (29) - (32) and by (35).

The minimum backlog occurs when the time variables t_{imc} and t^E_{imc} are set to the earliest feasible times. Once these variables are defined, separation over (39) - (40) is trivial since the right hand side is fixed. So we focus now on finding tight bounds for the time variables. First observe that the starting and ending times of each operation are established either from the (maximum) inventory levels (inventory constraints) or from the duration of the several operations the ships perform (time constraints). In the first case we need to ensure that the inventory capacity is not exceeded. Hence we have:

$$
\begin{aligned}
t_{imc} &\geq \frac{S^O_{ik} + \sum_{n\leq m-1}\sum_{v\in V} q_{invk} - \overline{S}_{ik}}{R_{ik}}, \qquad (i,m) \in S^A, k \in K : J_{ik} = -1, c \in \Omega, && (42)\\
t^E_{imc} &\geq \frac{S^O_{ik} + \sum_{n\leq m}\sum_{v\in V} q_{invk} - \overline{S}_{ik}}{R_{ik}}, \qquad (i,m) \in S^A, k \in K : J_{ik} = -1, c \in \Omega. && (43)
\end{aligned}
$$

Constraints (42) and (43) follow from (29) - (31), (33), (34) (adding alternately (31) and (30) from (i,m) to $(i,1)$ and then (29), and then lower bounding the backlog variables from the non-negativity requirements (34) and the inventory variables from the storage tank capacity constraints (33)). For a given feasible solution for the first stage, the right hand sides of (42) and (43) are constant.

For the second stage, the time variables are determined from the time constraints (22) - (25). For a feasible solution of the first stage, and for each scenario, most of the constraints (22) - (25) are not tight and many variables do not need to be considered. We can see that the t^E_{imc}-variables are bounded by (22) while the t_{imc}-variables are bounded by (23) (from the end time of the previous visit to the same port), by (24) (from the last ship operation), and by (25) (from the origin). These cases can be represented in a network $\mathcal{G} = (N, A, W)$, where N is the set of nodes, A is the set of arcs and W is the set of weights. The set of nodes P is given by the origin of each ship, represented by O_v, a node (i, m) representing the starting time of each port visit and a node $\overline{(i, m)}$ representing the end time of the visit. Each arc in A corresponds to a routing variable set to one. That is, there is an arc from node O_v to node (i, m) if $x^O_{imv} = 1$, and there is an arc from node $\overline{(i, m)}$ to node (j, n) if $x_{imjnv} = 1$ for some v. The arcs have weights $T^O_{ivc} + T^W_{imc}$ and $T_{ijvc} + T^W_{jnc}$, respectively. There is an arc from node (i, m) to node $\overline{(i, m)}$ with weight $\sum_{k\in K_v} T^S_{ik} o_{imvk} + \sum_{k\in K_v} T^Q_{ik} q_{imvk}$, and there is an arc from node $\overline{(i, m)}$ to node $(i, m + 1)$ with weight T^B_i. Finally, we consider an arc from O_v to each node visited by ship v. The weight from O_v to (i, m) is given by the right hand side of (42) and the weight from O_v to $\overline{(i, m)}$ is given by the right hand side of (43).

The weight of each path from one origin to a node gives a lower bound for the time variable corresponding to that node. Hence the earliest time associated to a node corresponds to the weight of the longest path from one origin to that node (one can always establish an artificial origin which is linked to all ship origins O_v and with null weight). Since the graph is acyclic, finding the longest path to each node can be done in polynomial time. However, for this particular graph, it is easy to derive a linear labeling correcting algorithm.

The time variables can then be restricted using these paths or sub-paths. For each (sub)path $\Pi^{(i,m)}_{(j,n)}$, from visit (j, n) to visit (i, m) of a ship v, we define the set of nodes (port visits) as $\mathcal{N}(\Pi^{(i,m)}_{(j,n)})$ and the set of arcs as $\mathcal{A}\left(\Pi^{(i,m)}_{(j,n)}\right)$. Let (i_v, m_v) denote the first visit of ship v after leaving the origin.

If the earliest time for a visit $(i, m) \in S^A$ is determined only by the schedule of operations for a given ship $v \in V$, then t_{imc} and t^E_{imc} are restricted as follows:

$$\begin{aligned}
t_{imc} \quad \geq \quad & \sum_{(\ell,u)\in\mathcal{N}(\Pi^{(i,m)}_{(i_v,m_v)})} T^W_{\ell uc} + \sum_{(\ell,u)\in\mathcal{N}(\Pi^{(i,m)}_{(i_v,m_v)})\backslash\{(i,m)\}} \sum_{k\in K} \left(T^S_{\ell k} o_{\ell uvk} + T^Q_{\ell k} q_{\ell uvk}\right) \\
& + T^O_{i_v vc} + \sum_{(\ell,u,t,w)\in\mathcal{A}(\Pi^{(i,m)}_{(i_v,m_v)})} T_{\ell tvc} \\
& - T\left(1 + \mid \mathcal{A}(\Pi^{(i,m)}_{(i_v,m_v)}) \mid - x^O_{i_v m_v v} - \sum_{(\ell,u,t,w)\in\mathcal{A}(\Pi^{(i,m)}_{(i_v,m_v)})} x_{\ell utwv}\right), \qquad (44) \\
t^E_{imc} \quad \geq \quad & \sum_{(\ell,u)\in\mathcal{N}(\Pi^{(i,m)}_{(i_v,m_v)})} T^W_{\ell uc} + \sum_{(\ell,u)\in\mathcal{N}(\Pi^{(i,m)}_{(i_v,m_v)})} \sum_{k\in K} \left(T^S_{\ell k} o_{\ell uvk} + T^Q_{\ell k} q_{\ell uvk}\right) \\
& + T^O_{i_v vc} + \sum_{(\ell,u,t,w)\in\mathcal{A}(\Pi^{(i,m)}_{(i_v,m_v)})} T_{\ell tvc} \\
& - T\left(1 + \mid \mathcal{A}(\Pi^{(i,m)}_{(i_v,m_v)}) \mid - x^O_{i_v m_v v} - \sum_{(\ell,u,t,w)\in\mathcal{A}(\Pi^{(i,m)}_{(i_v,m_v)})} x_{\ell utwv}\right). \qquad (45)
\end{aligned}$$

Validity of (44) and (45) is implied by (22) - (24). In Appendix 1 we provide a list of the remaining inequalities defined for each possible subpath.

