Computers in Industry 64 (2013) 128–135

Contents lists available at SciVerse ScienceDirect

Computers in Industry

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

Bi-objective optimization of the intermodal terminal location problem as a policy-support tool

Kenneth Sörensen, Christine Vanovermeire *

University of Antwerp Operations Research Group ANT/OR, Faculty of Applied Economics, Belgium

ARTICLE INFO

Article history:
Received 24 January 2012
Received in revised form 4 June 2012
Accepted 18 October 2012
Available online 16 January 2013

Keywords:
Intermodal terminal location
GRASP
Metaheuristics
Bi-objective optimization

ABSTRACT

Determining the optimal layout of an intermodal terminal network, more specifically the optimal locations of the terminals, is a complicated matter that requires adequate decision support tools. In this paper, a bi-objective model is developed, minimizing both the transportation cost for the users of the terminal network, as well as the location cost for the terminal operators. A problem-specific GRASP (greedy randomized adaptive search procedure) is developed to solve the bi-objective terminal location problem efficiently. The algorithm only has a single parameter, that determines the allowed calculation time and can be used to improve the quality of the Pareto set approximation.

1. Introduction

Multimodal transportation is defined as the movement of goods by at least two different modes of transport [1]. Intermodal transportation is a particular type of multimodal transportation, where goods, although transported by different modes of transport, use the same transport unit throughout their transportation. Switching to a different mode can only be done at a specific point, called a *terminal*. The best known example of an intermodal transport unit is the standardized (ISO-)container. Once packed in containers, goods can be carried across the globe using trucks, trains, barges, and ocean liners without every requiring repackaging. Terminals exist between any pair of transport modes, allowing for a fast distribution of goods around the globe. Without a doubt, the standardization of intermodal shipping containers has caused an enormous increase in worldwide shipping.

Although intermodal transportation is one of the most promising solutions to reduce greenhouse gasses in the European Union [2], Fig. 1 shows that road transportation is still preferred over more sustainable transport modes such as rail, pipelines or barge transportation. From 2000 to 2010, the share of goods transported with the most polluting mode (road transportation) has even increased. Furthermore, projections for 2020 do not predict a substantial increase in the use of the multimodal network (Table 1).

According to Pedersen et al. [3], the most important factor contributing to an increase of the attractiveness of the intermodal alternative to road transport, is the location of the intermodal terminals. Bowersox et al. [4] list speed, frequency and availability (or connectivity, the ability to service any pair of locations) as the main factors influencing the modal choice. More recent research [5] also lists these criteria, as well as flexibility and reliability. The correct number and location of terminals can thus significantly increase the usage of the intermodal network.

The decision where to locate terminals is dependent on many criteria, many of which are difficult to estimate, or even to express them at all in monetary values [6]. Therefore, a need exists for decision support tools, that can help a decision-maker or analyst determine the optimal lay-out of the intermodal network given different scenarios and costs. Such decision-makers could be large companies examining the optimal terminals (e.g., ports) to use in order to transport their goods, institutions such as the European Commission setting up a road map to support investments in the intermodal infrastructure, and terminal operators determining where to invest in new terminal capacity.

An important point in this respect is that not all costs in the intermodal terminal location problem are paid by the same stakeholders. In most cases, the fixed costs associated with building and exploiting terminals are carried by the terminal operators, whereas the transportation costs through the intermodal terminal network are paid for by the users of the networks, i.e., the companies transporting their goods, and, to some extent, by governments and other authorities supporting intermodal terminals by building port infrastructure. Most optimization models for locating intermodal terminals calculate a single "total

* Corresponding author at: University of Antwerp, Prinsstraat 13, 2000 Antwerp, Belgium. Tel.: +32 3 265 40 40; fax: +32 3 265 49 01.
E-mail address: christine.vanovermeire@ua.ac.be (C. Vanovermeire).

0166-3615/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.compind.2012.10.012

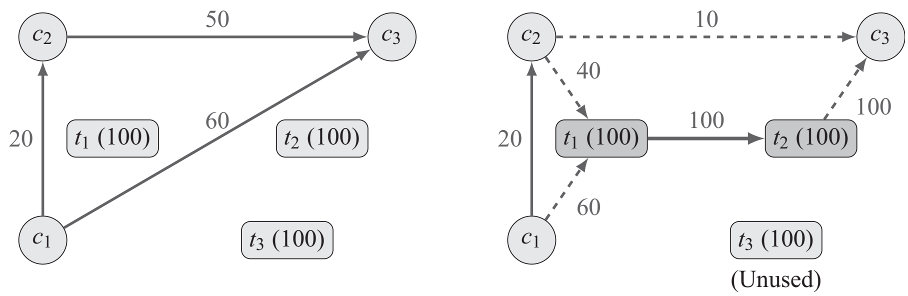

Fig. 1. Example of the intermodal terminal location problem.

cost", which is the sum of location and transportation cost. An argument can be made in favor of not combining these two costs in a single measure of the total cost of the network, but rather keep them separated and solve the intermodal terminal location problem as a bi-objective problem.

In this paper, the bi-objective terminal location problem is modeled and an algorithm is developed that generates a Pareto set, i.e., a set of mutually non-dominated solutions. In Section 2, an overview is given of the literature on models and optimization techniques for terminal location problems. The bi-objective model presented in Section 3 is an adaptation of an existing (single-objective) model, originally developed by Arnold et al. [7]. Given the strategic importance and the complexity of the terminal location problem, the aim of this model (and the problem-specific GRASP metaheuristic that is developed in Section 4 to solve it) is to give the decision-maker a clear insight in the trade-offs between location and transportation costs that exist. The resulting set of non-dominated solutions (see Section 5) may be used as a starting point for a more in-depth analysis of the characteristics of several different terminal network layouts.

