International Journal of
Operations Research

International Journal of Operations Research Vol. 5, No. 1, 1−12 (2008)

A Survey of Solution Methods for the Continuous Location-Allocation Problem

Jack Brimberg[1, *], Pierre Hansen[2], Nenad Mladenović[3], and Said Salhi[4]

[1]Department of Business Administration, Royal Military College of Canada, Kingston, Ontario, Canada, and GERAD

[2]GERAD and HEC Montreal, University of Montreal, Canada

[3]School of Mathematics, Brunel University, United Kingdom, and GERAD

[4]Centre for Heuristic Optimisation, Kent Business School, University of Kent at Canterbury, Kent, United Kingtom

Received November 2006; *Revised March 2007*; *Accepted April 2007*

Abstract—In this survey, we examine an important class of facility location problems known as the multisource Weber problem (also referred to as the continuous location-allocation problem). We also show how recent advances in the use of metaheuristic rules have significantly improved the ability to solve problems of this type. The new solution methods are discussed for both the well known multisource Weber problem and its counterpart the capacitated case. Research issues which we believe to be worthwhile exploring in future are also highlighted.

Keywords—Location, Continuous space, Heuristics

1. INTRODUCTION

The objective in continuous facility location problems is to "generate" sites in continuous space for locating new facilities relative to a set of existing facilities situated at given points in the space. Thus, continuous location models are also referred to as "site-generating" models in contrast to their discrete counterpart where candidate sites are identified beforehand as nodes on a graph ("site selection" models). The existing facilities typically represent customers or markets, and are labeled as demand points or fixed points. Another fundamental difference between continuous and discrete location models is that in the former case a function must be selected to estimate distances (costs) between points in the space, whereas actual travel distances (costs) may be used in the second case. On the one hand continuous models are easy and fast to set up in practice since large databases of travel distances, and computations of shortest paths, are eliminated from the start; on the other hand, they are inherently less precise than their discrete versions since travel distances (times, costs, ...) can only be approximated by the distance function.

In several situations, the loss in accuracy may be quite acceptable to the practitioner (decision-maker). For example, in an urban environment the rectangular distance function (also known as Manhattan or city-block distances) may be a highly-accurate predictor of actual distance. Cities whose road networks are characterized by a rectangular grid fall into this category. In geographical regions with dense, highly-developed transportation networks, the Euclidean (or straight line) distance multiplied by an inflation factor may provide an adequate measure of travel distance. Examples here include countries such as England and France that have highly sophisticated road and rail networks in place. In other cases, a tailored distance function may be acceptable. The fitting of empirical functions to actual travel data has received wide attention in the literature, and as result, some sophisticated distance functions have been developed for this purpose. The point made is that the price to pay for lost accuracy may be very small indeed when weighing the other advantages of the continuous location model.

Some software packages make use of distance functions to compute the distance between different locations on a network. Examples include truck dispatching models for vehicle routing and Geographic Information Systems (e.g., see Vine et al. (1997)). It makes sense in these cases to complement the software with algorithms for continuous location models. Another aspect is that the system under consideration does not always comprise a network. Air travel is an obvious case in point; here models that use great circle distances would be a natural choice (e.g., see Wesolowsky (1982), Hansen et al. (1995), Das et al. (2001), and Patel and Chidambaram (2002) for site generation models with geodesic distances). Other examples arise in related fields of study such as clustering (e.g., Taillard (2003)), where continuous models with distance functions may be applied (including the one we will be investigating) in higher dimensional spaces. Such applications are far removed from the traditional problem of locating

[*] Corresponding author's email: Jack.Brimberg@rmc.ca

warehouses (or distribution centers) that gave the original impetus to the study of location models. It should also be noted that cases arise where physical assets may be located anywhere in a continuous region. Love et al. (1988) give the example of locating wells in an irrigation system where the objective is to minimize the cost of connecting pipes. This problem is formulated as a continuous minisum location model of the type given below. Other examples include the location of transmission towers, relay stations, search and rescue headquarters, helicopter serviced emergency medical units, long range facilities location planning, and a host of other applications (e.g., see Ostresh (1975) and Hodgson et al. (1993)). As noted by Ostresh (1975), in the case of long range planning, "today's costs and times are probably a less reliable estimate of future costs than measures derived from distances".

Large scale location problems readily lend themselves to modeling as continuous location problems, because of the relatively modest requirements for input data. Essentially the information required consists of the coordinates of the fixed points, the weights or demands associated with each one, and a properly selected distance function. Existing large scale applications with real data include the transshipment center location problem with 20 new facilities and 1700 customers in Bhaskaran (1992), and the districting problem with 170 new facilities and 1400 existing ones in Fleischmann and Paraschis (1988). Both cases were modeled as continuous location-allocation problems, the focus of our review. Hidaka and Okano (2003) provide a very large scale study dealing with spare parts logistics for a Japanese manufacturing company with 6000 customers and 380,000 potential warehouse sites. They are able to model the network in detail using a digital map with hundreds of thousands of nodes and branches and actual distances of each branch. Transportation costs are estimated by the shortest path distance between two points, but the authors note that this requires too much computing time. As a result, they scale down the problem in a pre-processing step that selects a small subset of the candidate sites. The scaling down or aggregation of data could have been avoided by using the continuous analogue model with "tailored" distance function, thus avoiding the need to calculate shortest paths and store a huge matrix of actual distances. There is a tradeoff here between aggregation errors and distance approximation errors that should be accounted for. In the end, the continuous models may be not only simpler to use, but just as accurate! It should also be noted that the "look up" advantage of discrete models where actual freight costs are used, appears to be lost when large scale applications are considered.