The overall separation procedure for each iteration works as follow:

Algorithm 2 Separation procedure

```
Construct the network $\mathcal{G} = (N, A, W)$
Determine the longest path from the origin to each node
Associate the corresponding time variables to each node
Set the backlog variables to the minimum value using (39)-(41)
for each node do
    if the corresponding backlog variable has a value strictly greater than its value in the current solution then
        add inequalities belonging to (39)-(41) determining its value
        add time constraints (42)-(45) corresponding to the weight of the longest path
        (Use subpaths of each ship that are not contained in other subpaths of the same ship in the critical path)
    end if
end for
```

Example 5.1. *Consider an instance with 2 ships, v_1 and v_2, and 3 ports and assume that there is only one scenario. Hence we omit the corresponding scenario index from all variables and parameters. Let the paths resulting from the first stage solution be* $x^O_{11v_1} = x_{1132v_1} = 1$ *and* $x^O_{21v_2} = x_{2131v_2} = x_{3112v_2} = 1$. *Assume the weights of the arcs are those given in Figure 5. For instance,* $T^B_1 = T^B_3 = 0.5$, $\sum_{k \in K} \left(T^S_{1k} o_{11v_1k} + T^Q_{1k} q_{11v_1k} \right) = 1$ *and* $T_{13v_1} = 6$.

For simplicity we omit arcs with weights resulting from (42) *and* (43).

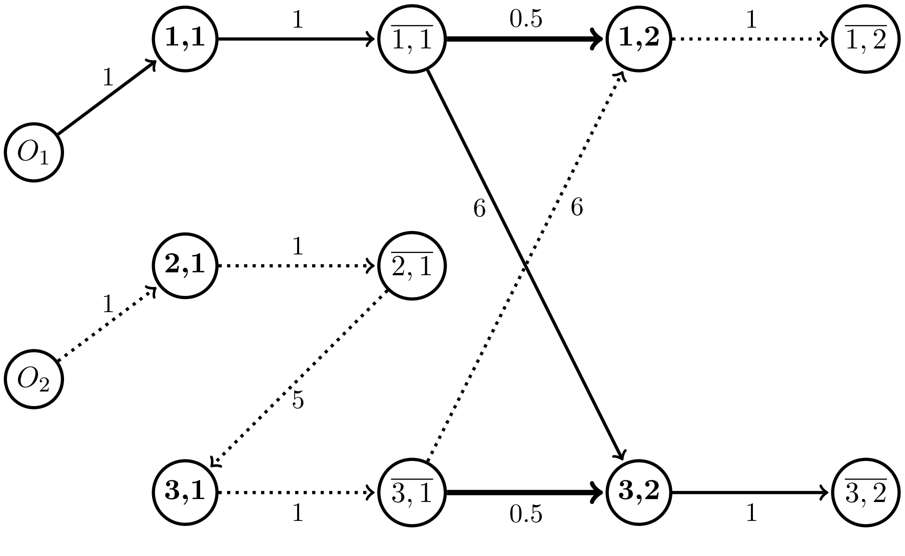

Figure 5: Example of a graph $\mathcal{G}$ for a set of three ports and two ships.

We can see that $t_{11} = 1, t^E_{11} = 2, t_{21} = 1, t^E_{21} = 2, t_{31} = 7, t^E_{31} = 8, t_{12} = 14, t^E_{12} = 15, t_{32} = 8.5, t^E_{32} = 9.5$.

Suppose there is backlog at nodes $(1,2)$ *and* $(3,2)$. *In addition to inequalities* (39) *defining* *the lower bound on the backlog for $(i,m) = (1,2)$ and $(i,m) = (3,2)$, the following inequalities, corresponding to the critical paths to nodes $(1,2)$ and $(3,2)$, are added to limit the time variables:*

$$\begin{aligned}
t^E_{31} &\geq \sum_{k\in K}\left(T^S_{2k}o_{21v_2k} + T^Q_{2k}q_{21v_2k}\right) + \sum_{k\in K}\left(T^S_{3k}o_{31v_2k} + T^Q_{3k}q_{31v_2k}\right)\\
&\quad +T^O_{2v_2} + T_{23v_2} - T(2 - x^O_{21v_2} - x_{2131v_2}),\\
t_{12} &\geq \sum_{k\in K}\left(T^S_{2k}o_{21v_2k} + T^Q_{2k}q_{21v_2k}\right) + \sum_{k\in K}\left(T^S_{3k}o_{31v_2k} + T^Q_{3k}q_{31v_2k}\right)\\
&\quad +T^O_{2v_2} + T_{23v_2} + T_{31v_2} - T(3 - x^O_{21v_2} - x_{2131v_2} - x_{3112v_2}),\\
t_{32} &\geq t^E_{31} + T^B_3.
\end{aligned}$$

6 Stochastic times and sample scenarios generation

In the SMIRP, the sailing and waiting times at ports are assumed to be random, following known probability distributions. We now describe the distributions used and how scenarios are generated for the stochastic programming model.

