European Journal of Operational Research 208 (2011) 86–94

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Production, Manufacturing and Logistics

Maritime inventory routing with multiple products: A case study from the cement industry

Marielle Christiansen [a,*], Kjetil Fagerholt [a,b], Truls Flatberg [c], Øyvind Haugen [a], Oddvar Kloster [c], Erik H. Lund [a]

[a] Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology, Trondheim, Norway
[b] Norwegian Marine Technology Research Institute (MARINTEK), Trondheim, Norway
[c] SINTEF ICT, Applied Mathematics, Oslo, Norway

ARTICLE INFO

Article history:
Received 18 September 2008
Accepted 24 August 2010
Available online 27 August 2010

Keywords:
Routing
Maritime transportation
Inventory management
Heuristics
Genetic algorithms

ABSTRACT

This paper considers a maritime inventory routing problem faced by a major cement producer. A heterogeneous fleet of bulk ships transport multiple non-mixable cement products from producing factories to regional silo stations along the coast of Norway. Inventory constraints are present both at the factories and the silos, and there are upper and lower limits for all inventories. The ship fleet capacity is limited, and in peak periods the demand for cement products at the silos exceeds the fleet capacity. In addition, constraints regarding the capacity of the ships' cargo holds, the depth of the ports and the fact that different cement products cannot be mixed must be taken into consideration. A construction heuristic embedded in a genetic algorithmic framework is developed. The approach adopted is used to solve real instances of the problem within reasonable solution time and with good quality solutions.

1. Introduction

Cement is a cheap but heavy product. This makes transportation costs high compared to other variable costs in the industry. In turn this means that transportation has a large impact on a cement company's bottom line and should therefore be an area to focus on in order to increase profits. Production of cement requires major investments in plant and equipment. Together with low operation margins it makes high volumes necessary. Producers strive to please their customers. Clever planning and good use of resources are therefore extremely important to obtain a profit in this low margin and high volume business.

This paper considers the tactical planning problem of a cement producer that is responsible for the inventory management at the production factories, routing and scheduling of bulk ships and finally, inventory management at the silos where customers consume the cement. There are several grades of cement and each of them has to be treated separately both in the ships' compartments and the inventories at cement factories (production ports) and consumption silos (consumption ports). The customer demand for cement varies considerably during the year and production equipment is often reset to produce new series of cement products, so both the production and consumption have a specified rate that varies in time. There are upper and lower limits for all inventories. To transport the cement between the production and consumption ports, the planners control a heterogeneous fixed fleet of bulk ships. The cargo hold of the ship is separated into several compartments, and each of the compartments has a defined capacity. There is no predetermined visiting pattern regarding the two types of ports: production and consumption. This means that a particular ship can load in several production ports in succession before several consumption ports are called in sequence. However, the geography and number of production ports in the particular case considered ensures that a call to a production port is always followed by a call to one or more consumption ports. Further, both the loaded and unloaded quantity at each call to a port and the number of calls during the planning horizon are variable. The ship fleet capacity is limited, so in peak seasons the silos may run empty. Hence the planning problem is to design routes and schedules for the fleet that ensures efficient transportation while fulfilling the customer demands as well as possible. This also includes determining the cargo sizes of each cement grade at each call and the allocation of cement grades to ship compartments.

Optimization in maritime transportation is a well-established field of research in transportation planning with reviews by Ronen (1983, 1993) and Christiansen et al. (2004, 2007). We have witnessed an accelerating amount of research in the literature during the last decade, indicating increased interest in these types of problems.

* Corresponding author. Tel.: +47 73593602; fax: +47 73593603.
E-mail address: Marielle.Christiansen@iot.ntnu.no (M. Christiansen).