2. The intermodal terminal location problem

The intermodal terminal location problem was introduced in Arnold et al. [7], in which the authors propose a mixed-integer programming (MIP) model that minimizes the total cost, i.e., the sum of uni- and intermodal *transportation* costs and fixed terminal *location* costs. In Sörensen et al. [8], we prove that the model of Arnold et al. [7] is NP-hard and develop the first metaheuristic to quickly obtain high-quality solutions.

A graphical representation for a simple problem with 3 customers, 3 origin–destination flows, and 3 terminals is shown in Fig. 1. All terminals are assumed to have a capacity of 100 units (indicated between brackets). The left part of the figure shows the (unimodal) transport streams if no terminals are open. The right part visualizes how these streams are rerouted when terminals t_1 and t_2 are open and t_3 is closed. The flow from c_1 to c_2 is still shipped by unimodal transport, whereas the flow between c_1 and c_3 is shipped entirely by intermodal transport. The flow between c_2 and c_3 is shipped only partially by intermodal transport, due to the limited capacity of the terminals.

One of the drawbacks of the model of Arnold et al. [7] is that it considers the problem from the viewpoint of a stakeholder that pays all costs, combining both transportation and terminal location costs in a single cost function. Clearly, such a stakeholder does not exist. Sirikijpanichkul and Ferreira [9] argue that the decision to locate a terminal has different perspectives. They distinguish three different stakeholders, namely the users, the service providers and the community as a whole, each having their own objectives. The authors state that the best way to improve the decision-making process is to incorporate multiple objectives, corresponding to different stakeholders.

A multi-objective approach to the intermodal location problem has not yet been developed in the literature. However, the use of multi-criteria evaluation methods such as EMOLITE [10], the evaluation procedure of Groothedde [11], PROMETHEE [12], and MAMCA [13], clearly demonstrate the usefulness of multi-objective optimization for terminal location problems.

The closely related facility location problem, on the other hand, has many multi-objective variants. Different objectives used in facility location problems can be found in Farahani et al. [14] (see Table 2). Many combinations of two or more of the objectives in this table have been studied in the literature.

The intermodal terminal location problem is closely related to problems such as the capacitated (fixed charge) facility location problem, the multicommodity capacitated (fixed charge) network design problem and the family of the hub location problems. However, we argue in Sörensen et al. [8] that the differences between all these problems and the intermodal terminal location problem are sufficiently large to warrant the development of new algorithms. Arnold et al. [7] solve the (single objective) intermodal terminal location problem with a branch-and-bound procedure, but such exact methods are not suitable to solve large real-life instances.

In Sörensen et al. [8], we therefore develop two different metaheuristics. Both methods work in a similar manner: they first construct a solution using some construction heuristic, and improve this solution in a second phase using local search. The main difference between the two methods lies in the heuristic used in the construction phase, which is a GRASP (greedy randomized adaptive search procedure) for the first metaheuristic and an ABHC

Table 1
Evolution of the modal split (in %) in Europe.

Mode	2000	2010	2020
Road	43	46	45
Rail	11	9	8
Inland waterways	4	3	3
Pipelines	3	3	3
Sea	39	39	41

Source: Ref. [19].

Table 2
Objectives influencing terminal/facility location decisions.

Objective	Example
Cost	
- Fixed	Installation and start-up cost
- Variable	Transportation cost
Environmental risks	Waste disposal
Coverage	Population coverage
Service level and effectiveness	Delivery frequency
Profit	Outcomes of the investment
Other	Resource accessibility

(attribute-based hillclimber) for the second. Both construction heuristics share a set of common components. They both start from a solution in which all terminals are closed. Terminals are then added (opened) one by one. The sequence in which the terminals are opened varies between the two different heuristics, although they are based on the same logic: *terminals that have the best location cost per unit of capacity should be opened first, as they are more likely to be part of the optimal solution*. However, to create diversity and search the entire solution space, the GRASP-based algorithm introduces randomness that allows to choose a different terminal than the best one according to this criterion. The attribute-based hillclimber on the other hand creates diversity by storing solutions in an archive that are the best solution yet encountered containing a specific attribute, where an attribute is defined as the status of a terminal (opened or closed). In this way, both algorithms find solutions that are of high-quality and diverse. The best performing solutions are then stored in an archive and further improved by local search. For a more in-depth discussion, we refer to Sörensen et al. [8].

A multi-objective algorithm to solve this problem however has not yet been developed. As genetic algorithms generally work well for the related class of multi-objective facility location problems, these are also proposed as a solution technique for the multi-objective intermodal terminal problem [14]. However, also well-performing heuristics incorporating tabu-search, particle swarm optimization, ant-colony optimization and greedy procedures have been reported [14].

3. A bi-objective model of the terminal location problem: operators vs. users

Given is the cost of opening a terminal at a certain location k, which is denoted as the location cost (or fixed cost) F_k. The capacity of terminal k is denoted by C_k. The quantity of goods q_{ij} that need to be shipped from a customer (location, or zone of activity) i to a customer j is assumed to be known, as well as the cost of shipping an item of goods to its destiny unimodally (c_{ij}), or multimodally (c_{ij}^{km} is the multimodal cost of shipping from customer i to customer j using terminals k and m).