For the sailing times we assume that there are three potential events that affect all the sailing times simultaneously. These correspond to “good weather”, “moderate weather” and “bad weather”. For good weather, the sailing times are obtained directly from the sailing distance and the ship speed. For moderate weather, the sailing times are 1.5 times the corresponding sailing times in good weather, and for bad weather the sailing times are 2.0 times those in good weather. From the historical data for the season we are considering, a probability is associated to each event.

Contrary to the sailing times, where the weather usually affects all the islands simultaneously, waiting times due to port occupancy depend only on the port. For each visit to each port, we assume that the random variable indicating whether the port is occupied or not follows a Bernoulli distribution with parameter $p \in [0,1]$ (p is the probability of the port being occupied). If the port is occupied then the random variable W indicating the waiting time is given by a truncated exponential distribution

$$F(w) = \begin{cases} 0, & w < 0,\\ (1-e^{-\lambda w})/A, & 0 \leq w \leq B,\\ 1, & w > B, \end{cases}$$

where $A = 1 - e^{-B\lambda}$, and λ is such that the expected value of waiting time is $1/\lambda - \frac{Be^{-B\lambda}}{A}$, and B represents the maximum waiting time (see [6]). Parameters p, λ, and B are obtained from historical data.

The weather events are trivially generated using the given probabilities. For each visit to each port the waiting times are randomly generated as follows: let $p \in [0,1]$ be the probability of the port being occupied. Generate an uniform random variable $U_1 \in [0,1]$. If $U_1 > p$ we assume that the port is not occupied. Otherwise we randomly generate a waiting time from the truncated exponential distribution using the inverse transformation method. The waiting time is given by

$$W = \begin{cases} \frac{ln(1-AU_2)}{-\lambda}, & \text{if } U_1 \leq p;\\ 0, & \text{if } U_1 > p. \end{cases}$$

where U_2 is an uniform random variable, $U_2 \in [0, 1]$.

To generate the set of scenarios, Ω, we first fix the number of scenarios $n =\mid \Omega \mid$ *a priori*. Then each scenario is generated separately, first by generating the sailing times at random and then by generating a random waiting time for each port visit.

7 Computational results

In this section we report the results from the computational experimentation conducted to test the stochastic model. All computations were performed using the optimization software Xpress Optimizer Version 20.00.05 with Xpress Mosel Version 3.0.0, on a computer with processor Intel Core 2 Duo 2.2GHz and with 4GB of RAM. In Algorithm 1 we use $\epsilon = 0.01$. Ten real-world instances are used in the main testing, considering two different ships, seven ports, four products, and a time horizon of eight days. The instances differ by the initial inventory levels. Detailed information on the instances is given in Appendix 2.

First we test the effectiveness of the decomposition method, through a comparison by solving the full stochastic programming model directly using commercial software. Then, we test the sample average approximation method using the decomposition method and we compute estimations of the Value of the Stochastic Solution and the Expected Value of Perfect Information. Finally, the scalability of this approach is tested by increasing the length of the time horizon, number of ships and number of ports.

7.1 Effectiveness of decomposed model

To test the effectiveness of the decomposed model we compared its performance with the use of Xpress Optimizer to directly solve the stochastic programming model with 10 scenarios. The results are reported in Table 1. The column “Opt” gives the optimal values, the columns “Nodes” indicate the number of branch and bound nodes, the columns “Seconds” report the running time in seconds to solve the instance. For the decomposed model we report additionally the number of cuts added in the column “Ncuts” and the number of iterations in the column “Iterations”, that is, the number of times we solve the separation problem to add backlog and time constraints.

Table 1: Effectiveness of the decomposition approach

		full model		**decomposed model**			
Instance	Opt	Nodes	Seconds	Nodes	Seconds	Ncuts	Iterations
1	16210	5888	1498	3503	390	20	3
2	17610	20292	5397	16436	1061	24	4
3	18500	8495	2111	3863	434	65	3
4	17248.6	9253	1644	5377	526	78	4
5	15410	8177	2284	5356	384	18	3
6	18576.8	42774	7799	7247	854	32	4
7	15362.3	20720	4546	6658	603	45	4
8	17008	27740	5564	7558	740	28	4
9	13330	1911	362	1462	146	16	3
10	14550	46407	9200	3821	351	25	4
Average	16380.57	19165.7	4040.5	6128.1	548.9	35.1	3.6

As expected, the running times of the decomposition method are much lower than the running times obtained by solving the complete model. Additionally, we can see that the number of times the separation problem is called is at most 4 and few cuts are added.

7.2 Testing different sizes of sets of scenarios

Next we follow the solution approach described in Section 5, see [9]. Each instance is solved for M independent sets of scenarios, each set i containing m scenarios.

We conducted tests for $m = 10$ and $m = 50$. In all cases we consider $M = 10$ and the solutions are evaluated using a bigger set of $n = 1000$ scenarios. For each value of m we give two tables (Tables 2 and 3 for $m = 10$ and Tables 4 and 5 for $m = 50$). In the first table we present, for each instance, $z_n(X^*)$, $\bar{z}_m$, GAP, $\hat{\sigma}_{z_n(X^*)}$, $\hat{\sigma}_{\bar{z}_m}$, $\hat{\sigma}_G$. In the second table we give, for each instance, the average running time to solve the M problems (one problem for each set of scenarios of size m) using the decomposition method, "Seconds M", and the average time to compute $z_n(X^k), k \in \{1, \ldots, M\}$ "Seconds n", the average number of iterations, "Iterations", to solve the M problems, that is, the average number of times we solve the separation problem, and the average number of cuts (39)-(41) added, "Cuts".

Table 2: Bounds and variances for $m = 10$.