0377-2217/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.ejor.2010.08.023

In maritime transportation, large quantities are usually loaded and unloaded at each call to a port. Both the loading operation and the transportation are time consuming and expensive. Thus the potential is great if the planning of the routing and the inventory management at each end of the sailing leg is synchronized. In reality, we can also find other maritime supply chains where one of the actors has the responsibility for both the transportation and the inventory management. However, we can only find references to this combined problem within the maritime sector in this last decade. Christiansen (1999) considers a problem faced by an ammonia producer. This producer is controlling a fleet of heterogeneous ships transporting ammonia between production and consumption ports, in addition to the inventory management at all ports. Opposed to the cement problem, this is a single product maritime inventory routing problem (IRP) and the production and consumption rates are constant during the planning horizon. Christiansen (1999) solves the problem by a Dantzig–Wolfe decomposition approach. Flatberg et al. (2000) study the same problem, but make use of an iterative improvement heuristic combined with an LP solver for finding arrival times and quantities for given routes. At the next stage in the fertilizer supply chain, the ammonia is further processed into different fertilizer products at the consumption ports, and these products are supplied to the agricultural market. Fox and Herden (1999) describe a mixed integer programming model to schedule ships from such ammonia processing plants to several ports in Australia. The model is solved by commercial optimization software in several steps. Compared to the cement problem, Fox and Herden (1999) do not consider the allocation of products on board the ship and treat routing issues quite differently because of discretization of time. Similar to the cement problem, Ronen (2002) presents a maritime IRP with multiple products. The problem is to decide when and how much of each product to ship, from which production facility to which destination and by which ship. Ronen (2002) assumes that it is possible to separate the shipments planning stage from the ship scheduling stage. This means that much of the flexibility inherent in IRPs is removed. In addition, the ship voyages have a single production and consumption port. This implies that no routing is necessary because every voyage has only one transportation leg and can be treated as a spot charter. Another maritime IRP with multiple products concerns the transportation of calcium carbonate slurry from a processing plant in Norway to European paper manufacturers (Dauzère-Pérès et al., 2007). The problem has, in contrast to our cement distribution problem, a vehicle routing problem structure with one depot (processing plant). Other differences compared to our problem are that the ships never unload in more than one port before returning to the processing plant and that no allocation of cargoes to compartments is considered. The problem is solved by a metaheuristic-based algorithm. Persson and Göthe-Lundgren (2005) have considered a planning problem integrating both the shipment planning of petroleum products from oil refineries to storage depots and the production scheduling at the refineries. In this study too, the allocation of cargoes to compartments is not taken into account. The problem is solved by a branch-and-price-and-cut approach. Another multi-product maritime IRP is studied by Al-Khayyal and Hwang (2007). However, unlike the cement problem, Al-Khayyal and Hwang (2007) assume that the ships have dedicated compartments for each product type. This simplifies the problem since no allocation needs to be performed. Each product can only be loaded in a specific compartment with a given capacity. Hwang (2005) shows that the problem can be decomposed into several subproblems by dualizing coupling constraints and suggests a solution method using Lagrangian relaxation and a randomized greedy heuristic. A combined supply chain and ship routing problem for a large Scandinavian producer of pulp is considered in Bredström et al. (2005) and Bredström and Rönnqvist (forthcoming). Compared to the cement case, a larger part of the supply chain is considered for the pulp industry problem. Also here, the capacity of the fleet of ships is limited compared to the demand in peak periods, so the company transports the pulp products by trains and trucks in addition to spot and owned ships. The problem is solved by use of a hybrid algorithm based on a genetic algorithm and linear programming in Bredström et al. (2005) and a rolling horizon heuristic in Bredström and Rönnqvist (forthcoming). More maritime IRPs are referred to in Christiansen and Fagerholt (2009).

The objective of this paper is to describe a new type of problem in maritime transportation to the OR-literature and to present an appropriate heuristic approach to the planning problem for real-life instances. The structure of the problem can be seen in many real-life planning situations. Compared to most traditional maritime IRPs, this problem includes some complex aspects such as (1) multiple non-mixable products, (2) partial loads allowing several cargoes to be onboard a ship at a time, (3) variable production and consumption rates of the products, and (4) limited ship capacity compared to production and consumption, requiring a decision on which demand to supply in peak periods.

The outline of this paper is as follows: Section 2 describes the real planning problem, while Section 3 is devoted to the solution method. Real-world cases based on planning problems from a major cement producer are described in Section 4 and computational results for these cases are also reported. Finally, concluding remarks and future research follow in Section 5.

2. Case description

The maritime inventory routing problem (maritime IRP) that was introduced in Section 1 will be described in detail in this section. Section 2.1 gives a background and overview of the problem, before elaborating upon the main elements of the problem; ports and inventories, ships and planning objective; in Sections 2.2, 2.3, 2.4, respectively.

2.1. Problem background and overview

The problem studied is a real-life planning problem faced by the major Norwegian producer of cement. It is natural to utilize ships for transport of heavy and high volume bulk commodities like cement, as Norway is a narrow country with a long coastline. The producer operates two factories in Norway, one in the south and one in the north, and controls 28 consumption ports with inventories located along the Norwegian coast. Fig. 1 shows the locations of all the production and consumption ports in Norway.

There are 10 cement grades, also called products, and these must be stored separately from each other. Therefore, each consumption port is equipped with up to five different silos. There are more than 60 silos in Norway, including 12 at the two production ports.

In addition to the production and consumption ports in Norway, shown in Fig. 1, the problem also includes a few *external locations* in neighbouring countries. These external locations can be considered as production ports, even though their inventories are not controlled by the cement producer. One such production port in Sweden acts as a complementary supplier of cement to the Norwegian market. In addition to that, raw materials (only one grade) for the two Norwegian cement factories are transported from external locations in Denmark and the Netherlands. In this sense, the two cement factories can be viewed both as production ports (regarding cement) and consumption ports (regarding raw materials).

Fig. 2 presents a simplified example with external locations, two production ports, four consumption ports and three cement

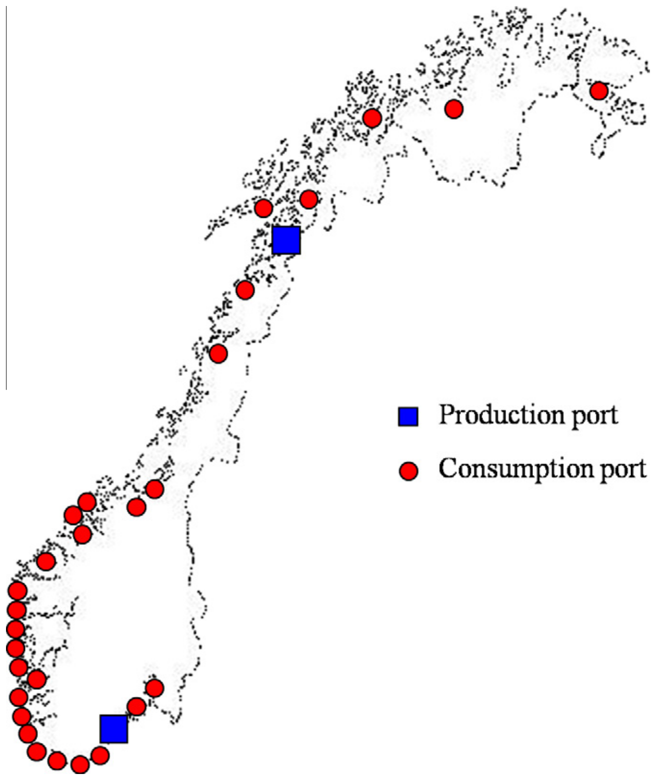

Fig. 1. The production and consumption ports of the cement producer.

grades. The arcs illustrate where the transport of cement and raw materials may be performed by the fleet of ships.