The values of the decision variables y_k determine whether a terminal k is opened ($y_k = 1$) or closed ($y_k = 0$). The routing of the goods through the network is determined by the decision variables x_{ij}^{km}, denoting the amount of goods that are shipped from customer i to customer j through terminals k and m, and w_{ij}, the amount of goods that is shipped unimodally from i to j.

$$\min \sum_{\substack{i,j \in I \\ i \neq j}} \sum_{\substack{k,m \in K \\ k \neq m}} c_{ij}^{km} x_{ij}^{km} + \sum_{\substack{i,j \in I \\ i \neq j}} c_{ij} w_{ij} \tag{1}$$

$$\min \sum_{k \in K} F_k y_k \tag{2}$$

s.t.

$$x_{ij}^{km} \leq q_{ij} y_k \quad \forall k, m \in K | k \neq m, \ \forall i, j \in I | i \neq j \tag{3}$$

$$x_{ij}^{km} \leq q_{ij} y_m \quad \forall k, m \in K | k \neq m, \ \forall i, j \in I | i \neq j \tag{4}$$

$$\sum_{\substack{k,m \in K \\ k \neq m}} x_{ij}^{km} + w_{ij} = q_{ij} \quad \forall i, j \in I | i \neq j \tag{5}$$

$$\sum_{\substack{i,j \in I \\ i \neq j}} \sum_{m \in K} x_{ij}^{km} + \sum_{\substack{i,j \in I \\ i \neq j}} \sum_{m \in K} x_{ij}^{mk} \leq C_k \quad \forall k \in K \tag{6}$$

$$w_{ij} \geq 0, \quad x_{ij}^{km} \geq 0, \quad x_{ij}^{kk} = 0 \quad \forall i, j \in I | i \neq j, \ \forall k, m \in K | k \neq m \tag{7}$$

$$y_k \in \{0, 1\} \quad \forall k \in K \tag{8}$$

The two objective functions (1) and (2) minimize the total transportation cost associated with all transportation flows within the network and the location cost of intermodal terminals, respectively. The constraints of the model are the same as those in the model of Arnold et al. [7]. A terminal should be open for goods to be transported through this terminal (constraints (3) and (4)). The total amount of goods for each origin/destination-pair should be transported, either multimodally or unimodally (5). Goods can only be transported multimodally through a certain terminal if its capacity is not yet reached (6). Constraints (7) ensure that no demand is transported making use of only one terminal and that this demand should be non-negative. Finally, constraints (8) enforce that a terminal is either open or closed.

As with all multi-objective optimization problems, the notion of optimality has to be abandoned in favor of the notion of *dominance*. A solution is said to dominate another solution if it scores at least as well on all objective functions and better on at least one. The bi-objective intermodal terminal location model therefore does not have a single optimal solution, but rather a set called the *Pareto set* of mutually non-dominated solutions. The projection of the Pareto set in the objective function space is called the *Pareto front*.

The Pareto set and the Pareto front provide decision makers with the means to study the intermodal network on a high level, showing the best possible network for a given level of investment in terminal infrastructure. It also allows examining the relationship between the level of investment and the shift towards multimodal transportation. Additionally, the Pareto set allows decision makers to study different non-dominated solutions and their similarities and differences. In this way they can compare the effects of different levels of investment, which will in turn allow to determine the different returns on investment. Additionally, the robustness of a preferred solution can be investigated with respect to possible additional investments in the future. If the preferred solution for a higher investment level differs completely from the solution for the current investment level, it is probably wise to investigate further in which way the future expansion plan might be executed, and it might be interesting to opt for a sub-optimal solution at present, safeguarding the future performance of the network.

4. A GRASP metaheuristic for the bi-objective intermodal terminal location problem

Determining the complete Pareto set and its projection in the objective function space, the Pareto front, is NP hard for most realistic optimization problems. For this reasons, heuristics or metaheuristics are generally used to approximate the Pareto set, resulting in a set of solutions called the *Pareto set approximation*.

Most multi-objective algorithms, including those to solve facility location problems, belong to the class of genetic/evolutionary algorithms. In this paper, a simple but effective constructive algorithm is developed, based on the uni-objective GRASP heuristic from Sörensen et al. [8]. The bi-objective GRASP heuristic developed in this paper, is specifically designed to solve the bi-objective intermodal terminal location problem and extensively uses the structure of this optimization problem.

The algorithm uses two basic concepts of the problem-specific uni-objective algorithm developed in Sörensen et al. [8]. First, given a set of open terminals, the transportation costs are determined using a greedy constructive heuristic. The motivation for this is that determining the transportation costs requires the calculation of the optimal routing of all goods through the network (variables x_{ij}^{km} and w_{ij}). As has been shown in Sörensen et al. [8], this subproblem can only be solved exactly in a reasonable amount of time for small instances. This fact warrants the use of an efficient heuristic to quickly calculate the approximate transportation cost associated with a specific set of open terminals, rather than an exact method to calculate the optimal cost. A pseudo-code version of this heuristic is shown in Algorithm 1.

Algorithm 1 (Heuristic evaluation procedure [8]).

Input: $\mathbf{y} = \{y_k\}$ (status of each terminal), R (list of origin-destination pairs, sorted in decreasing order of regret),
$\mathbf{P} = \{P_{ij}\}$ (list of all possible paths for all possible terminal pairs, sorted in increasing order of cost)
Output: Routing cost RC
create list C with capacity C_k of each terminal;
$RC \leftarrow 0$;
foreach $(i, j) \in R$ **do**
    allocated $\leftarrow$ false ;
    $p \leftarrow 0$;
    $q'_{ij} \leftarrow q_{ij}$;
    while *allocated* $=$ *false and* $p < |P_{ij}|$ **do**
        $t1 \leftarrow$ terminal with lowest C_k in route p ;
        $t2 \leftarrow$ terminal with highest C_k in route p ;
        if $t1$ *and* $t2$ *are open* **then**
            if C_{t1} *and* $C_{t2} \geqslant q'_{ij}$ **then**
                allocated $\leftarrow$ true;
                $C_{t1} \leftarrow C_{t1} - q'_{ij}$;
                $C_{t2} \leftarrow C_{t2} - q'_{ij}$;
                $RC \leftarrow RC + c_{ij}^{p} q'_{ij}$;
            else
                $RC \leftarrow RC + C_{t1} \times c_{ij}^{p}$;
                $q'_{ij} \leftarrow q'_{ij} - C_{t1}$;
                $C_{t2} \leftarrow C_{t2} - C_{t1}$;
                $C_{t1} = 0$;
        $p \leftarrow p + 1$;
    if *allocated* $=$ *false* **then**
        $RC \leftarrow RC + c_{ij} q'_{ij}$;