With the preceding rather lengthy discussion as a motivating force, we turn now to the main topic of the review, the multisource Weber problem, also referred to as the continuous location-allocation problem. This well-known and much studied model is the analogue of the discrete p-median problem (see Mladenovic et al. (2007) for a recent survey of the discrete model). The continuous problem requires the generation of a given number m of facility sites in the plane $(\mathbb{R}^2)$ to serve the demands of n customers or fixed points in order to minimize the total transportation (or service) cost. Under the assumption that transportion cost is proportional to distance traveled, the uncapacitated version of the multisource Weber problem (MWP) may be formulated as follows (e.g., Love et al. (1988)):

$$\text{(MWP)} \min_{w,x} \sum_{i=1}^{m} \sum_{j=1}^{n} w_{ij} \cdot \left\| X_i - A_j \right\|$$

subject to

$$\sum_{i=1}^{m} w_{ij} = w_j, \; j = 1,2,\ldots,n, \tag{1}$$

$$w_{ij} \geq 0 \quad i = 1,2,\ldots,m \quad j = 1,2,\ldots,n, \tag{2}$$

where $A_j = (a_j, b_j)$ is the known location of customer j, $j = 1, \ldots, n$; $X = (X_1, \ldots, X_m)$ designates the vector of location decision variables, with $X_i = (x_i, y_i)$ being the unknown location of facility i, $i = 1, \ldots, m$; $w_j > 0$ is the given demand rate of customer j, $j = 1, \ldots, n$; $W = (w_{ij})$ designates the vector of allocation decision variables, where w_{ij} gives the flow to customer j from facility i, $i = 1, 2, \ldots, m$, $j = 1, 2, \ldots, n$; and

$$\left\| X_i - A_j \right\| = \sqrt{(x_i - a_j)^2 + (y_i - b_j)^2}$$

is the Euclidean distance between facility i and customer j.

The objective function is seen to represent the total service cost. The constraint set ensures that all the customer demands are satisfied. The basic model assumes Euclidean travel distance, but other distance functions, such as the Manhattan (or l_1) norm or, more generally, the l_p norm, have also been employed (e.g., see Love et al. (1988), Francis et al. (1992) and Drezner (1995) for a review). An advantage to using block norms as l_1 is that the candidate facility locations may be reduced to a finite number of points formed by the intersection of reference lines through the fixed points. In this case, the problem is equivalent to the discrete m-median problem. With round norms, the facility locations may vary continuously in the plane as the parameters, such as the weights, are varied. The multisource Weber problem formulated above and its several variants constitute one of the most-studied problems in continuous location theory.

The main difficulty in solving MWP arises from the non-convexity of the objective function and the existence of multiple local minima. As a result, this problem falls in the realm of global optimization. Equivalently, the problem may be viewed as an enumeration of the Voronoi partitions of the customer set which is shown to be NP-hard by Megiddo and Supowit (1984). In their well-known 50 customer problem, Eilon et al. (1971) were able to generate 61 local minima from 200 random restarts of Cooper's heuristic (discussed below) for the case $m = 5$. It was proven much later by Krau (1997) that their best solution was indeed the global optimum. The fact that the worst solution deviated from the best by some 40% showed quite dramatically the danger of being "trapped" in a local optimum. The same problem is investigated in Brimberg et al. (2004a): with 10,000 random iterations for m = 5, 10, 15 the different local solutions obtained numbered 272, 3008, 3363, respectively. It was further observed that the worst deviation from the optimal solution was in the respective cases, 47%, 66% and 70%; also the optimal solution was obtained 690 times for m = 5, 34 times for m = 10 and only once for m = 15. These relatively small instances are able to demonstrate another important characteristic of this class of problem, namely, a tendency for the number of local minima to increase exponentially with problem size. The existence of degenerate local solutions in which some of the facilities have no customers assigned to them further complicates the problem (see Brimberg and Mladenović (1999)).

The remainder of the paper is organized as follows. In the next section we give an overview of exact solution methods for the multisource Weber problem. This is followed in section 3 with a discussion of approximate methods (heuristics). In section 4 we present the case where the facilities have limited capacities. As there is a shortage of published work in this area we summarise the research in one section only. Our conclusions and possible research avenues are given in section 5. In this paper, since the exact methods are limited in application to relatively small problem sizes, the emphasis is on heuristic approaches. We aim to provide the reader with an appreciation of the "richness" in OR/MS/CS techniques applied here, and where we may be headed in terms of new approaches to solving ever-larger problem sizes.

2. EXACT METHODS

Exact solution methods have limited applicability due to the non-polynomial (*NP*) nature of the problem. However, as in other areas, some specific cases have been identified that allow for a significant simplification. For example, an efficient solution procedure has been developed for the case of two facilities (m = 2), based on the property that the customer set may be partitioned by a straight line intersecting at least two of the fixed points (Ostresh (1975) and Drezner (1984)). Thus, $O(n^2)$ candidate solutions are identified that must contain an optimal solution. Computation times are reported in Drezner (1984) for up to n = 100. A more efficient procedure based on difference of convex functions programming, known as d-c programming for short, is given by Chen et al. (1998) for the m = 2 case, with near linear computation times observed for up to n = 1000 customers.

Another special case occurs when all the fixed points are located on a straight line, or main travel is restricted to one dimension as along a river or highway. Such problems may be solved efficiently by a dynamic programming algorithm by Love (1976). A special class of two-dimensional problems is identified by Brimberg and Love (1998) that also permits a dynamic programming approach.

Branch-and-bound algorithms have been developed for the general problem (Kuenne and Soland (1972) and Ostresh (1973)), but these have been used to solve only very small instances, of the order of n = 15, m = 4 and n = 50, m = 3. Rosing (1992) is able to incorporate improvements in the methodology that allow problems with n = 30, m = 5 and n = 25, m = 6 to be solved exactly. The general idea is to develop all convex hulls (Voronoi partitions) and then cover the fixed points with a set of disjoint convex hulls. The problem may then be converted to a large partitioning problem where each customer must belong to exactly one subset. The problem size is very limited due to the exponential rate of increase in the number of subsets. In the case of rectangular distances, the problem may be converted to a discrete m-median problem, and the set of candidate solutions may be further reduced by examining the hull properties (Love and Morris (1975)).

A significant advance is reported in Krau (1997), who uses a column generation approach combined with global optimization and branch-and-bound. The author is able to obtain exact solutions of instances with up to n = 287 customers (ambulance problem from Bongartz et al. (1994)) and up to m = 100 facilities by utilizing a dual formulation which is also equivalent to a concave minimization problem. A bundle method in the l_1-norm (du Merle et al. (1997)) is added to stabilize solution of the dual, leading to a very effective algorithm in Hansen et al. (1997) that successfully solves problems up to size (n = 1000, m = 100). Both methods, column generation and l_1-norm bundle, are highly sensitive to the starting solution, and so, state-of-the-art heuristics are used to obtain the best possible initial solution.

3. HEURISTIC APPROACHES

As seen in the previous section, the exact methods are generally restricted to small problem size. Recent developments in this area have significantly improved the ability to find global solutions, but even so, good heuristics are required to initialize the methods. For large scale problems that exist in practice, the only reasonable approach is solving by approximate methods.