The cement producer controls a heterogeneous fleet of five ships through long-term time-charter contracts. The fleet is used for transporting the cement between silos at the production and consumption ports. The raw materials, as well as the cement from Sweden, are also transported by the fleet.

The maritime IRP is a short-term planning problem with a typical planning horizon of two weeks. Planning follows a rolling horizon principle, where the plan is updated more or less on a daily basis (or when new relevant information appears).

The problem deals with deciding efficient routes and schedules for the fleet of ships, as well as determining loading and unloading quantities at the different ports and silos. An important driving force in the planning is to maintain inventory levels within defined upper and lower limits. Keeping the inventory level below upper limits at the factories is important to protect against overflow, which may result in costly production shutdowns. Moreover, staying above the inventory lower limits at the consumption ports is the key to fulfilling customer demand.

2.2. Ports and inventories

As mentioned, each production and consumption port has one or several silos for storing the cement (and raw material). Each silo has predefined upper and lower inventory limits. These limits may include some safety margins to protect against different uncertain factors, such as variations in demand/production and ship delivery delays. Each silo is dedicated to a single grade of cement. However, two silos at the same port can store the same grade. It is assumed that the cement grade stored in a particular silo is the same throughout the planning horizon.

In the beginning of the planning horizon each silo will have a given initial inventory level. We assume that the production plan for the various cement grades is given as input to the maritime IRP. The same applies to the consumption rates at the different silos, where forecasts are continuously updated by the sales department at the cement producer. This means that there is a production or consumption rate for each product at each port (i.e. for each silo). It should be noted that the production and consumption rates may fluctuate considerably within the planning horizon, and therefore cannot be assumed to be constant.

Due to shortage of ship transport capacity, it is sometimes impossible to find a plan that satisfies the demand at all silos. Sometimes it is possible to organize transport by trucks between the production and consumption ports to reduce or eliminate this shortage. However, the truck transport is expensive and can in practice only be done for small volumes (one truck can only carry about 30 tonnes) and for short distances. Therefore, each consumption port and silo is given a priority. High priority is given to (major) consumption ports that are strategically important, and here, the supply has to be satisfied. Low priority is given to (minor) consumption ports where the consequence of a stock-out is considered smaller, and/or for where the volumes can be delivered by trucks due to short distances.

Only one ship can unload at each consumption port at the same time due to berth capacity constraints.

2.3. Ships

The cement producer controls and operates a fleet of five ships, with two to eight cargo compartments and total loading capacities from 3200 to 5800 tonnes.

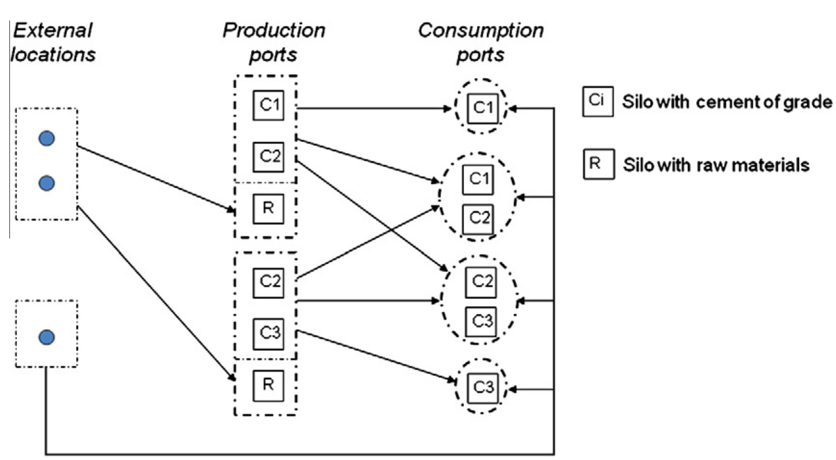

Fig. 2. Product flow.

Each ship has the following attributes that are relevant for the planning process:

– Total load capacity.
– Number of compartments and capacity of each.
– Service speed (defines the sailing times between ports).
– Loading and unloading rates (tonnes/h).
– Physical measures (length, draft and beam).
– Fuel and diesel oil consumption rates (for sailing, (un) loading and when idle).
– Cost of entering any given port (port costs).

It should be noted that in this problem the sailing times are relatively short (from a few hours up to a few days), so the time spent on loading or unloading in ports becomes a substantial part of the total time. The port time depends on the quantity loaded or unloaded. Therefore, we model each ship with a loading and unloading rate for each combination of cement grade and port as it may be dependent on both these factors.

Since different cement grades cannot be mixed, we have to keep track of the cargo levels in each compartment. Even though the final stowage decisions onboard the ship are made at loading time by the ship's master, stowage must be considered in the maritime IRP as it influences the feasibility of the plans. We assume in this paper that vessel stability on partially loaded legs is not an issue in this planning. However, this must also be dealt with by the ship's master before voyage execution.

The ship's physical measures, especially the draft, often result in ship–port compatibility constraints that must be satisfied in the planning. For instance, the largest ships cannot enter particular ports due to draft restrictions in the port. Sometimes these constraints can be dependent on the load onboard the ship, such that a ship can only enter or leave a given port if the total load onboard is less than e.g. 50% of its capacity.