The so-called *heuristic evaluation procedure* allocates a flow of goods to either two terminals (intermodal flow) or directly between customers (unimodal flow). The main idea behind the heuristic is to give preference to the origin/destination-pairs that have the highest opportunity cost of not being allocated to their preferred route. The *regret* is used, calculated as the difference in cost between the route (i.e., pair of terminals) with the lowest cost and that with the second lowest cost, for any pair of customers. The algorithm then attempts to allocate (partial) flows to the best route possible (i.e., a route for which (1) the terminals are opened and (2) sufficient capacity is left in both terminals). If no multimodal route is possible, or if the cheapest possible multimodal route is more expensive than unimodal shipment, the goods are shipped unimodally.

A second concept aims to use GRASP to construct good solutions. GRASP [15], or *greedy randomized adaptive search procedure*, is a multi-start metaheuristic of which the objective is to build a feasible solution by starting from an empty solution and adding one element at a time. Typical of the GRASP metaheuristic is that it balances greediness (selecting the best element at each iteration) and randomness (selecting a random element at each iteration) and is able to combine the advantages of both approaches. The GRASP heuristic for the terminal location problem starts from a solution in which all terminals are closed and opens one terminal at a time. A pseudo-code description of the GRASP algorithm is given in Algorithm 2.

Like its uni-objective counterpart, the bi-objective GRASP performs a fixed number of iterations maxiter. In each iteration, the heuristic starts from an empty solution (with all terminals closed, implying the lowest possible location cost) and opens one terminal at a time. The final solution in each iteration is the one in which all terminals have been opened, which is therefore the solution (with the lowest transportation cost). Moving from the

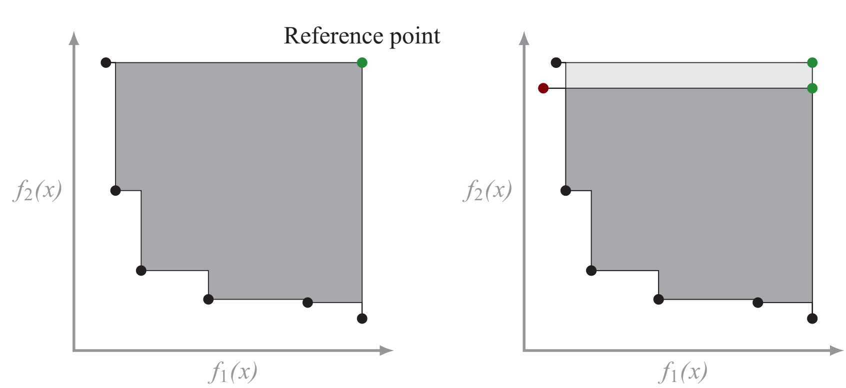

Fig. 2. The hypervolume, using the maximum value for each of objective as the reference point.

solution with the lowest location cost to the one with the lowest transportation cost implies that the procedure automatically generates solutions with different location cost/transportation cost ratios at each GRASP iteration.

The quality of the solutions produced by the heuristic is therefore only determined by the order in which the terminals are opened. Since high capacity and low location cost are both favorable qualities of a terminal, the GRASP procedure favors opening terminal that have these attributes.

Algorithm 2 (GRASP algorithm pseudo-code).

Output: Archive of non-dominated solutions
create and sort list L with all terminals k according to $r_k = \dfrac{F_k}{C_k}$ in ascending order;
add solution with all terminals closed to archive;
$i = 0$;
while $i < maxiter$ **do**
  $i \leftarrow i + 1$;
  $s \leftarrow$ solution s with all terminals closed;
  foreach *terminal* $k \in K$ **do**
    create Restricted Candidate List $\bar{L}$ of α best terminals (closed, with smallest value of r_k);
    randomly select terminal $t \in \bar{L}$;
    remove t from L ;
    open t in s;
    if *s is non-dominated* **then**
      add s to archive;
    remove all solutions dominated by s from archive;

The order in which the terminals are opened by the GRASP algorithm, depends on the location cost/capacity-ratio of each terminal. For each terminal $k \in K$, this ratio is given by $r_k = F_k/C_k$. At each step of the GRASP procedure, the RCL (or restricted candidate list) consists of the α best terminals according to the r_c ratio. α is a parameter of the heuristic that determines the balance between greediness and randomness and that can be changed at each iteration. The reason for changing α is to allow the search to explore different parts of the solution space by changing the greediness–randomness balance.

At each step, a terminal is randomly selected from the restricted candidate list and the transportation cost of the newly constructed solution is calculated using the heuristic evaluation procedure described earlier in this paper.

The solution is then compared to every other solution in the archive and added if it is not dominated by one of those solutions. If the solution is dominated by a solution already present in the archive, it is discarded. In the next step, the archive is checked for solutions that are dominated by the new solution, and these solutions are discarded as well. At each step of the GRASP procedure, the archive therefore only contains non-dominated solutions.

Finally, the restricted candidate list is updated to make sure that the element under consideration is no longer available for selection and the above process is repeated, until all terminals have been added.