For purposes of presentation we classify the heuristics developed for this problem into two standard categories which we term "classical" heuristics and "metaheuristics". Each of these groups may be further subdivided. In the multisource Weber problem, the first group is characterized predominantly by the property that the search terminates at a single "local" solution; that is, the heuristic examines only a narrow region in the solution space. The second group uses metaheuristic rules to expand the search in a broader region of the solution space, and thus, escape "local optima traps". Using terminology from metaheuristics, these methods are able to "diversify" the search. In fact, several of these methods have been shown in theory to be globally convergent. Since the classical methods generally comprise the older heuristics, we will discuss them first.

3.1 Classical heuristics

The first heuristic approach to solve MWP was proposed by Cooper (1964). This popular heuristic is based on the simple observation that the two components of the problem – location of the facilities and allocation of demands – are easy to solve in isolation. That is, given the customer allocations (or partition of the customer set), the problem reduces to m convex single facility location problems which may be solved efficiently using the well-known Weiszfeld iterative procedure (e.g., see Weiszfeld (1937), Kuhn (1973), Rosen and Xue (1991), Drezner (1992), and for extensions to l_p distances, Brimberg and Love (1993) and Frenk et al. (1994); also see Wesolowsky (1993) for an interesting historical review of the single facility (Fermat-Weber) problem), or by standard descent methods. Since there are no capacity constraints on the facilities, once their locations are given, each customer is allocated to the nearest one, with ties broken arbitrarily. The Cooper heuristic alternates between location and allocation steps until no further improvement is possible. Computational experience shows that a local minimum is attained typically after a small number of location-allocation iterations. A multi-start version repeats Cooper's method from random starting points until a stopping condition such as a limit on execution time is met, and retains the best local minimum as the final solution. This early method remains a basis of comparison for other heuristics; the steps are outlined in Figure 1. Variants of Cooper's method may be found in Scott (1970) and Baxter (1981), while Sullivan and Peters (1980) propose a method to cluster customers into mutually exclusive subsets, and then locate a facility in each one.

The heuristic by Love and Juel (1982) is the first method that imposes a neighborhood structure on the problem. Here a given neighborhood is defined as the set of points around a current solution that are obtained by exchanging a specified number of assignments of customers from their current facilities to new ones. The authors consider up to two exchanges, and show that the two-exchange neighborhood may be used successfully to "jump out" of a local optimum trap in the one-exchange neighborhood. This comes, of course, at a computational cost, and the method will not be able to escape from deeper troughs. To illustrate, consider a simple example given in Brimberg and Mladenović (1996a) with n = 6 customers at randomly-generated points, m = 4 facilities, and all weights w_j = 1. For convenience we reproduce the example in Figure 2. The partition of the customer set into subsets {1}, {3, 5}, {2, 4}, and {6}, allocated respectively to facilities 1, 2, 3, and 4, is a local minimum in the one-exchange neighborhood. (It is also a local minimum in the (X, W) solution space.) However, the partition {{1, 6}, {3, 5}, {2}, {4}} obtained by two exchanges, gives a better solution and the global optimum. In this method, the facilities are always located optimally with respect to a given allocation of the customers.

Step 1.	Choose m facilities at random, $X = (X_1, X_2, ..., X_m)$, as an initial solution.
Step 2.	Fixing X, assign each customer to its nearest facility (breaking ties arbitrarily) to obtain the corresponding allocation vector $W = (w_{ij})$.
Step 3.	Fixing W, find new improved sites Y for the facilities, by solving the corresponding m single facility location problems. If at least one facility site has changed ($Y \neq X$), set $X \leftarrow Y$, and return to *Step* 2; otherwise the current solution (Y, W) is a local minimum.
Step 4.	Repeat *Steps* 1 to 3 K times (or until a limit on execution time or given number of iterations without improvement is reached); retain the best solution from all local minima thus obtained.

Figure 1. Multi-start Cooper alternate heuristic.

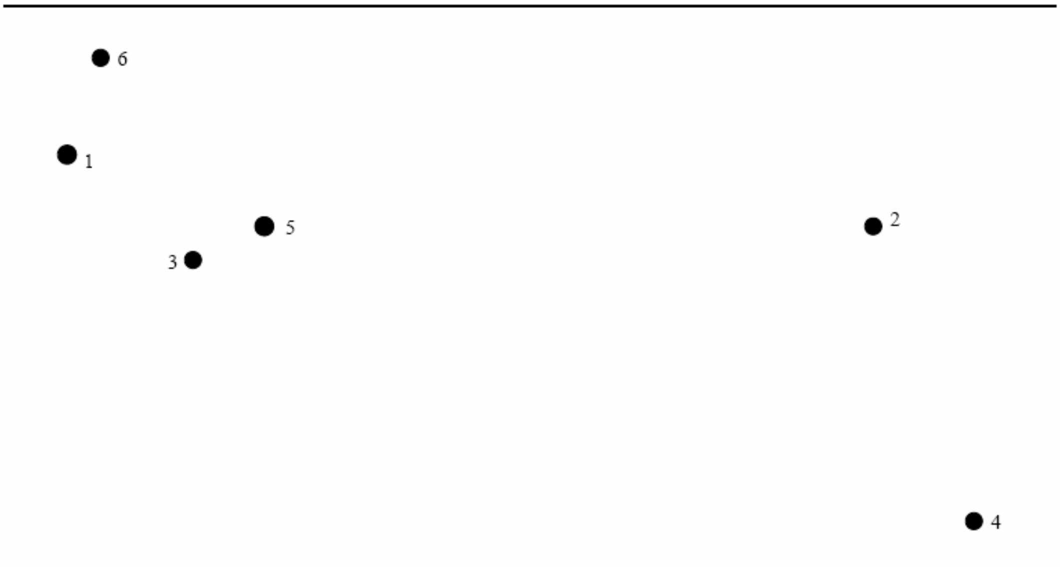

Figure 2. Example with n = 6.

Murtagh and Niwattisyawong (1982) use MINOS, a large-scale nonlinear programming package, to solve simultaneously both the location and allocation variables. As the iterations proceed, their heuristic fixes allocation variables (w_{ij}) that either attain a value of 0 or w_j. The update on the remaining free variables uses a quasi-Newton approximation of the Hessian matrix, and at nondifferentiable points, a subgradient given by Kuhn (1973).

Chen (1983) develops an interesting approach based on an approximation scheme by Charalambous and Bandler (1976). The distances between customers and facilities are each raised by an exponent ($-N$), and an exponent ($-1/N$) is then applied to each sum of modified distances from all facilities to a customer. For large enough N, these sums approach the distance between each customer and its closest facility, thus eliminating the allocation variables from the formulation. The reduced problem is then solved by a quasi-Newton method, and good results, but not always the best, are obtained with $N = 100$.