At the beginning of each planning period, each ship will have a given initial position and possibly load onboard with known destinations. The initial position can be a port or a point at sea.

2.4. Objective function

There are multiple objectives in the planning of cement distribution:

(1) fulfil customer demands (i.e. have little or no shortages),
(2) minimize transportation costs, and
(3) achieve appropriate inventory end levels at the consumption ports.

The last objective is important in order to reduce end effects since the planning follows a rolling horizon principle.

In peak periods there is a shortage in ship transport capacity relative to demand, so objective (1) is modelled as a cost for not fulfilling the demand or violating lower inventory limits in the consumption ports.

Since the planning follows a rolling horizon principle, it is not given that minimizing costs for the next planning period will minimize costs in the long run. Cost minimization for a given planning period will result in a plan that seeks to transport as little as possible. This will negatively influence the number of shortages in following planning periods. Finding the optimal balance between cost minimization and the quantity transported will greatly depend on the planning situation and the planner's preferences. We will present some computational experiments regarding this issue in Section 4.

3. A heuristic solution approach

To solve the maritime IRP, we must design a plan for how to use the ships to service the silos. The plan describes the route of each ship as a sequence of *port calls*, each indicating a port and a period of time that the ship is in the port. In addition, a port call may contain one or more *actions*. An action consists of either loading or unloading a certain amount of cement from or to a silo.

The following terminology is used in describing the algorithm. A shortage or an overflow in a silo is called a *capacity breach*. A *shipment* is a pair comprising a single load action and a corresponding single unload action for the same ship, transporting an amount of cement from one silo to another. The positions of the load action and the corresponding unload action in the route are specified for the shipment. The unload action of a shipment does not need to follow immediately after the load action, as actions for other shipments may be performed in-between.

In the following, we first give an overview of the solution approach, before discussing the components in more detail.

3.1. Overview of heuristic solution approach

The solution method for the studied maritime IRP has two components: a heuristic construction algorithm and a genetic algorithm. The construction algorithm builds a plan from scratch. It is deterministic, but has parameters that can be varied to produce different plans. The genetic algorithm is used to search for parameters that produce good plans by the construction heuristic. It evolves a population of individuals that represent different parameter settings. The fitness of an individual is evaluated by running the construction heuristic using the individual's parameter settings and calculating the objective value of the produced plan.

The construction algorithm starts with an initial plan. In this plan, each ship starts either empty, with the initial time and position given by a port call, or it starts loaded and the plan contains a sequence of port calls and unload actions that lead to an empty ship. As little, if any, cement is transported in the initial plan, this plan normally implies a capacity breach within the planning horizon for most silos. We then start the process of iteratively improving the plan. In each iteration, the silo with the most critical capacity breach is identified. Then a shipment between the critical silo and some other silo is added to the plan, thus postponing or eliminating the critical breach. If all goes well, the procedure eventually leads to a plan with no capacity breaches. However, the final plan may also contain capacity breaches, if at some point during construction it is impossible to find a feasible way to add a shipment that improves the critical breach. An overview of the algorithm is illustrated in Fig. 3.

In each iteration, we can typically choose between many different shipments that can service the critical silo (using different ships, different silos and various placements of the actions in the

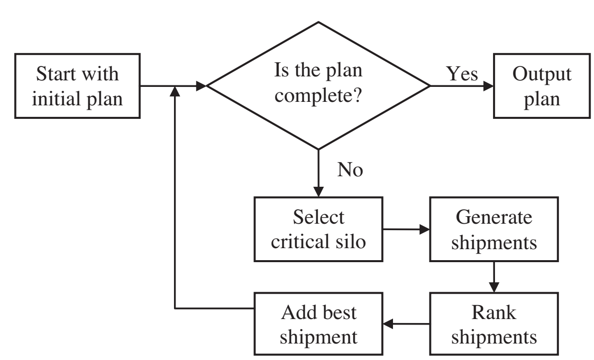

Fig. 3. Overview of the construction heuristic.

ship's plan). The algorithm evaluates each shipment according to a number of criteria and chooses the one that maximizes a weighted sum of the criteria scores. The weights used in this sum are the algorithm parameters that the genetic algorithm seeks to optimize.

3.2. The construction heuristic

Here, the construction heuristic is described in more detail. As mentioned, the algorithm starts with an initial plan. The plan specifies the ships at known initial positions and times, but apart from possibly unloading their initial load, the ships have no load or unload actions. The state of each silo at a given initial time and their subsequent production or consumption profiles are also known. The algorithm then repeats the following steps until the plan is complete:

1. Select the critical silo
2. Select a counterpart silo to get one producer and one consumer
3. For each ship
4. For each possible pair of insertion points in the ship's plan
5. Insert a new shipment for the ship, serving the silos
6. Assign time and quantity to make the plan feasible
7. If successful, record the shipment
8. Remove the shipment and restore the plan
9. If no feasible shipments were found, exclude the silo
10. Otherwise, select the best recorded shipment and add it to the plan

The steps are described in more detail below.

1. Critical silo selection

For each silo (except those that have been excluded in step 9), the algorithm finds the earliest time at which there is a capacity breach, i.e. when the inventory level breaks the allowed upper or lower limit, given the plan constructed so far. If the inventory levels at the end of the planning horizon are not within the desired end level limits, this also counts as a breach. The silos that have a breach in the planning horizon are sorted, firstly on silo priority, and secondly on the time of the breach. The highest priority silo with the earliest breach time is chosen as the critical silo.