Instance	$z_n(X^*)$	$\bar{z}_m$	**GAP** (%)	$\hat{\sigma}_{z_n(X^*)}$	$\hat{\sigma}_{\bar{z}_m}$	$\hat{\sigma}_G$
1	16956.8	16358.0	3.7	24.12	76.68	80.38
2	19080.3	18516.9	3.0	47.19	173.77	180.06
3	21150.9	19660.2	7.6	95.57	265.06	281.76
4	19613.9	18750.8	4.6	198.39	293.10	353.93
5	18813.2	16658.5	12.9	72.32	194.24	207.27
6	21182.1	19743.3	7.3	104.68	210.69	235.26
7	16694.8	16509.5	1.1	76.59	196.75	211.13
8	19325.2	18664.2	3.5	70.96	227.79	238.59
9	14335.6	14139.0	1.4	115.00	238.07	264.47
10	17636.8	16721.6	5.5	127.03	324.51	348.49
Average	18479.0	17572.2	5.1	93.20	220.10	240.1

Table 3: Average times, average number of iterations and average number of cuts for $m = 10$.

Instance	**Seconds M**	**Seconds n**	**Iterations**	**Cuts**
1	77.0	35.0	3	16.4
2	180.5	38.5	4	16.4
3	95.6	45.7	3	16.4
4	104.9	42.5	3.4	18.4
5	118.3	41.2	3	15.2
6	181.7	50.9	3	14.4
7	133.5	29.3	3	13.9
8	117.7	54.3	3.2	19.6
9	37.1	42.9	3	20.2
10	123.1	40.7	3	16.3
Average	116.9	42.1	3.2	16.7

Table 4: Bounds and variances for $m = 50$.

Instance	$z_n(X^*)$	$\bar{z}_m$	**GAP** (%)	$\hat{\sigma}_{z_n(X^*)}$	$\hat{\sigma}_{\bar{z}_m}$	$\hat{\sigma}_G$
1	16498.1	16352.4	0.9	11.51	44.18	45.65
2	19080.3	18375.8	3.8	47.19	219.28	224.30
3	20839.2	19542.5	6.6	97.57	141.96	172.26
4	19401.7	18027.2	7.6	198.39	241.19	312.30
5	18490.2	16160.0	14.4	66.21	119.28	136.42
6	20177.1	19651.1	2.7	104.48	96.41	142.17
7	16358.8	16204.4	1.0	76.59	98.48	124.76
8	18148.3	18015.3	0.7	70.96	198.24	210.56
9	13877.7	13876.4	0.0	115.18	116.46	163.80
10	16875.0	16336.1	3.3	127.03	286.35	313.26
Average	17974.6	17254.1	4.1	91.5	156.2	184.5

Table 5: Average times, average number of iterations and average number of cuts for $m = 50$.

Instance	**Seconds M**	**Seconds n**	**Iterations**	**Cuts**
1	426.5	51.0	3.3	18.8
2	838.2	64.8	3	16.5
3	658.3	59.4	3.1	18.9
4	682.8	57.8	3	20.1
5	627.2	63.9	3	16.8
6	751.9	50.2	4.2	16.8
7	573.3	76.1	3	16.3
8	791.0	49.5	3	19.8
9	174.5	51.5	4	18.8
10	524.3	56.0	3	17.8
Average	604.8	58.0	3.3	18.1

We can see that increasing m, the cost of the selected solution ($z_n(X^*)$) decreases in average by 2.8%. Also, the standard deviations $\hat{\sigma}_{z_n(X^*)}$, $\hat{\sigma}_{\bar{z}_m}$, $\hat{\sigma}_G$ and the gaps have a reduction. The price to pay for the improvement of the solution and reduction of variability is an increase in the average running times. The running time is, on average, approximately 2 minutes for $m = 10$, and increases to 10 minutes for $m = 50$.

7.3 Importance of a stochastic approach

To evaluate the importance of the stochastic approach we compute estimations of the Value of the Stochastic Solution (VSS) and the Expected Value of Perfect Information (EVPI). The results are given in Table 6. In column $z_n(X^*)$ we give the corresponding value for $m = 50$. To compute the VSS we solve the model with one scenario, where the stochastic parameters are set to their expected values. We used the sample average values (considering the larger sample), which are very similar to the theoretical expected values. Solving this deterministic model we obtain the well known expected value solution (VSS). The cost of this solution is given in column "EVS". In column "VSS" we give an estimation of the Value of the Stochastic Solution which is the difference between EVS and $z_n(X^*)$. In column "PI" we give the average value of the $n = 1000$ deterministic models, one model for each scenario, and in column "EVPI" we give an estimation of the Expected Value of Perfect Information which is the difference $z_n(X^*) - PI$.

Table 6: Estimating the VSS and EVPI.

Instance	$z_n(X^*)$	EVS	VSS	PI	EVPI
1	16498.1	42049.2	25551.1	16210.1	288.0
2	19080.3	33020.1	13939.8	17620.2	1460.1
3	20839.2	39871.2	19032.0	18572.6	2266.6
4	19401.7	49582.5	30180.8	17285.9	2115.8
5	18490.2	58053.6	39563.4	15461.7	3028.5
6	20177.1	41503.3	21326.2	18942.0	1235.1
7	16358.8	32256.6	16052.2	15497.8	706.6
8	18148.3	64144.0	45995.7	17023.8	1124.5
9	13877.7	41125.7	27248.0	13354.0	523.7
10	16875.0	25623.9	8748.9	15066.9	1808.1
Average	17974.6	42723.01	24763.8	16503.5	1455.7

We can see, from Table 6, that the gains for using stochastic programming instead of the deterministic model based on expected values are in general very high. In average, the value of the best solution is only 9% above the value of Perfect Information (PI).