A constructive type heuristic is proposed in Moreno et al. (1991), in which a solution with N clusters, where N is selected between m and $2m$, is obtained initially. Then the surplus facilities are dropped in a "stingy" manner until exactly m are left. The term "stingy" is used to imply that in each drop we attempt to minimize the increase in the objective value. Comparable results to the Cooper method are obtained for problem sizes up to 900 customers and 10 facilities. An opposite constructive approach would be to add facilities one at a time in some greedy fashion (i.e., attempting to maximize the decrease in the objective value in each move) until m facilities are in place. See Brimberg et al. (2000) for a discussion on various strategies for the "drop" and "add" moves.

Bongartz et al. (1994) develop a projection method which solves, as in Murtagh and Niwattisyawong (1982), simultaneously for location and allocation variables. However, simple projection formulas on subspaces are derived (instead of solving the system of equations in general), and used to find descent directions. The authors compare a multi-start version of their method, where initial solutions are generated randomly or by partitioning customers in successive sets along a traveling salesman tour, with several existing methods, and obtain favorable results.

A different heuristic approach is to formulate and solve a discrete version of the problem. In our case, this given the well-known p-median problem where p (= m) refers to the number of facilities or median points to be located on a graph representation. The idea was first proposed by Cooper (1963). Hansen et al. (1998) construct the graph as n nodes representing both the fixed points and the potential facility sites. The edges are given lengths equal to the Euclidean distance between the connected points. Thus, the associated m-median problem may be formulated as follows:

$$(PM)\ \min \sum_{i=1}^{n}\sum_{j=1}^{n} w_j d_{ij} x_{ij}$$

subject to

$$\sum_{i=1}^{n} x_{ij} = 1,\ \forall\, j, \tag{3}$$

$$\sum_{i=1}^{n} y_i = m, \tag{4}$$

$$y_i - x_{ij} \ge 0,\ \forall\, i, j, \tag{5}$$

$$x_{ij} \ge 0,\ y_i \in \{0,\ 1\},\ \forall\, i, j; \tag{6}$$

where x_{ij} denotes the fraction of the demand (w_j) of customer (or node) j assigned to a facility at node i; $y_i = 1$, if a facility is opened at node i, and 0, otherwise; d_{ij} = the distance between nodes i and j. Constraint (3) ensures that the demand of each customer is satisfied, similarly as constraint (1) in (MWP). Constraint (4) imposes the condition that exactly m facility sites be opened, and finally, (5) ensures that $y_i = 1$ (or a facility is opened at node i) if any positive flow leaves node i. The discrete formulation takes advantage that the optimal solution of the continuous problem often has facilities located at or near the customer sites. Hansen et al. (1998) solve (PM) exactly using the efficient code of Hanjoul and Peeters (1985). A continuous improvement is then performed in a second step by solving m independent single facility problems resulting from the partition of the customer set found in the first step.

Mladenović and Brimberg (1995) develop a composite heuristic that combines the allocation neighborhood structure in Love and Juel (1982) with Cooper's method. A set number b of points are selected at random in the k-neighborhood of the current solution, obtained by reassigning k customers to different facilities, where b and k are parameters set by the analyst. Cooper's method is then applied to each of these b points. In a strictly descent version, a move is made if one of the local minima thus obtained is better than the current solution; otherwise, the process is repeated. This "fixed neighborhood" (or "iterated local") search may be considered a precursor of the soon to come "variable neighborhood" search.

In Brimberg et al. (2000), a new neighborhood structure is proposed that is based on relocation of facilities instead of reallocation of customers. This leads to a simple, yet very powerful local search procedure where the facilities are relocated one at the time to an unoccupied fixed point (i.e., a customer that does not have a facility coincident with it). The one-exchange neighborhood is constructed from all such possible single moves. This type of neighborhood has been used before in a discrete setting for network location problems (e.g., see Teitz and Bart (1968) for the p-median problem). A local search using Cooper's method is then conducted from all or selected points in the one-exchange neighborhood. An "interchange" version of the heuristic visits all the points in the neighborhood. In this case, an efficient updating procedure is adapted using ideas from Whitaker (1983) for the p-median problem to significantly reduce the computation time. Various strategies are also examined for visiting selected points in the neighborhood in a "drop and add" version of the heuristic. The new "interchange" and "drop and add" heuristics are found to perform very well compared to earlier methods, obtaining superior results in a fraction of the time.

Gamal and Salhi (2001) present constructive heuristics which produce good but slightly inferior results than the best ones from the literature. Initial solutions are found using a furthest distance criterion while avoiding having facilities inserted too close to already found locations. Gamal and Salhi (2003) also produce a two phase heuristic, known as a cellular-type method, where in phase one several local solutions are obtained from random starting points, and in the second phase, cells are constructed with their centers of gravity represented by the number of locations generated. A simple selection rule is adopted to choose the cells with a higher frequency. Some variations in the implementation of the rule are also attempted and encouraging results found.

3.2 Metaheuristics

Despite the number of different heuristics, the state-of-the-art did not advance far beyond the early random multi-start version of Cooper's algorithm until the application of metaheuristics to this problem. Some of the first attempts using metaheuristics may be attributed to Brimberg and Mladenovic (1996a, 1996b). The methodlogies proposed there all adopt the allocation interchange neighborhood structure of Love and Juel (1982). The method given in Brimberg and Mladenovic (1996a) uses basic tabu search rules (e.g., see Glover (1989) and Glover and Laguna (1997)). All $n(m-1)$ points in the one-exchange neighborhood of the current solution are examined, and a move is made to the best one. Meanwhile the reverse move is added to the Tabu list to prevent cycling. A FIFO rule is used to remove elements from the list once a specified length is exceeded. The search terminates when a stopping criterion, such as a maximum number of moves without improvement, is reached. Once a local minimum is attained (by steepest descent), the search will attempt to climb out of the trough (by mildest ascent). To illustrate, let us return to the small example in Figure 2, and the local optimum represented by the partition $\{\{1\}, \{3, 5\}, \{2, 4\}, \{6\}\}$. The next move is to the best point in the one-exchange neighborhood, given by $\{\{1, 6\}, \{3, 5\}, \{2, 4\}, \emptyset\}$, which of course is a worse solution. However, from here the procedure moves downhill to $\{\{1, 6\}, \{3, 5\}, \{2\}, \{4\}\}$, which is a better solution than the preceding local minimum, and as noted before, the global optimum.