The algorithm terminates at this point if only the excluded silos have a capacity breach within the planning horizon.

2. Counterpart silo selection

Usually, the critical silo is a consumption silo, and this is assumed in the following. The case where the critical silo is a production silo is symmetric.

Having selected a critical consumption silo, we need to select a production silo from which the same product can be loaded. These silos are prioritized according to two criteria: the distance between the production and consumption silos, and the time at which there is a breach at the production silo. The silo with the minimal sum of ranks in these two lists is chosen.

Instead of selecting a single critical silo, we may choose to generate shipments for all silos that have a capacity breach. This behaviour is controlled by a parameter.

3. Ship selection

Steps 4–8 are done for all ships except those that are not allowed to enter the production or consumption port, e.g. due to draft limits.

4 and 5. Insertion point generation

The loading action may be inserted after any existing action for the ship, including at the end of the ship's plan. The unloading action may then be inserted after the loading action or any later action. For each combination of insertion points, we generate a shipment that is evaluated in steps 6–8.

An insertion pointis immediately rejected if one of the new actions would occur so late in the plan that it cannot be performed before the silo's capacity has been breached.

6. Time and quantity assignments

Having inserted the two new actions (with as yet unknown start time and quantity), the algorithm proceeds through a number of steps to find legal time and quantity assignments for the shipment. In this process, the algorithm may change the timing of other actions for the chosen ship, and thereby also the inventory levels at the silos that it services, but the plans for the other ships are left unchanged. The procedure is complex, partly because the values chosen can influence large parts of the plan, and partly because the time and quantity cannot be chosen independently, since the duration of an action depends on the quantity loaded or unloaded. The procedure is given below, in an imperative mode, since the form of the text is close to pseudocode.

The algorithm has a parameter that gives the smallest quantity of cement allowed in a shipment. Assuming this quantity for the new shipment, time is propagated through the ship's plan. That is for each port call and action; the new start and end times are calculated based on the end time of the previous port call/action, the required sailing time and the action's duration. This establishes approximate new times for the ship's actions. Note that existing actions may now be significantly postponed due to the newly inserted shipment. If the end time of the plan exceeds the planning horizon, reject this shipment.

Examine the plan for the production silo from the new action onwards, and find the maximum quantity that the ship can load in the new action without causing a shortage to occur at any time in the silo. Similarly, limit the quantity so that there will be no overflow at the consumption silo.

Limit the maximum quantity to ensure that the ship obeys any draft restrictions in the ports of the production and consumption silos, and those visited in-between.

Limit the maximum quantity to ensure that the ship's load does not exceed its capacity at any time.

Assign the compartments in which to load the shipment, as follows. The available compartments are those that are empty or carry only the same cement grade between the load and discharge of the shipment. Examine all subsets of these compartments and select the subset whose total free size is closest to but smaller than the maximum quantity. Set the loaded quantity for the shipment equal to the total free size of the chosen compartments.

Using the now assigned quantity for the shipment, propagate the time and load through the ship's plan. If the end time of the plan exceeds the planning horizon, reject this shipment. Also propagate the new inventory levels for the silos that the ship visits.

Check if any of the ship's actions now overlap in time with some other ship's action at the same consumption port. If so, postpone the action so that the overlap is eliminated. Then restart the time and quantity assignment procedure for this shipment, but with a twist: from now on, time propagation may never assign an earlier start time to this action than the postponed start time just arrived at, in any plan that contains the shipment currently being evaluated. Thus this particular overlap of actions will not occur again.

Check if the changes made in the plan have caused problems at some silos. If any silo serviced by this ship has a capacity breach (except after the last action added so far), reject this shipment.

7 and 8. Record shipment and restore plan

At this point in the algorithm, there are legal time and quantity assignments for the shipment. We record the shipment, remove it to restore the plan as it was and proceed to evaluate the next shipment for the next pair of insertion points.

9. Excluding the silo

If no feasible shipment is found that improves the critical breach, the algorithm excludes the critical silo. This means that for the remainder of the construction, breaches for this silo are ignored. Consequently, the resulting plan will end up containing a breach for this silo, unless the breach is accidentally fixed by a shipment created for a different critical silo.

10. Shipment selection

If the algorithm finds more than one feasible shipment, just one can be added to the plan. We would like the chosen shipment to be efficient, in order to reduce the cost locally and enable the addition of further shipments later. As it is unlikely that a single efficiency measure works best for all problems, we use several, and tune them using adjustable weights. Note that the algorithm does not simply use the objective function as the efficiency measure. While the objective function is authoritative for the quality of a complete plan, it is not given that it is good at estimating which partial plans will lead to good complete plans.

For each shipment, the following statistics are evaluated:

– q: the quantity transported in the shipment.
– Q: the capacity of the chosen ship.
– t: the increase in sailing time caused by the shipment.
– c: the increase in plan cost caused by the shipment.

Then a set of lists is created, sorting all shipments according to the following criteria:

(1) Quantity: sorted from large to small q.
(2) Extra sailing time: sorted from small to large t.
(3) Capacity utilization: sorted from large to small q/Q.
(4) Quantity per time: sorted from large to small q/t.
(5) Cost per quantity: sorted from small to large c/q.
(6) Pseudorandom: shuffled using a random generator.

Each list has an associated weight. The score of a shipment is found by summing its rank in each list, multiplied by the associated weight. The shipment whose score is minimal is added to the plan.