7.4 Scalability study

To study scalability of the stochastic approach proposed we tested the decomposition method for larger instances, that can be seen as future realities in Cape Verde. For these tests we used an intermediate dimension of $m = 30$ scenarios. The results obtained for the 10 unmodified instances are given in Table 7. Table 8 shows the results for modified instances where the number of ports is increased from 7 to 9, corresponding to a situation where all the inhabited islands are considered simultaneously in the distribution planning. The average running time and the gaps increase, but nevertheless, each master problem was solved to optimality. One of the reported gaps is negative, which can happen as the optimality gap is only an estimation [9].

In Table 9 we report the results for another possible future reality, where the demand rates are increased by 50% at each port and four ships instead of two are considered. As a result, the set of possible port visits is approximately doubled. As the problem gets harder we added a time limit of 600 seconds to each iteration and the overall running time for each instance, and for each set of the $M = 30$ scenarios, was limited to 2400 seconds. Additionally, we also needed to introduce a new stopping criteria in Algorithm 1 by bounding the number of iterations to five, since for the previous instances solved we observed that in general 3 - 4 iterations were enough. In this case none of the solutions obtained are optimal, and both the running time limit of 2400 seconds and the maximum number of iterations for the subproblems are reached. Hence, the procedure proposed acts as a heuristic, and in these cases the estimated optimality gaps are negative and the standard deviations are very large. As the gaps and standard deviations provide no insight, we report the EVS, VSS, PI, and EVPI in addition to the gaps but omit the standard deviations. The average running times given previously in columns Seconds M and the number of iterations are also omitted since the corresponding maximum values are attained in all instances. Finally, Table 10 shows results for a situation considering a time horizon of 15 days instead of eight days. Again, the number of possible port visits is doubled, and the problem is much harder to solve. In both cases with an increased number of port visits, the solution method only provides feasible solutions. Further research, including improving the mathematical model by including valid inequalities and test new heuristic approaches could be followed, but such lines of research are out of the scope of this paper.

Table 7: Computational results considering $m = 30$, with unmodified instances.

Inst.	$z_n(X^*)$	$\bar{z}_m$	**GAP** (%)	$\hat{\sigma}_{z_n(X^*)}$	$\hat{\sigma}_{\bar{z}_m}$	$\hat{\sigma}_G$	**Sec. M**	**Sec. n**	**It.**	**Cuts**
1	16940.0	16353.1	3.6	0.0	218.1	218.1	88.0	42.8	3	15.6
2	19080.3	18467.6	3.3	118.9	187.3	221.8	153.3	50.9	3	16.5
3	19890.5	19611.6	1.4	136.3	61.0	149.3	175.5	52.2	3	15.8
4	19376.3	18205.0	6.4	198.5	237.6	309.6	188.2	48.4	3	21
5	18746.7	16607.2	12.9	77.3	71.1	105.0	192.5	49.8	3	17.2
6	21879.4	19704.2	11.0	63.4	71.5	95.6	123.1	48.7	3.2	14.2
7	16538.3	16380.4	1.0	48.2	77.6	91.3	82.0	49.2	3	15.3
8	18648	18198.1	2.5	164.7	204.3	262.4	94.2	60.1	3.2	21.2
9	14183.2	14112.2	0.5	74.0	209.4	222.1	49.7	49.1	3	20
10	17335.5	16651.7	4.1	40.3	399.7	401.7	79.4	49.4	3	16.2
Av.	18261.8	17429.1	4.7	92.2	173.8	207.7	122.6	50.1	3.0	17.3

Table 8: Computational results considering $m = 30$ and 9 ports.

Inst.	$z_n(X^*)$	$\bar{z}_m$	**GAP** (%)	$\hat{\sigma}_{z_n(X^*)}$	$\hat{\sigma}_{\bar{z}_m}$	$\hat{\sigma}_G$	**Sec. M**	**Sec. n**	**It.**	**Cuts**
1	16331.1	16275.5	0.3	33.8	66	74.2	168.2	43.4	3.1	18.7
2	19821	17343.31	14.3	179.8	217	281.8	428.3	50.6	3	16.7
3	21778.1	17368.51	25.4	147.8	67	162.3	4729.8	45	3.6	26.8
4	21668.6	16699.71	29.8	150.7	162.8	221.8	1308.2	43.2	3.6	25.2
5	17387.8	16918.69	2.8	53	210.5	217.1	449.2	55.4	3.2	18.9
6	20685.8	16812.74	23.0	91.5	74.8	118.2	426.5	53.4	3.1	14.9
7	17591.1	16789.83	4.8	18.8	17.9	26	379.2	55.6	3.1	11.9
8	24740.6	23474	5.4	296.3	207.6	361.8	2038.4	39.4	4.5	34.6
9	15306.2	15476.5	-1.1	99.3	231.3	251.7	76.3	38.5	3	21.5
10	17255.3	17172.3	0.5	7.3	13.2	15.1	200.1	45	3	15.6
Av.	19256.6	17433.1	10.5	107.8	126.8	173.0	1020.4	47.0	3.3	20.5

8 Conclusions

We have presented a two-stage stochastic programming model with recourse for a maritime inventory routing problem where sailing times and waiting times in port are random. The model has relatively complete recourse, indicating that for each scenario, a feasible solution to the second stage can always be found as long as the first stage solution is feasible. We proposed a decomposition method where, for a given first stage solution, optimality is checked for the complete model through an efficient separation method.

Ten instances based on real data from oil distribution at the Cape Verde islands are solved using the sample approximation method. Computational tests have shown the effectiveness of the decomposition method, and the importance of stochastic programming instead of a deterministic approach. In addition, larger instances are tested study the scalability of the solution approach. The instances become much harder to solve with increased demand and time horizon probably due to the increased number of visits to each port.