Brimberg and Mladenovic (1996b) introduce a variable neighborhood search heuristic for solving MWP. The basic idea is to methodically increase the distance, or number of interchanges, k, defining the neighborhoods around a current solution, and conduct searches from random points in these neighborhoods until a better solution is found (see Mladenovic and Hansen (1997) and Hansen and Mladenovic (2001)). The number of points selected in a neighborhood is set to the parameter, b, as in the fixed neighborhood search. The local search conducted at these points is, once again, given by Cooper's method. The index k cycles repetitively through the sequence 1, ..., k_{max} (a second parameter), until an improved solution is found. Each time this occurs, k is re-set to 1. The iterations proceed until a stopping criterion (such as execution time, or limit on number of cycles (1, ..., k_{max}) without improvement) is reached. The strength of this approach appears to reside in the combination of random search superimposed on a systematic structure of neighborhoods over the solution space. For the purposes of this survey, it may also be interesting to note that this is the first paper to appear with variable neighborhood search (at the time referred to as a variable neighborhood algorithm); see Figure 3.

Houck et al. (1996) propose a genetic algorithm to solve MWP. The basic steps are outlined in Figure 4 (e.g., see Goldberg (1989)). A genetic algorithm is also implemented in Brimberg et al. (2000) using this general framework. The initial population is obtained by repeatedly running Cooper's algorithm from random starting points until N different local minima are found. The selection process picks solutions at random from the population using a nonuniform distribution that gives higher probabilities to the better solutions in accordance with a "survival-of-the-fittest" strategy. The cross-over operation selects facility sites at random from the two chosen solutions while maintaining a minimum separation distance between the sites. The new solution obtained from the cross-over is improved by applying the Cooper method to obtain a local minimum. Finally, when the population size exceeds a specified limit, the culling operator removes the worst solution from the list. The stopping criterion is given by a limit on the number of iterations or the execution time; alternatively, it may be based on an observed convergence of the population (i.e., the solutions in the population stop changing).

A significant advancement in the state-of-the-art occurs in Brimberg et al. (2000) with the location interchange neighborhood described previously. As noted, a highly effective local search results from the relocation of a single facility to an unoccupied customer site (or fixed point). All possible relocations of a single facility provides an interchange neighborhood of $O(nm)$ points. Various efficient drop/add strategies are also investigated which reduce the neighborhood size to $O(m+n)$ points.

The authors investigate several different heuristics that use interchange and drop/add local search within tabu and variable neighborhood search frameworks. In the latter case, a neighborhood structure is obtained by varying the number of facilities ($k = 1, ..., k_{max}$) to be displaced in the shaking operation. That is, a point is selected in the k-neighborhood by relocating k facilities in random fashion at unoccupied fixed points. Discretization strategies are also investigated where the related m-median problem is solved heuristically, and adjustment to continuous space is carried out at specified intervals. The aim is to make the best use of limited computer time. For example, the full interchange neighborhood is found to be

Step 1. { initialization. } Specify an initial solution, and conduct a local search using Cooper's method to obtain current solution (X_c). Set $k = 1$.

Step 2. { neighborhood search. }
- (a) Select b points at random in the k-neighborhood of X_c.
- (b) Conduct a local search at each of these points to obtain solutions X_i, $i = 1, ..., b$.
- (c) Retain the best solution X^* among them.
- (d) If X^* is better than X_c move there ($X_c \leftarrow X^*$), set $k = 1$, and return to the beginning of *Step* 2. Else continue to *Step* 3.

Step 3. { augmenting the neighborhood. }
- (a) If $k < k_{max}$, set $k = k + 1$ and return to *Step* 2.
- (b) Else if the stopping condition is not met, set $k = 1$ and return to *Step* 2
- (c) Else STOP (final solution is X_c).

Figure 3. A variable neighborhood search for MWP.

Step 1. Generate N different initial solutions.

Step 2. Sort the population in non-increasing order of solution quality (measured by the objective funtion).

Step 3. Repeat
- (a) select two solutions from the population;
- (b) mix these two solutions with a crossover operator to obtain a new solution;
- (c) Modify this solution with a mutation operator;
- (d) Insert the modified solution in the population and sort;
- (e) Remove solutions from the population with a culling operator;

Until a stopping criterion is reached.

Figure 4. A genetic algorithm for MWP.

best on the smaller problems tested, but the faster drop/add strategies take over on the larger problems.

Salhi and Gamal (2003) propose a genetic algorithm where three classes of chromosomes are constructed (e.g., good, mediocre and poor) and the selection process is based on the class rather than the individual entities. The assigned probabilities are systematically changed to reflect the convergence criteria within the genetic algorithm. In addition, a regular injection of new chromosomes is activated to provide diversity and avoid early convergence. Their method is able to obtain better results than the genetic algorithm in Brimberg et al. (2000). Meanwhile Taillard (2003) presents a decomposition heuristic where the number of facilities k to be merged in the subproblem is fixed beforehand. In each iteration, the subproblem is formed by choosing a central facility at random, and selecting ($k - 1$) closest facilities to it. More recently Brimberg et al. (2006) develop a variable neighborhood decomposition search (VNDS) heuristic where the size of the subproblem is allowed to vary in a systematic fashion. The subproblem is solved by VNS or reduced VNS depending on its size, giving rise to a two-level variable neighborhood scheme. It is interesting to note that among the different decomposition strategies examined, a mixed strategy for merging facilities outperformed pure random and deterministic strategies. The main outcome appears to be that VNDS is a competitive heuristic for solving large scale problems; see Hansen et al. (2001) and also Hansen and Mladenovic (2003) for a review. A summary of the references for the multisource Weber problem is given in Table 1.

4. THE CAPACITATED MULTISOURCE WEBER PROBLEM

In this section we present some research work for the case where the facilities have limited capacities within MWP. This capacitated continuous location-allocation problem is also known as the capacitated multisource Weber problem and seems to suffer from a shortage of published papers. This problem may be stated as follows: Given the location of each fixed point (customer), the demand at each fixed point, the transportation cost for the area of interest as a function of distance, the maximum number of facilities that can be opened and the capacity of each of these facilities, the aim is to determine the locations of the facilities to be opened, and the allocation of customers to each one.

We introduce the following additional input parameters and decision variables:

M : an upper bound on the number of facilities to be located,
b_i : the given capacity of the ith facility, $i = 1, ..., M$,
F_i : the given fixed cost of the ith facility, $i = 1, ..., M$,
Y_i = 1, if the ith facility is opened, and 0 otherwise, $i = 1, ..., M$.