3.3. The genetic algorithm

The solution produced by the construction heuristic depends on the weight chosen for each list in the shipment selection step. A random generator is used to produce the pseudorandom lists, and the seed chosen for this generator can also influence the solution. These parameters form a search space, within which we should search for values that produce a good solution. As a simple and robust means for performing this search, we have chosen to use a genetic algorithm. The genetic algorithm has some adjustable parameters, e.g. population size, which have been chosen by limited experimentation to get robust behaviour. Fine-tuning these parameters has not been a focus in the present work. Genetic algorithms have been used for solving maritime IRPs before; see e.g. Bredström et al. (2005).

The genetic algorithm works with a population of *individuals*. An individual has a *genome*, which consists of a collection of weights and a seed for the pseudorandom ranking. The genome is represented directly as a collection of numbers, and not in binary chromosome form. The individual's *phenotype* is the plan generated by the construction algorithm using the genome's settings. The individual's *fitness* is equal to the objective value of the phenotype plan.

The algorithm starts with a population of 20 randomly generated individuals. The following steps are then repeated:

– Generate a new population of 40 individuals. Each new individual's genome is created by recombining the genomes of two randomly selected individuals in the old population, using the method explained below.
– Preserve the four best individuals from the old population by adding them to the new population. Preserving the best individuals is known as *elitism*.
- Evaluate the fitness of all individuals in the new population and reduce it by keeping only the 20 best individuals. This population is the starting point for the next iteration.

Recombination works as follows: each weight in the genome of a new individual is determined from the corresponding weights, w_1 and w_2, in the two 'parent' individuals. Assuming $w_1 \leqslant w_2$, let $L = w_2 - w_1$. The new weight is drawn randomly from a uniform distribution over $[\max(0, w_1 - 0.55L), \min(1, w_2 + 0.55L)]$. The value 0.55 is chosen so that by itself, the recombination procedure yields a slight increase in the variance of the weights over time, but this is counteracted by the elimination of individuals in inferior regions of the solution space. There is also a 20% chance that a randomly selected weight in the new genome is mutated to a random value between 0 and 1. The random seed for the new genome is drawn randomly.

The genetic algorithm can be run for as long as desired in the hope of finding better solutions.

4. Computational study

In this section we present results from running the algorithm on a real-world test case. The main goal of this section is to establish how well the plans produced by the construction heuristic satisfy the overall objectives of Section 2.4. The critical aspect is whether the ranking criteria used in the shipment selection locally, lead to good plans for the complete period. To assess this we have done experiments with each ranking criterion individually and experimented with the genetic algorithm as a method for combining the various criteria.

4.1. Test cases

Based on input from planners at the cement company, we have made a test case based on a real planning situation. The case is

Table 1
The characteristics of the real-world case.

	#	Capacity (tonnes)	Initial inventory (tonnes)
Ships	5	23,300	
Production silos	12	164,960	81,837
Consumption silos	49	144,395	71,892
Products (incl. raw material)	11		
Planning period	14 days		

Table 2
Results for the 14-day case with different ranking criteria.

Criterion	Quantity (tonnes)	Cost	Cost per q	Shortage (tonnes)	Below end level (tonnes)
1. Quantity (q)	58,094	282,520	4.86	3378 (3;8)	3393 (4;18)
2. Extra sailing time (t)	68,564	276,079	4.03	2032 (3;6)	3513 (6;14)
3. Capacity util. (q/Q)	63,548	292,785	4.61	2040 (3;6)	3482 (3;19)
4. Quantity per time (q/t)	68,348	269,399	3.94	2011 (3;6)	2188 (5;13)
5. Cost per quantity (c/q)	71,588	277,190	3.87	1802 (1;5)	2748 (3;11)
6. Random	60,670	315,715	5.20	3365 (2;9)	4084 (3;22)

based on a period with hot markets where the company does not have the necessary ship capacity to satisfy all demand. The required end levels were set to 40% of the silo capacity. Table 1 summarizes the characteristics of the case.

Based on the 14-day case we have created test cases for 21 and 28 days by duplicating the production and consumption profiles from the first half and the whole period, respectively. These test cases will not be as realistic, but will provide information about the algorithm's behaviour as the length of the time period increases.

4.2. Computational results

The algorithms described in Section 3, have been implemented as part of a generic C++ software library for solving maritime IRPs. The library, called *Invent*, is based on a generic conceptual model for maritime routing with a corresponding XML format for data input and plan output. All test results referred to in this section were obtained on a 2.6 GHz Intel Core 2 Duo PC with 4 GB of RAM.

Table 3
Results using a genetic algorithm on the 14-day case with different objective functions.

Objective	Quantity (tonnes)	Cost	Cost per q	Shortage (tonnes)	Below end (tonnes)
Quantity	76,541	297,313	3.88	1772 (2;3)	3511 (4;10)
Cost	66,489	239,613	3.60	1119 (1;7)	2328 (2;15)
Cost per quantity	72,806	255,512	3.51	2026 (1;6)	3145 (3;14)
Shortage	67,048	278,734	4.16	981 (1;6)	3097 (2;16)

In the construction algorithm there are five different criteria for ranking shipments, as described in Section 3.2. To compare their individual efficiency, we ran the construction heuristic using only one insertion criterion at a time, including the random ranking. The result for each criterion is reported in Table 2. For each criterion we report the total quantity transported, the cost of the plan, the cost per quantity transported, the total shortage with regard to minimum limits and the total breach of the required end levels. All cost values appearing in this paper are normalized. For the breaches we report the number of silos affected in parenthesis – the first number being high priority silos and the second, lower priority silos. For the 14-day case there were no overflows at the production silos.