Table 9: Computational results considering $m = 30$, 4 ships and an increase of demand of 50%.

Inst.	$z_n(X^*)$	$\bar{z}_m$	**GAP** (%)	**EVS**	**VSS**	**PI**	**EVPI**	**Sec. n**	**Cuts**
1	20881.2	24094.7	-13.3	79305.7	58424.5	19284.3	1596.9	65	41.1
2	26410.5	43286.8	-39.0	36906.8	10496.3	20952.5	5458.0	68.6	51.3
3	30398.3	41360.5	-26.5	43685.2	13286.9	23224.1	7174.2	69.1	51.5
4	29909.3	41529.0	-28.0	86177.4	56268.1	23899.8	6009.5	70.3	60.5
5	26842.4	39513.0	-32.1	36572.7	9730.3	21661.1	5181.3	71.1	47.9
6	30115.5	55669.2	-45.9	37622.1	7506.6	26644.3	3471.2	73	60.2
7	26193.5	46872.4	-44.1	48606.5	22413.0	21580.2	4613.3	68.3	38.3
8	23560.2	24961.4	-5.6	21713.4	-1846.8	20493.1	3067.2	76	45.1
9	26365.6	50617.8	-47.9	59275.9	32910.3	21408.1	4957.5	67.9	48.9
10	26845.5	44631.7	-39.9	35867.3	9021.8	23800.8	3044.7	74.8	55.8
Av.	26752.2	41253.7	-32.2	48573.3	21821.1	22294.8	4457.4	70.4	50.1

Table 10: Computational results considering $m = 30$, and a time horizon of 15 days.

Inst.	$z_n(X^*)$	$\bar{z}_m$	**GAP** (%)	**EVS**	**VSS**	**PI**	**EVPI**	**Sec. n**	**Cuts**
1	44214.1	76173.8	-42.0	117381.2	73167.1	36153.5	8060.6	66.2	52.4
2	33447.2	42104.6	-20.6	57057.7	23610.5	33170.1	277.1	91.9	39.6
3	36646.7	86994.6	-57.9	95638.2	58991.5	33770.2	2876.5	81.6	44
4	32948.5	55697.3	-40.8	78427.9	45479.4	32942.1	6.4	87.4	40.6
5	37187.7	53273.1	-30.2	76047.2	38859.5	34530.0	2657.7	75.8	48.8
6	39283.3	55622.8	-29.4	92705.1	53421.8	33510.0	5773.3	79.6	53.4
7	38454.5	182426	-78.9	72285.6	33831.1	35610.0	2844.5	80.6	43.1
8	37167.2	47438.1	-21.7	44431.8	7264.6	36046.5	1120.7	74.5	49.7
9	35518.5	42995.2	-17.4	52212.1	16693.6	33185.1	2333.4	77.5	56.5
10	37810.8	87438.4	-56.8	63678.2	25867.4	36957.5	853.3	74.1	57.4
Av.	37267.9	73016.4	-39.6	74986.5	37718.6	34780.9	2487.0	78.9	48.6

Acnowledgements

The work of the first and third authors was supported by the Portuguese Foundation for Science and Technology, FCT, and by European funds from FEDER through COMPETE, within projects PEst-OE/MAT/UI4106/2014 and EXPL/MAT-NAN/1761/2013. The second and forth authors were supported financially from the Research Council of Norway through the DOMinant II-project.

Appendix 1

The following inequalities, for each $(i,m) \in S^A$, $v \in V$ and $c \in \Omega$, are implied by (22) - (24):

$$\begin{aligned} t_{imc} \;\geq\;& t_{jnc} + \sum_{(\ell,u)\in\mathcal{N}(\Pi^{(i,m)}_{(j,n)})\setminus\{(j,n)\}} T^W_{\ell uc} + \sum_{(\ell,u)\in\mathcal{N}(\Pi^{(i,m)}_{(j,n)})\setminus\{(i,m)\}} \sum_{k\in K} \left(T^S_{\ell k} o_{\ell uvk} + T^Q_{\ell k} q_{\ell uvk}\right) \\ &+ \sum_{(\ell,u,t,w)\in\mathcal{A}(\Pi^{(i,m)}_{(j,n)})} T_{\ell tvc} - T\left(\mid \mathcal{A}(\Pi^{(i,m)}_{(j,n)}) \mid - \sum_{(\ell,u,t,w)\in\mathcal{A}(\Pi^{(i,m)}_{(j,n)})} x_{\ell utwv} \right), \end{aligned} \tag{46}$$

$$\begin{aligned} t_{imc} \;\geq\;& t^E_{jnc} + \sum_{(\ell,u)\in\mathcal{N}(\Pi^{(i,m)}_{(j,n)})\setminus\{(j,n)\}} T^W_{\ell uc} + \sum_{(\ell,u)\in\mathcal{N}(\Pi^{(i,m)}_{(j,n)})\setminus\{(j,n),(i,m)\}} \sum_{k\in K} \left(T^S_{\ell k} o_{\ell uvk} + T^Q_{\ell k} q_{\ell uvk}\right) \\ &+ \sum_{(\ell,u,t,w)\in\mathcal{A}(\Pi^{(i,m)}_{(j,n)})} T_{\ell tvc} - T\left(\mid \mathcal{A}(\Pi^{(i,m)}_{(j,n)}) \mid - \sum_{(\ell,u,t,w)\in\mathcal{A}(\Pi^{(i,m)}_{(j,n)})} x_{\ell utwv} \right), \end{aligned} \tag{47}$$