A capacitated version of the problem may now be formulated as follows:

$$\text{(CMWP)} \min_{W,X,Y} \sum_{i=1}^{M} \sum_{j=1}^{n} w_{ij} \cdot \| X_i - A_j \| + \sum_{i=1}^{M} F_i Y_i$$

subject to

Table 1. Reference list for heuristics of the Multisource Weber problem

Heuristic	Type	References
Greedy	constr.	Brimberg et al. (2000), Gamal and Salhi (2001)
Stingy	constr.	Moreno et al. (1991)
Alternate	local *s.*	Cooper (1964), Scott (1970), Sullivan and Peters (1980), Baxter (1981), Bongartz et al. (1994)
Interchange	local *s.*	Love and Juel (1982), Brimberg et al. (2000)
Gradient	local *s.*	Murtagh and Niwattisyawong (1982), Chen (1983)
Fixed neigh.	composite	Mladenovic and Brimberg (1995)
p-median	discrete	Hansen et al. (1998)
Hybrids	constr.	Gamal and Salhi (2003)
Tabu search	Metah.	Brimberg and Mladenovic (1996a), Brimberg et al. (2000)
Variable *n.* search	Metah.	Brimberg and Mladenovic (1996b), Brimberg et al. (2000)
Genetic algorithms	Metah.	Houck et al. (1996), Brimberg et al. (2000), Salhi and Gamal (2003)
Hybrids	Metah.	Taillard (2003), Brimberg et al. (2006)

$$\sum_{j=1}^{n} w_{ij} \le b_i Y_i, \quad i = 1,\ldots,M \tag{7}$$

$$\sum_{i=1}^{M} w_{ij} = w_j, \quad j = 1,\ldots,n \tag{8}$$

$$w_{ij} \ge 0, \quad i = 1,\ldots,M, \quad j = 1,\ldots,n \tag{9}$$

$$Y_i \in \{0,1\}, \quad i = 1,...,M, \tag{10}$$

where the other symbols are the same as in the original version of MWP.

The objective function now represents the sum of the transportation costs and facility opening costs; (7) ensures that capacity constraints of the facilities are not violated; (8) and (9) are the same as (1) and (2) in the uncapacitated MWP; and (10) denotes the binary constraints on the Y_i, that is, a facility i is either open completely ($Y_i = 1$) at fixed cost F_i, or not at all. The above problem is a multi-modular mixed non-linear integer problem. In addition, if the problem is restricted to use exactly m facilities as in MWP, the following constraint needs to be added: $\sum_{i=1}^{M} Y_i = m.$ To guarantee feasibility, the m largest facilities must have a total capacity exceeding the total demand.

It should be noted that once the set of open facilities and their locations have been determined, the resulting problem reduces to the standard Transportation Problem (TP) which can be solved optimally in polynomial time. In short, the problem is therefore to determine the best facility configuration as in the uncapacitated MWP. Most research on capacitated facility location is on the discrete problem. The techniques used include a dual-ascent based method (Khumawala (1974)), cross decomposition method (Van Roy (1986)), constructive-type heuristics (Jacobsen (1983), Domschke and Drexl (1985)) and Lagrangian relaxation heuristics (Beasley (1993), Agar and Salhi (1998)).

Other related work on the continuous location problem include Brimberg et al. (2001) and Eben-Chaim et al. (2002) who studied the special case of capacitated facility location on a line. However, to our knowledge, the literature on the general problem given by CMWP is very scarce.

For instance, Cooper (1972) was the first to attempt this location problem using $F_i = 0$, and fixing the number of facilities to open. He presented exact and approximate methods for solving the resulting transportation-location problem. The heuristic method is known as the alternating transportation-location method (ATL). The idea is that firstly, m facility sites are randomly chosen from the fixed points. Then, the TP using these m open facilities is solved to find the allocation for the capacitated problem. (Recall in the uncapacitated case that the customers are simply allocated to the nearest facility. Due to capacity constraints this is no longer possible.) For each of the m independent sets of allocations just obtained, containing n_i fixed points, $i = 1, ..., m$, and where $\sum_{i=1}^{m} n_i \ge n$, since some customers may have their demand split between two or more facilities, the new facility location is found using the Weiszfeld iterative procedure as in MWP. The location problem and the TP are alternately solved until there is no improvement in cost. According to Cooper (1972), ATL yields a convergent monotone non-increasing sequence of values for the objective function. However, for similar reasons as in MWP, there is no guarantee that it will converge to the global minimum.

Gong et al. (1997) also studied this problem by designing a hybrid evolutionary method. Here, a genetic algorithm is used to find the locations of the facilities and a Lagrangian relaxation technique is adopted to solve the allocation subproblem. Their approach was tested on random data sets based on n = 20 to 180 customers, m = 2 to 6 facilities, and capacity b fixed to set values to cover demand feasibility. An improvement between 3% and 25% over the classical alternate method of Cooper is observed.

One simple way to solve the capacitated problem is to generate several local solutions for the uncapacitated case and choose the best configuration (i.e. configuration with the minimum transportation cost) out of all the runs to be the starting configuration for the capacitated problem. In other words, the capacitated problem is only solved once. As the best cost for the capacitated problem does not necessarily originate from the initial solution that yields the best cost for the uncapacitated problem, another way of solving this problem would be to consider each configuration for the uncapacitated problem. In this case, the capacitated problem needs to be solved several times. A compromise can be achieved by solving the capacitated problem only a few times using a smaller number of configurations, say k, extracted from all those found initial configurations. There are obviously schemes for selecting these k configurations. For instance, a rule that combines the total cost with other criteria could be used as a measure to differentiate between dissimilar configurations.

Zainuddin and Salhi (2007) developed a perturbation-based scheme which acts as a post optimization procedure. In this approach, the locations of the facilities found are perturbed by taking into account the clustering of the borderline customers. These are defined as the customers that lie in between their nearest facility and their second nearest facility. The size of the clusters as well as the number of customers in a given cluster is important as it has an effect on the choice of the new facilities. The customers of these clusters are temporarily assigned in entirety to their nearest facilities during the allocation stage (i.e., while the TP is solved). This task is performed by temporarily removing these customers from the system during the resolution of the TP stage and then re-introducing them back at the location stage. This restriction is imposed in order to bring the locations of their 'best' facilities nearer to these customers. When using these new locations in the next TP iteration, it is likely that some of these customers will be allocated to their nearest facility as in the uncapacitated case. This scheme is repeated starting with the recent best configuration for the capacitated problem until there is no reduction in cost or when there are no borderline customers that are being served by their second best facility. Though the TP is solved in polynomial time, the use of such a procedure so many times renders the whole exercise computationally unattractive. For instance, instead of applying the TP to the entire problem, at each iteration, a smaller part of the original problem is considered using a subset of facilities and their respective customers only. These are the facilities that are likely to be affected by the change. The construction of such a reduced neighborhood and its computational benefits can be found in the PhD thesis of Zainuddin (2004).