The GRASP algorithm has two parameters that need to be set: the threshold parameter α that determines the balance between randomness and greediness in the construction phase, and the maximum number of iterations maxiter. Given the fact that a large diversity in solutions found by the GRASP heuristic is desirable, the α parameter is allowed to vary between 0 (completely greedy construction) and $|K|$ (completely random). At each iteration iter, the value of α is determined as

$$\alpha = \lceil \frac{\text{iter}}{\text{maxiter}} |K| \rceil . \tag{9}$$

By setting α as a function of maxiter the number of parameters of the GRASP algorithm is reduced to one.

5. Computational results

5.1. Algorithm analysis

To quantify the quality of the output of a multi-objective algorithm is notoriously difficult, and several measures exist in the literature, each with their own properties. A complete review is far beyond to the scope of this paper, and we refer to Zitzler et al. [16] for an overview of some of the best-known measures, including a discussion of their properties, drawbacks and advantages. In general, the aim of a multi-objective performance indicator is to measure how closely the Pareto set approximation approximates the true Pareto set.

To assess the performance of the GRASP heuristic proposed in this paper, the compliance to the three goals of a multi-objective search [16] is measured.

A first goal stipulates that the number of non-dominated solutions found by the heuristic should be as large as possible. Secondly, the distance of the Pareto front approximation found by the algorithm should approximate the true Pareto front as much as possible. Thirdly, the solutions should be evenly spread along the Pareto front approximation.

To measure the size, we simply record the number of non-dominated solutions found by the heuristic (i.e., the *cardinality* of the Pareto set).

Secondly, the *hypervolume* indicator is used. This measure was introduced by Zitzler and Thiele [17] and is defined as the hypervolume enclosed by the hyperplane defined by the Pareto front and some reference point. For a bi-objective optimization problem, this boils down to the surface area defined by the Pareto front and some reference point. The reference point is often taken to be the point defined by the maximum value found on each objective function (see Fig. 2). This has the advantage of not requiring an arbitrary reference point to be defined. One of the drawbacks of this approach is that finding new solutions that dominate existing solutions can actually make the value of the hypervolume indicator decrease. Another drawback of the hypervolume indicator is its computational complexity. The calculation of the hypervolume indicator has been shown to be exponential in the number of objectives [20]. For bi-objective problems, like the one discussed in this paper, this is relatively harmless. It should be noted that the hypervolume will also improve if more solutions are found and the points are distributed evenly.

Thirdly, the quality of the distribution of the obtained solutions along the Pareto front is measured using Schott's spacing metric [18], an indicator of how evenly the solutions are distributed. Schott's spacing metric is defined as the standard deviation of the Manhattan distances between all solution and their closest solution. The smaller the value of this metric, the less deviation and thus the more evenly spread the Pareto front is.

Let f_m^i be the objective function value of the ith solution on the mth objective, and let N be the number of non-dominated solutions on the Pareto front, then Schott's spacing metric is defined as:

$$\sqrt{\frac{1}{N}\sum_{i=1}^{N}(d_i - \bar{d})^2} \tag{10}$$

with

$$d_i = \min_{j \in [1,N] \wedge i \neq j} \sum_m |f_m^i - f_m^j| \tag{11}$$

5.2. Efficient solutions

As mentioned, the algorithm only has a single parameter, the number of GRASP iterations maxiter. In general, a larger number of GRASP iterations leads to better efficient solutions, since new solutions will generally be found that dominate solutions already in the archive. To illustrate this, Fig. 3(a), (b) and (c) shows the efficient fronts found for 10, 100, 1000, and 10,000 iterations of the GRASP procedure, for test sets with 30 customers and 30 potential locations, 50 customers and 50 potential locations, and 100 customers and 100 potential locations, respectively.

Visual inspection of Fig. 3(a)–(c) reveals that the quality of the Pareto front does indeed improve with an increasing number of GRASP iterations. Note that the four fronts in each of the figures were generated with different runs of the GRASP heuristic. This means that some solutions found by the GRASP heuristic with, say, 10 runs, might actually dominate some solutions found by the GRASP heuristic with 100 runs. However, visual inspection of the figures clearly shows that the overall quality of the Pareto can be said to improve as a function of the number of GRASP iterations.

Quantitative analysis of the quality of the Pareto set approximations confirms these conclusions. Figs. 4–6 show the evolution during a run of 10,000 GRASP iterations of the number of non-dominated solutions, the hypervolume, and Schott's spacing metric for two different problem instances (30 customers/30 locations and 50 customers/50 locations).

The number of non-dominated solutions on the Pareto front can be seen to evolve towards a stable number of around 100–120 (see Fig. 4(a) and (b), for the cases of 30 customers/30 locations and 50 customers/50 locations, respectively).

As shown in Fig. 5(a) and (b), the hypervolume clearly increases as a function of the number of iterations. The "jumps" in Fig. 5(b) illustrate how the discovery of a new solution that dominates one of the solutions that determine the reference point, can make the value of the hypervolume decrease or increase abruptly. In between such jumps, however, the hypervolume steadily increases as better solutions are found by the heuristic.

Fig. 6(a) and (b) shows the evolution of Schott's spacing metric as a function of the number of iterations. Since this metric is a standard deviation, it has the same units as a distance in the objective function space. Clearly, a good spread corresponds to a small value, relative to the average value of a distance. For 30 customers and 30 locations (see Fig. 2), the difference between the maximum and the minimum transportation cost is approximately 1×10^8. For the location cost this is 3.5×10^7. Distances in the objective space, as well as the standard deviation of those distances, should therefore be evaluated on the 10^7–10^8 scale. The value of Schott's spacing metric for this data set is around 6×10^5, which is several orders of magnitude smaller. This fact shows that the solutions are nicely spread out along the Pareto front. The same holds for the case with 50 customers and 50