$$\begin{aligned} t^E_{imc} \;\geq\;& t_{jnc} + \sum_{(\ell,u)\in\mathcal{N}(\Pi^{(i,m)}_{(j,n)})\setminus\{(j,n)\}} T^W_{\ell uc} + \sum_{(\ell,u)\in\mathcal{N}(\Pi^{(i,m)}_{(j,n)})} \sum_{k\in K} \left(T^S_{\ell k} o_{\ell uvk} + T^Q_{\ell k} q_{\ell uvk}\right) \\ &+ \sum_{(\ell,u,t,w)\in\mathcal{A}(\Pi^{(i,m)}_{(j,n)})} T_{\ell tvc} - T\left(\mid \mathcal{A}(\Pi^{(i,m)}_{(j,n)}) \mid - \sum_{(\ell,u,t,w)\in\mathcal{A}(\Pi^{(i,m)}_{(j,n)})} x_{\ell utwv} \right), \end{aligned} \tag{48}$$

$$\begin{aligned} t^E_{imc} \;\geq\;& t^E_{jnc} + \sum_{(\ell,u)\in\mathcal{N}(\Pi^{(i,m)}_{(j,n)})\setminus\{(j,n)\}} T^W_{\ell uc} + \sum_{(\ell,u)\in\mathcal{N}(\Pi^{(i,m)}_{(j,n)})\setminus\{(j,n)\}} \sum_{k\in K} \left(T^S_{\ell k} o_{\ell uvk} + T^Q_{\ell k} q_{\ell uvk}\right) \\ &+ \sum_{(\ell,u,t,w)\in\mathcal{A}(\Pi^{(i,m)}_{(j,n)})} T_{\ell tvc} - T\left(\mid \mathcal{A}(\Pi^{(i,m)}_{(j,n)}) \mid - \sum_{(\ell,u,t,w)\in\mathcal{A}(\Pi^{(i,m)}_{(j,n)})} x_{\ell utw} \right). \end{aligned} \tag{49}$$

Appendix 2

In this appendix we provide detailed information on the data used in the paper.

Table 11: Distances, in nautical miles, between ports.

Port \ Port	1	2	3	4	5	6	7
1	0	8	48	121	142	152	133
2	8	0	44	119	130	160	130
3	48	44	0	86	88	119	92
4	121	119	86	0	37	116	144
5	142	130	88	37	0	83	121
6	152	160	119	116	83	0	61
7	133	130	92	144	121	61	0

Table 12: Ship capacities C_{vk}.

Ship \ Product	1	2	3	4
1	983	960	0	0
2	645	0	818	480

Other parameters:

Sailing time between i and j, for the best scenario, is given by the distance between i and j divided by 8. $T_i^B = 0.5, i \in N$. $Q_{vk}^O = 0, k \in K, v \in V$, and $\overline{Q}_{ik} = \min\{max_{v \in V}\{C_{vk}\}, \bar{S}_{ik}\}$. $C_{ijv}^T = 10 * T_{ijv}, C_{ik}^O = 80$ and $C_{ik}^P = 1$. Furthermore, $\mu_i = 3$ for ports with supply storage tanks, and $\mu_i = 1$ for ports which have consumption storage tanks only.

References

[1] D. Adelman. A price-directed approach to stochastic inventory/routing. *Operations Research*, 52:499–514, 2004.

[2] E.-H. Aghezzaf. Robust distribution planning for supplier-managed inventory agreements when demand rates and travel times are stationary. *Journal of the Operational Research Society*, 59:1055–1065, 2008.

[3] A. Agra, M. Christiansen, and A. Delgado. Mixed integer formulations for a short sea fuel oil distribution problem. *Transportation Science*, 47:108–124, 2013.

[4] A. Agra, M. Christiansen, A. Delgado, and L. Simonetti. Hybrid heuristics for a maritime short sea inventory routing problem. *European Journal of Operational Research*, 236:924–935, 2014.

[5] A. Agra, M. Christiansen, R. Figueiredo, L.M. Hvattum, M. Poss, and C. Requejo. The robust vehicle routing problem with time windows. *Computers and Operations Research*, 40:856–866, 2013.

[6] F. M. Al-Athari. Estimation of the mean of truncated exponential distribution. *Journal of Mathematics and Statistics* 4:284–288, 2008.

[7] N. Ando and E. Taniguchi. Travel time reliability in vehicle routing and scheduling with time windows. *Networks and Spatial Economics*, 6:293–311, 2006.

Table 13: Data for each port i and each product k.

Port	Product $k = 1$					Product $k = 2$					Product $k = 3$					Product $k = 4$				
	J_{ik}	R_{ik}	$\underline{S}_{ik}$	$\overline{S}_{ik}$	$\underline{Q}_{ik}$	J_{ik}	R_{ik}	$\underline{S}_{ik}$	$\overline{S}_{ik}$	$\underline{Q}_{ik}$	J_{ik}	R_{ik}	$\underline{S}_{ik}$	$\overline{S}_{ik}$	$\underline{Q}_{ik}$	J_{ik}	R_{ik}	$\underline{S}_{ik}$	$\overline{S}_{ik}$	$\underline{Q}_{ik}$
1	-1	8	80	380	80	0	0	0	0	0	0	0	0	0	0	-1	2	12	50	25
2	1	0	0	0	0	1	0	0	0	0	-1	5	60	520	50	-1	10	120	586	100
3	-1	4	40	170	70	0	0	0	0	0	0	0	0	0	0	-1	1	10	50	20
4	-1	17	200	930	170	-1	40	300	1164	200	1	0	0	0	0	1	0	0	0	0
5	-1	17	170	600	170	-1	8	200	1200	150	-1	15	300	1200	150	-1	2	15	50	20
6	-1	100	600	3800	480	-1	100	500	960	300	-1	60	240	587	240	-1	15	200	777	200
7	-1	15	150	500	100	0	0	0	0	0	0	0	0	0	0	-1	5	30	100	60

[8] H. Andersson, A. Hoff, M. Christiansen, G. Hasle, and A. Løkketangen. Industrial aspects and literature survey: Combined inventory management and routing. *Computers and Operations Research*, 37:1515–1536, 2010.