Luis et al. (2007) recently presented two constructive heuristics to generate the initial facility locations. These take into account the sparsity of the customers as well as the information gathered from previously found locations. Here, some guidance based on forbidden regions when selecting subsequent facility locations is followed. Encouraging preliminary results are obtained with a significant reduction in computational effort when compared to those originally found by the perturbation procedure discussed above.

5. CONCLUSIONS

This paper presents a review of one of the fundamental models in location theory known as the continuous location-allocation problem, also referred to as the multisource Weber problem. We attempt to introduce the unfamiliar reader to the rich history of this problem that spans several decades following the initial work by Cooper in the early 1960s. During this time period, the basic model and variants thereof have been the subject of considerable research. Numerous exact and approximate methods have been devised to solve the original uncapacitated problem, and subsequent capacitated versions of it. These methodologies cover a broad range of techniques from Operations Research. This survey will hopefully encourage students and practitioners of OR to study this exciting area in more depth.

The elegant method by Cooper of alternating between location and allocation phases remained the state-of-the-art for several years until the advent of metaheuristics. Whereas the performance of the older heuristics tends to deteriorate with problem size, due to an explosion in the number of local optima, the general frameworks provided by the metaheuristics allow a diversification of the search. We note in this overview the substantial improvement in solution quality obtained by methodologies such as Tabu search, genetic algorithm, and variable neighborhood search. Hybrid methods that include, for example, a decomposition step are able to improve even further the solution quality for large scale instances. Also much larger problem instances are being solved exactly due to the high quality initial solutions found by these heuristics.

For future research we believe the following avenues are worthwhile exploring. We classify them under two headings, namely, heuristic design and issue-related problems.

Heuristic Design—As for the problem being investigated here, we believe that the current trend in research is only at an infant stage, and many questions remain to be answered:

1) How far can these new metaheuristics be improved by experimenting with their parameter settings? Will some dominate the others? If so, why?
2) How is solution quality affected by problem size? Since we don't know the optimal solution in most cases, how can tight bounds be obtained to estimate the solution quality?
3) Can these new heuristics provide insights into the topology of the solution space for this class of problem? What is the effect of different neighborhood structures on this topology?
4) Can the information provided by these methods be used in a continuous self-improving process?

It follows that much work remains on the design side. The effect of different neighborhood structures on the search process as measured by the finite time performance of the various heuristics will have to be better understood. As for the measurement of solution quality, we suggest that dual-based heuristics should be investigated here. Recent work in this direction has been successful in guaranteeing solution quality of very large instances of the simple plant location problem (Hansen et al.(2007)).

Issues Related—For instance, for the capacitated case, instead of starting the Transportation Problem (TP) from scratch, the current location and allocation could be used as the initial basic feasible solution. Such a simple modification in the implementation of the TP will obviously reduce the computing time and hence may be worth the consideration. The application of a metaheuristic based on variable neighbourhood search, tabu search (TS) or genetic algorithm (GA), as successfully implemented in the uncapacitated problem, could be another way forward. The book of Glover and Laguna (1997) for TS and the one by Goldberg (1989) for GA provide useful platforms for further work. An overview of heuristic search in general can also be found in Salhi (2006). A number of variants such as the capacitated continuous location problem with an unknown number of facilities, or the inclusion of facility fixed costs that can be zone dependent or through-put related are worthwhile investigating. For instance, Brimberg and Salhi (2005) explore the case of facility fixed costs that are zone dependent, whereas Brimberg et al. (2004b) introduce a constant opening cost at any site. In both cases, the number of facilities is treated as an extra decision variable to determine the best tradeoff between fixed and variable costs. These ideas may also be adopted for other types of continuous location problems. For instance see Welch et al. (2006) and Wei et al. (2006) for some challenging related centre problems.

As a final note, we suggest that the relevance of the continuous location-allocation problem in the real world is increasing and will continue to do so. One of the main reasons is that much larger problem instances are being dealt with in practice. The advantages in terms of reduced data collection and data input for the continuous model are much more substantial as a result, and the trade-off between accuracy and modeling effort is shifting in favor of this approach. In fact, the required data is often already available from geographic information or positioning systems. Furthermore, as pointed out earlier, the "look up" advantage of discrete models where actual freight costs can be used becomes too onerous or infeasible for these larger networks. A second reason is that new applications are evolving where the continuous model is idealy suited. A case in point occurs in the area of data mining for instance. Undoubtedly, other applications will materialize in the future.

ACKNOWLEDGEMENT

We would like to thank all the referees for their constructive comments that improved both the content as well as the presentation of the paper.