The results show that cost per quantity is the best criterion if used on its own, followed by quantity per time and extra sailing time. Still, a weighted combination of criteria may produce better results. We also note the poor performance when using the quantity transported or the capacity utilization to rank shipments. This can be explained by the fact that these criteria do not include time as a limiting resource, and will prefer long and costly shipments if they transport large quantities.

To see if a weighted combination of ranking criteria can lead to improved results, we used the genetic algorithm on the 14-day case. For each test run we ran the algorithm until 500 solutions had been generated. As discussed in Section 2.4 the objective to use is not clearly defined. Table 3 shows the results obtained using different objectives for the genetic algorithm.

These results clearly outperform those obtained using a single ranking criterion with better values obtained for all quality measures, except the total breach quantity of the required end levels. If the heuristic is unable to generate a journey for a critical silo,

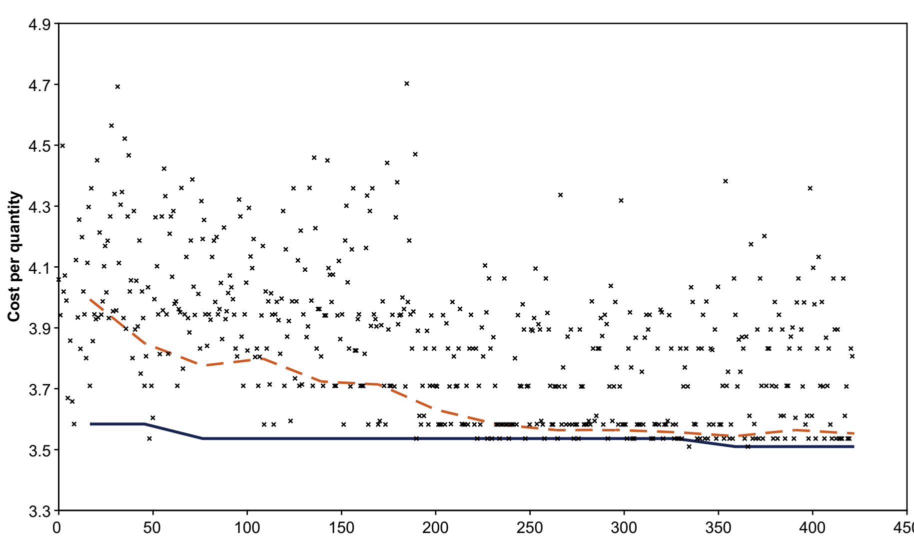

Fig. 4. The development of the best objective value (solid line) and the average objective value (dashed line) of the population as a function of time. In addition, the graph shows the objective value of each generated solution throughout the search.

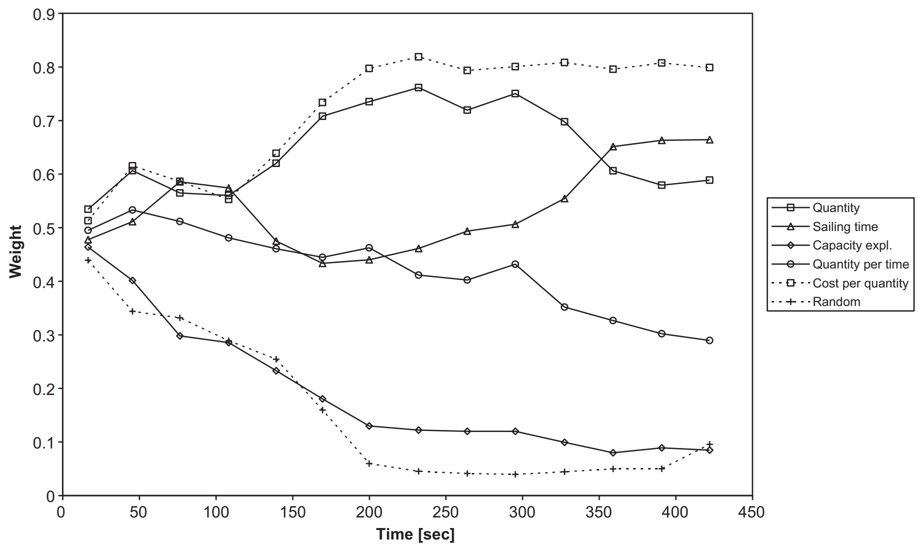

Fig. 5. The average weight of each ranking criterion in the population as a function of time.

the silo is excluded for the rest of the search and will end up below its lower limit. These silos will have to be served by trucks at additional cost or left unserviced with a reduction in consumer satisfaction. As we do not have cost data for these situations, it is not clear which of the plans in Table 3 is best for operational purposes. Both focusing on quantity and reducing the shortage leads to higher cost, while a focus on cost leads to an increased shortage.

The results of using the genetic algorithm with cost per quantity transported as the objective value are summarized in Figs. 4 and 5. Similar results were obtained for the other objectives.

The results in Fig. 5 are consistent with the results in Table 2, showing a preference for the cost per quantity as the search progresses. As the search progresses there will be less variation in the population and the same solution may be generated multiple times. To see whether the genetic algorithm is better than a random exploration of the space of ranking weights, we performed a set of 500 test runs with random weights. The results are summarized in Table 4.

As can be seen from Table 4, the genetic algorithm produces better results for one of the objectives and worse for another, while similar results are obtained for the remaining two objectives. The difference between the two methods is marginal and does not favour either of them.

Using the extended versions of the basic test case, we tested the effect of increasing the problem size. The results are summarized in Table 5. The reported time is the time to generate a complete solution, averaged over 100 runs with different ranking weights. The running time increases considerably, because the number of algorithm iterations and the work of each propagation both grow roughly in proportion to the length of the planning horizon, while the number of possible insertion points in each iteration grows quadratically. It is evident that for larger cases, where planning speed becomes an issue, more effort must be spent to eliminate inferior insertion points quickly.