[9] B. Verweij B, B. Ahmed, A.J. Kleywegt, G. Nemhauser, and A. Shapiro. The sample average approximation method applied to stochastic routing problems: a computational study. *Computational Optimization and Applications*, pages 289–333, 2003.

[10] L. Bertazzi, A. Bosco, F. Guerriero, and D. Laganà. A stochastic inventory routing problem with stock-out. *Transportation Research Part C*, 27:89–107, 2013.

[11] J.R. Birge and F. Louveaux. *Introduction to Stochastic Programming.* Series in Operations Research and Financial Engineering. Springer, New York, 2nd edition, 2011.

[12] 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.

[13] M. Christiansen, K. Fagerholt K., B. Nygreen, and D. Ronen. Ship routing and scheduling in the new millennium. *European Journal of Operational Research*, 228:467–483, 2013.

[14] M. Christiansen and B. Nygreen. Robust inventory ship routing by column generation. In G. Desaulniers, J. Desrosiers and M. Solomon, editors, *Column Generation*, pages 197–224. Springer-Verlag, New York, 2005.

[15] E.E. Halvorsen-Weare and K. Fagerholt. Robust supply vessel planning. In J. Pahl, T. Reiners, and S. Voß, editors, *Network Optimization*, volume 6701 of *Lecture Notes in Computer Science*, pages 559–573. Springer, Berlin, 2011.

[16] E.E. Halvorsen-Weare, K. Fagerholt, and M. Rönnqvist. Vessel routing and scheduling under uncertainty in the liquefied natural gas business. *Computers and Industrial Engineering*, 64:290–301, 2013.

[17] J. Higle. Stochastic programming: Optimization when uncertainty matters. In H.J. Greenberg and J. Cole Smith, editors, *Tutorials in operations research*, Tutorials in operations research, pages 30–53. INFORMS, New Orleans, 2005.

[18] L.M. Hvattum and A. Løkketangen. Using scenario trees and progressive hedging for stochastic inventory routing problems. *Journal of Heuristics*, 15:527–557, 2009.

[19] L.M. Hvattum, A. Løkketangen, and G. Laporte. Scenario tree based heuristics for stochastic inventory routing problems. *INFORMS Journal on Computing*, 21:268–285, 2009.

[20] M. Kaut and S. Wallace. Evaluation of scenario-generation methods for stochastic programming. *Pacific Journal of Optimization*, 3:257–271, 2007.

[21] A.S. Kenyon and D.P. Morton. Stochastic vehicle routing with random travel times. *Transportation Science*, 37:69–82, 2003.

[22] A.J. Kleywegt, V.S. Nori, and M.W.P. Savelsbergh. The stochastic inventory routing problem with direct deliveries. *Transportation Science*, 36:94–118, 2002.

[23] A.J. Kleywegt, V.S. Nori, and M.W.P. Savelsbergh. Dynamic programming approximations for a stochastic inventory routing problem. *Transportation Science*, 38:42–70, 2004.

[24] V. Lambert, G. Laporte, and F. Louveaux. Designing collection routes through bank branches. *Computers and Operations Research*, 20:783–791, 1993.

[25] G. Laporte, F. Louveaux, and H. Mercure. The vehicle routing problem with stochastic travel times. *Transportation Science*, 26:161–170, 1992.

[26] H. Lei, G. Laporte, and B. Guo. A generalized variable neighborhood search heuristic for the capacitated vehicle routing problem with stochastic service times. *TOP*, 29:99–118, 2012.

[27] X. Li, P. Tian, and S.C.H. Leung. Vehicle routing problems with time windows and stochastic travel and service times: Models and algorithms. *International Journal of Production Economics*, 125:137–145, 2010.

[28] J.G. Rakke, M. Stålhane, C.R. Moe, H. Andersson H., M. Christiansen, K. Fagerholt K., and I. Norstad. A rolling horizon heuristic for creating a liquefied natural gas annual delivery program. *Transportation Research, Part C*, 19:896–911, 2011.

[29] R.A. Russell and T.L. Urban. Vehicle routing with soft time windows and erlang travel times. *Journal of the Operational Research Society*, 59:1220–1228, 2008.

[30] H.D. Sherali and S.M. Al-Yakoob. Determining an optimal fleet mix and schedules: Part I - single source and destination. In J.K. Karlof, editor, *Integer Programming Theory and Practice*, pages 137-166, CRC Press, Boca-Raton, 2006.

[31] H.D. Sherali and S.M. Al-Yakoob. Determining an optimal fleet mix and schedules: Part II - multiple sources and destinations, and the option of leasing transshipment depots. In J.K. Karlof, editor, *Integer Programming Theory and Practice*, pages 167-193, CRC Press, Boca-Raton, 2006.

[32] D. Taş, N. Dellaert, T. van Woensel, and T. de Kok. Vehicle routing problem with stochastic travel times including soft time windows and service costs. *Computers and Operations Research*, 40:214–224, 2013.

[33] D. Taş, M. Gendreau, N. Dellaert, T. van Woensel, and T. de Kok. Vehicle routing with soft time windows and stochastic travel times: A column generation and branch-and-price solution approach. *European Journal of Operational Research*, 236(3): 789-799, 2014.

[34] S.Y. Teng, H.L. Ong, and H.C. Huang. An integer L-shaped algorithm for time-constrained traveling salesman problem with stochastic travel and service times. *Asia-Pacific Journal of Operational Research*, 21:241–257, 2004.