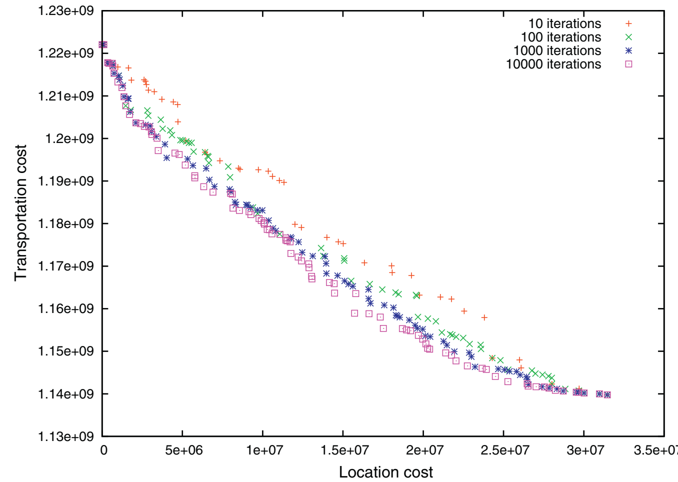

(a) 30 customers and 30 potential locations

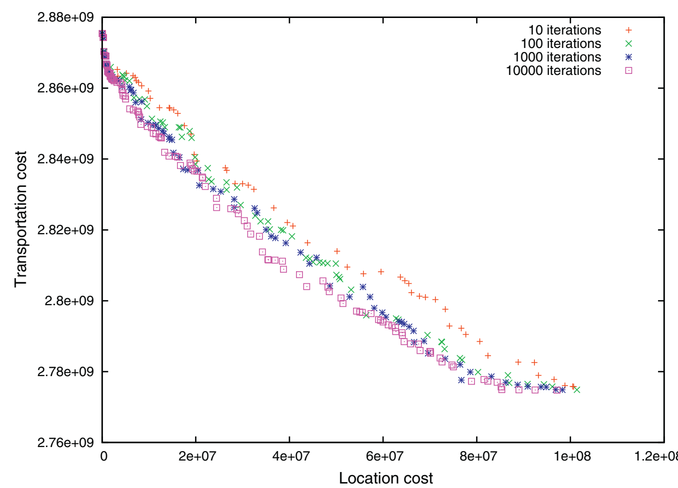

(b) 50 customers and 50 potential locations

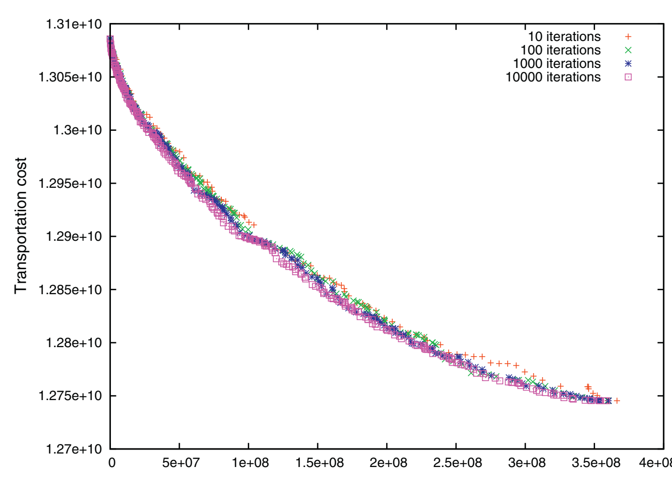

(c) 100 customers and 100 potential locations

Fig. 3. Pareto front approximations for 10, 100, 1000, and 10,000 iterations of the GRASP heuristic.

objectives (compare the order of magnitude of the objectives in Fig. 3(b) with that of Fig. 6(b)).

The values in Fig. 6(b) show the same jumps as the hypervolume, indicating that Schott's spacing metric is similarly

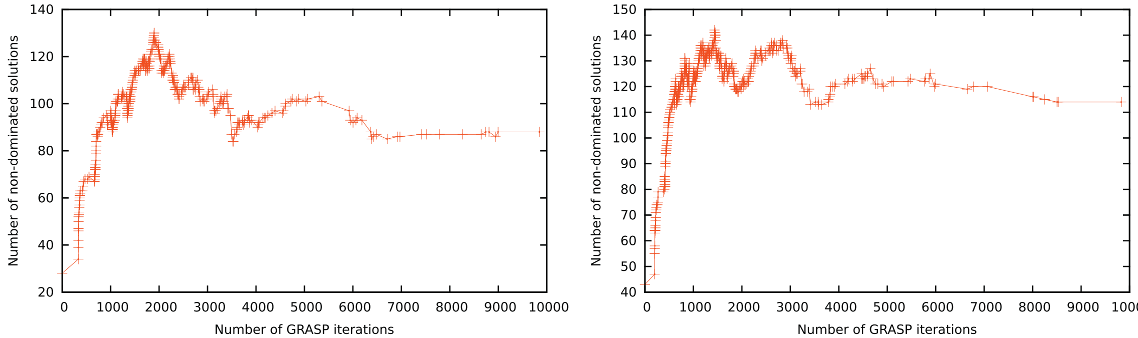

(a) 30 customers and 30 potential locations

(b) 50 customers and 50 potential locations

Fig. 4. Evolution of the number of non-dominated solutions on the Pareto front for a run of 10,000 iterations of the GRASP algorithm.