Table 4
Comparison of the genetic algorithm vs. multiple runs for the different objectives.

Objective	Genetic algorithm	Multiple runs
Cost	239,613	239,613
Quantity	76,258	76,541
Cost per quantity	3.51	3.54
Shortage	981 (1.6)	981 (1.6)

Table 5
Running times for different planning horizons.

Period	Time (s)
14 days	0.71
21 days	7.81
28 days	31.38

5. Concluding remarks

This paper has presented a new maritime inventory routing problem (IRP) with multiple products allocated into different silos at ports and different compartments on ships. The problem is based on a real-world problem from the cement industry. The demand for cement products varies considerably throughout the year, so in peak periods the demand exceeds the ship fleet capacity. The planning problem consists of designing routes and schedules for the fleet that ensure efficient transportation while fulfilling the customer demands as well as possible. This also includes determining the cargo sizes of each cement product at each port call, the allocation of cement products to ship compartments and which silos to let go short in peak periods.

Even though this is a case study from the cement industry, we can find many similar IRPs in other businesses such as the oil and gas industry. Within maritime transportation, there is a particular potential in combining the routing and the inventory management at each end of the sailing leg. This comes from the fact that large quantities are loaded and unloaded at each call to a port, and both the loading operation and the transportation are time consuming and expensive.

A construction heuristic embedded in a genetic algorithm framework is developed to solve the real problem. In real planning of this type of problem it is important to get solutions quickly, and the heuristic finds a solution for a relevant case within a second. The heuristic is flexible in the way that it is easy to change the objective function, change the weights of the different quality measures and produce several different good solutions. For the real planning situation in the particular company, these characteristics of the solution method are very important due to the diffuse objective and a high probability of shortage at some consumption silos. The solutions produced by the heuristic have been presented to the company involved, and they regard the quality of the solutions as better than the ones they achieved by manual planning methods.

This is a crucial planning problem for the shipping industry and an interesting problem for researchers, so solution methods for this type of problem should be focused on even more in the future. By introducing an innovative adaptive construction with learning mechanisms, the heuristic can be improved. An improvement heuristic on top of the constructive heuristic has to be studied, but it is challenging to develop efficient neighbourhood operators because a move can result in considerable changes in the rest of the solution. It is obvious that the real planning problem contains several uncertain parameters like the demand and sailing times. How to consider these uncertain parameters is a topic of future research. Finally, it is important to produce alternative solutions for the particular company, but also as part of a solution approach. Therefore, it is relevant to develop good distance measures and diversification mechanisms for these types of problems.

Acknowledgements

This work was partly supported by the DESIMAL, OPTIMAR and DOMinant projects funded by the Research Council of Norway. We would also like to thank the management at Norcem. Here, we are grateful to Gunnar Syvertsen (Chief Operating Officer) and Birger Pedersen (Logistic Manager) for their involvement in the project.

References

Al-Khayyal, F., Hwang, S.-J., 2007. Inventory constrained maritime routing and scheduling for multi-commodity liquid bulk. Part I: Applications and model. European Journal of Operational Research 176, 106–130.

Bredström, D., Rönnqvist, M., forthcoming. Supply chain optimization in pulp distribution using a rolling horizon solution approach. International Journal of Operational Research.

Bredström, D., Carlsson, D., Rönnqvist, M., 2005. A Hybrid Algorithm for Distribution Problems. IEEE (July/August, pp. 19–25).

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

Christiansen, M., Fagerholt, K., 2009. Maritime inventory routing problems. In: Floudas, C.A., Pardalos, P.M. (Eds.), Encyclopedia of Optimization, second ed. Springer-Verlag, pp. 1947–1955.

Christiansen, M., Fagerholt, K., Nygreen, B., Ronen, D., 2007. Maritime transportation. In: Barnhart, C., Laporte, G. (Eds.), Handbooks in Operations Research and Management Science, Transportation, vol. 14. North-Holland, Amsterdam, pp. 189–284.

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

Dauzère-Pérès, S., Nordli, A., Olstad, A., Haugen, K., Koester, U., Myrstad, P.O., Teistklub, G., Reistad, A., 2007. Omya hustadmarmor optimizes its supply chain for delivering calcium carbonate slurry to european paper manufacturers. Interfaces 37, 39–51.

Flatberg, T., Haavardtun, H., Kloster, O., Løkketangen, A., 2000. Combining exact and heuristic methods for solving a vessel routing problem with inventory constraints and time windows. Ricerca Operativa 29 (91), 55–68.

Fox, M., Herden, D., 1999. Ship scheduling of fertilizer products. OR Insight 12 (2), 21–28.

Hwang, S.-J., 2005. Inventory Constrained Maritime Routing and Scheduling for Multi-Commodity Liquid Bulk. Ph.D. Thesis, Georgia Institute of Technology, Atlanta.

Persson, J.A., Göthe-Lundgren, M., 2005. Shipment planning at oil refineries using column generation and valid inequalities. European Journal of Operational Research 163, 631–652.

Ronen, D., 1983. Cargo ships routing and scheduling: survey of models and problems. European Journal of Operational Research 12, 119–126.

Ronen, D., 1993. Ship scheduling: the last decade. European Journal of Operational Research 71, 325–333.

Ronen, D., 2002. Marine inventory routing: shipments planning. Journal of the Operational Research Society 53, 108–114.