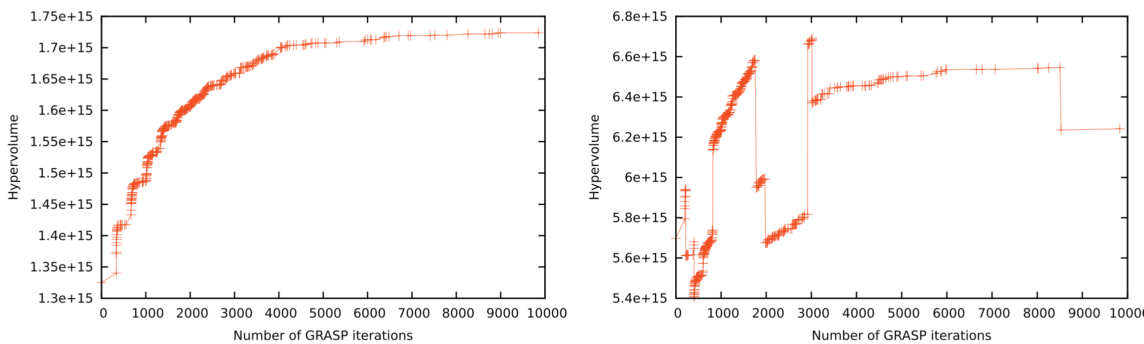

(a) 30 customers and 30 potential locations

(b) 50 customers and 50 potential locations

Fig. 5. Evolution of the hypervolume of the Pareto front for a run of 10,000 iterations of the GRASP algorithm.

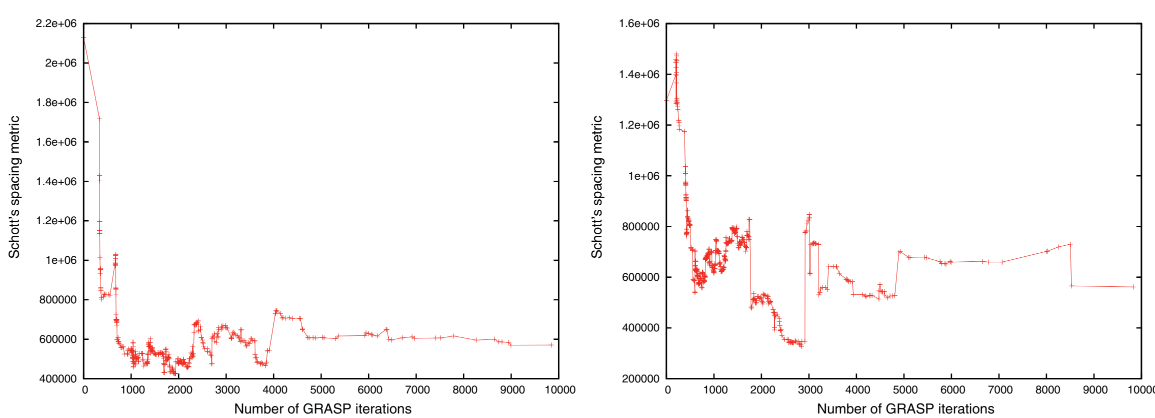

(a) 30 customers and 30 potential locations

(b) 50 customers and 50 potential locations

Fig. 6. Evolution of Schott's spacing metric for a run of 10,000 iterations of the GRASP algorithm.

affected by finding solutions that dominate one of the solutions that define the hypervolume.

5.3. CPU time analysis

An analysis of the CPU time required by the bi-objective GRASP algorithm can be found in Fig. 7, that shows the CPU time required for 10 iterations for different instance sizes. Given the fact that different iterations of the GRASP algorithm are independent of each other, the calculation time for longer or shorter runs can be easily estimated from the corresponding data point in this graph. Given the strategic importance of the intermodal terminal location decision, computing times are very reasonable, especially given the computational complexity of the underlying optimization problem. This is largely due to the fact that the GRASP metaheuristic efficiently exploits the specificities of the intermodal terminal location problem. Especially the use of an approximate evaluation function strongly reduces the computing time (see Sörensen et al. [8], for a more elaborate discussion).

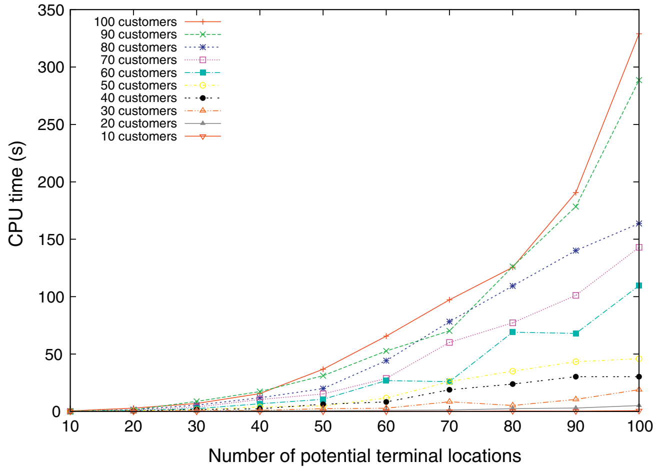

Fig. 7. CPU time (in seconds) for 10 iterations of the GRASP heuristic for problems of different size.

6. Conclusions and future research

In this paper, a bi-objective version of the intermodal terminal location problem is developed. An efficient bi-objective GRASP algorithm was developed, that generates a set of mutually non-dominated solutions. The usefulness of this approach in analyzing an intermodal terminal network is twofold.

First of all, it allows the decision-maker to make a distinction between the costs paid by the users (companies shipping goods through the network) and the costs paid by the terminal operators. It therefore allows an estimation of the effects of different levels of investments. It also allows a more thorough analysis of the cost structure of the intermodal terminal location problem, by allowing the decision maker to analyze different solutions that are located in each other's vicinity on the Pareto front.

Secondly, the algorithm is user-friendly, simple, and has only a single parameter: the number of iterations of the GRASP procedure. By adjusting this parameter, the decision-maker can choose the right balance between calculation time and quality of the Pareto front.

In the future the model could be adjusted to include more realistic cost structures. An analysis of the robustness of the solutions found is also necessary: some solutions may drastically change when changes in demand and supply in the network occur and may for that reason be disregarded in favor of other solutions. Costs could also include greenhouse gas emissions, noise levels, etc.

References

[1] UNECE, Illustrated Glossary for Transport Statistics. United Nations Economic Commission for Europe, 2009.

[2] G. Bühler, P. Jochem, CO_2 emission reduction in freight transports: how to stimulate environ mental friendly behaviour? ZEW discussion papers, ZEW – Zentrum für EuropäischeWirtschaftsforschung/Center for European Economic Research, 2008.

[3] M.B. Pedersen, O.B.G. Madsen, O.A. Nielsen, Optimization models and solution methods for intermodal transportation, Ph.D. Thesis, Technical University of Denmark, Centre for Traffic and Transport, Traffic Modelling Group, 2005.

[4] D.J. Bowersox, D.J. Closs, O.K. Helferich, Logistical Management: A Systems Integration of Physical Distribution, Manufacturing Support, and Materials Procurement, MacMillan, Michigan, 1986.

[5] F. Witlox, E. Vandaele, Determining the monetary value of quality attributes in freight transportation using a stated preference approach, Transportation Planning and Technology 28 (2) (2005) 77–92.

[6] European Commission, Cost–benefit and multi-criteria analysis for nodal centres for goods, Technical Report, Office for Official Publications of the European Community, 1996.

[7] P. Arnold, D. Peeters, I. Thomas, H. Marchand, Pour une localisation optimale des centre de transbordement intermodaux entre re'seaux de transport: formulation et extensions, The Canadian Geographer 45 (3) (2001) 427–436.

[8] K. Sörensen, C. Vanovermeire, S. Busschaert, Efficient metaheuristics to solve the intermodal terminal location problem, Computers & Operations Research 39 (9) (2012) 2079–2090.

[9] A. Sirikijpanichkul, L. Ferreira, Multi-objective evaluation of intermodal freight terminal location decisions, in: Proceedings of 27th Conference: Australian Institute of Transport Research, 2005.

[10] A.M. Moreira, R.A. Ribeiro, E. Declercq, O. Schinas, P. Guerreiro, G.K. Janssens, Optimal location of intermodal terminals in Europe: an evaluation model, in: Proceedings of the 6th European Congress on Intelligent Techniques and Soft Computing, EUFIT'98, vol. 2, September, (1998), pp. 1057–1061.

[11] B. Groothedde, Collaborative Logistics and Transportation Networks: a modeling approach to hub network design, Ph.D. Thesis, Technische Universiteit Delft, Delft, November 2005.

[12] C. Macharis, A. Verbeke, Een multicriteria-analyse methode voor de evaluatie van intermodale terminals, Tijdschrift Vervoerswetenschap 4 (1999) 323–341.

[13] C. Macharis, A. De Witte, J. Ampe, The multi-actor, multi-criteria analysis methodology (MAMCA) for the evaluation of transport projects: theory and practice, Journal of Advanced Transportation 43 (2) (2009) 183–202.

[14] R.Z. Farahani, M. SteadieSeifi, N. Asgari, Multiple criteria facility location problems: a survey, Applied Mathematical Modelling 34 (7) (2010) 1689–1709. , http://dx.doi.org/10.1016/j.apm.2009.10.005, ISSN 0307-904X.

[15] T.A. Feo, M.G.C. Resende, A probabilistic heuristic for a computationally difficult set covering problem, Operations Research Letters 8 (2) (1989) 67–71.

[16] E. Zitzler, J. Knowles, L. Thiele, Quality assessment of pareto set approximations, Multiobjective Optimization (2008) 373–404.

[17] E. Zitzler, L. Thiele, Multiobjective optimization using evolutionary algorithms – a comparative case study, in: A.E. Eiben, T. Bäck, M. Schoenauer, H.-P. Schwefel (Eds.), PPSN 1998, Volume 1498 of LNCS, Springer, Heidelberg, 1998, pp. 292–301.

[18] J.R. Schott, Fault Tolerant Design Using Single and Multicriteria Genetic Algorithm 360 Optimization, Ph.D. Thesis, Masters Thesis, Department of Aeronautics and Astronautics, MIT Cambridge, MA, USA, 1995.

[19] A. Goel, A roadmap for sustainable freight transport, in: F. Heyde, A. Löhne, C. Tammer (Eds.), Methods of Multicriteria Decision Theory and Applications, Shaker, Maastricht, 2009, pp. 47–56.

[20] L. While, A new analysis of the lebmeasure algorithm for calculating hypervolume, in: C.A. Coello Coello, A. Hern'andez Aguirre, E. Zitzler (Eds.), EMO 2005, Volume 3410 of LNCS, Springer, Heidelberg, 2005, pp. 326–340.

Kenneth Sörensen is professor at the Faculty of Applied Economics of the University of Antwerp, where he also obtained his PhD. He chairs the research group ANT/OR (University of Antwerp Operations Research Group), a group that focuses on applications of Operations Research in diverse domains, and on the development of optimization algorithms, especially metaheuristics. Kenneth Sörensen has published extensively on applied (heuristic) optimization and is a leading expert on the development of metaheuristics. He also founded the largest working group on metaheuristics: EU/ME - the metaheuristics community.

Christine Vaovermeire graduated as a commercial engineer in management of information systems at the University of Antwerp, in 2009. She started immediately as a supply chain management trainee at Cargill N.V., but when she was offered to prepare her PhD at the University of Antwerp, she decided to return. After meeting TRI-VIZOR, supply chain orchestrator for horizontal collaboration, she decided that this would be her topic of research. In January 2011, she obtained a scholarship of the agency for Innovation through Science and Technology. Since then, she is working on various topics on horizontal collaboration, such as cost allocation, partner selection, planning of the collaboration and so on.