REFERENCES

1. Agar, M.B. and Salhi, S. (1998). Lagrangean heuristics applied to a variety of large capacitated plant location problems. *Journal of the Operational Research Society*, 49: 1072-1084.
2. Baxter, J. (1981). Local optima avoidance in depot location. *Journal of the Operational Research Society*, 32: 815-819.
3. Beasley, J.E. (1993). Lagrangean heuristics for location problems. *European Journal of Operational Research*, 65: 383-399.
4. Bhaskaran, S. (1992). Identification of transshipment center locations. *European Journal of Operational Research*, 63: 141-150.
5. Bongartz, I., Calamai, P.H., and Conn, A.R. (1994). A projection method for l_p norm location-allocation problems. *Mathematical Programming*, 66: 283-312.
6. Brimberg, J., Hansen, P., and Mladenović, N. (2004a). Convergence of variable neighborhood search. Les cahiers du GERAD, G-2002-21, University of Montreal.
7. Brimberg, J., Hansen, P., and Mladenović, N. (2006). Decomposition strategies for large-scale continuous location-allocation problems. *IMA Journal of Management Mathematics*, 17: 307-316.
8. Brimberg, J., Hansen, P., Mladenović, N., and Taillard, É. (2000). Improvements and comparison of heuristics for solving the multisource Weber problem. *Operations Research*, 48: 444-460.
9. Brimberg, J., Korach, E., Eben-Chaim, M., and Mehrez, A. (2001). The capacitated p-facility location problem on the real line. *International Transactions in Operational Research*, 8: 727-738.
10. Brimberg, J. and Love, R.F. (1993). Global convergence of a generalized iterative procedure for the minisum location problem with l_p distances. *Operations Research*, 41: 1153-1163.
11. Brimberg, J. and Love, R.F., (1998). Solving a class of two-dimensional uncapacitated location-allocation problems by dynamic programming. *Operations Research*, 46: 702-709.
12. Brimberg, J. and Mladenović, N. (1996a). Solving the continuous location-allocation problem with Tabu search. *Studies in Locational Analysis*, 8: 23-32.
13. Brimberg, J. and Mladenović, N. (1996b). A variable neighborhood algorithm for solving the continuous location-allocation problem. *Studies in Locational Analysis*, 10: 1-12.
14. Brimberg, J., and Mladenović, N. (1999). Degeneracy in the multi-source Weber problem. *Mathematical Programming*, 85: 213-220.
15. Brimberg, J., Mladenović, N., and Salhi, S. (2004b). The multi-source Weber problem with constant opening cost. *Journal of the Operational Research Society*, 55: 640-646.
16. Brimberg, J. and Salhi, S. (2005). A continuous location-allocation problem with zone-dependent fixed cost. *Annals of Operations Research*, 136: 99-115.
17. Charalambous, C. and Bandler, J.W. (1976). Non linear optimization as a sequence of least pth optimization with finite values of p. *International Journal of System Science*, 7: 377-391.
18. Chen, R. (1983). Solution of minisum and minimax location-allocation problems with Euclidean distances. *Naval Research Logistics Quarterly*, 30: 449-459.
19. Chen, P.C., Hansen, P., Jaumard, B., and Tuy, H. (1998). Solution of the multisource Weber and conditional Weber problems by d-c. programming. *Operations*

Research, 46: 548-562.
20. Cooper, L. (1963). Location-allocation problems. *Operations Research*, 11: 331-343.
21. Cooper, L. (1964). Heuristic methods for location-allocation problems. *SIAM Review*, 6(1): 37-53.
22. Cooper, L. (1972). The transportation-location problem. *Operations Research*, 20: 94-108.
23. Das, P., Chakraborti, N.R., and Chaudhuri, P.K. (2001). Spherical minimax location problem. *Computational Optimization and Applications*, 18: 311-326.
24. Domschke, W. and Drexl, A. (1985). ADD-heuristics starting procedures for capacitated plant location models. *European Journal of Operational Research*, 21: 47-53.
25. Drezner, Z. (1984). The planar two-centre and two median problem. *Transportation Science*, 18: 351-361.
26. Drezner, Z. (1992). A note on the Weber location problem. *Annals of Operations Research*, 40: 153-161.
27. Drezner, Z. (Ed.) (1995). *Facility Location. A Survey of Applications and Methods*, Springer, New York.
28. Eben-Chaim, M., Mehrez, A., and Markovich, G. (2002). Capacitated location allocation problem on a line. *Computers and Operations Research*, 29: 459-470.
29. Eilon, S., Watson-Gandy, C.D.T., and Christofides, N. (1971). *Distribution Management*, Hafner, New York.
30. Fleischmann, B. and Paraschis, J. (1988). Solving a large scale districting problem: A case report. *Computers and Operations Research*, 15: 521-533.
31. Francis, R.L., McGinnis Jr., L.F., and White, J.A. (1992). F*acility Layout and Location: An Analytical Approach*, 2nd ed., International Series in Industrial and Systems Engineering, Prentice Hall, Englewood Cliffs, NJ.
32. Frenk, J.B.G., Melo, M.T., and Zhang, S. (1994). A Weiszfeld method for a generalized l_p distance minisum location model in continuous space. *Location Science*, 2: 111-127.
33. Gamal, M.D.H. and Salhi, S. (2001). Constructive heuristics for the uncapacitated location-allocation problem. *Journal of the Operational Research Society*, 52: 821-829.
34. Gamal, M.D.H. and Salhi, S. (2003). A cellular type heuristic for the multisource Weber problem. *Computers and Operations Research*, 30: 1609-1624.
35. Glover, F. (1989). Tabu Search-Part I. *ORSA Journal on Computing*, 1: 190-206.
36. Glover, F. and Laguna, M. (1997). *Tabu Search*, Kluwer Academic Publishers, Boston.
37. Gong, D., Gen, M., Yamazaki, G., and Xu, W. (1997). Hybrid evolutionary method for capacitated location-allocation problem. *Computers and Industrial Engineering*, 33: 577-580.
38. Goldberg, D.E. (Ed.) (1989). *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley, New York.
39. Hanjoul, P. and Peeters, D. (1985). A comparison of two dual-based procedures for solving the *p*-median problem. *European Journal of Operational Research*, 20: 387-396.
40. Hansen, P., Brimberg, J., Urosevic, D., and Mladenović, N. (2007). Primal-dual variable neighborhood search for the simple plant location problem. *INFORMS Journal on Computing*, 19: 552-564.
41. Hansen, P., Jaumard, B., and Krau, S. (1995). An algorithm for Weber's problem on the sphere. *Location Science*, 3: 217-237.
42. Hansen, P., Jaumard, B., Krau, S., and du Merle, O., (1997). A column generation algorithm for the multisource Weber problem. Les cahiers du GERAD, University of Montreal.
43. Hansen, P. and Mladenović, N. (2001). Variable neighborhood search: Principles and applications. *European Journal of Operational Research*, 130: 449-467.
44. Hansen, P. and Mladenović, N. (2003). Variable neighborhood search. In: F. Glover and GA Kochenberg (Eds.), *Handbook of Metaheuristics*, Kluwer Academic Publishers, Boston/Dordrecht/London, pp. 145-184.
45. Hansen, P., Mladenović, N., and Perez-Brito, D. (2001). Variable neighborhood decomposition search. *Journal of Heuristics*, 7(4), 335-350.
46. Hansen, P., Mladenović, N., and Taillard, E. (1998). Heuristic solution of the multisource Weber problem as a *p*-median problem. *Operations Research Letters*, 22: 55-62.
47. Hidaka, K. and Okano, H. (2003). An approximation algorithm for a large-scale facility location problem. *Algorithmica*, 35: 216-224.
48. Hodgson, M.J., Rosing, K.E., and Shmulevitz, F. (1993). A review of location-allocation applications literature. *Studies in Locational Analysis*, 5: 3-29.
49. Houck, C.R., Joines, J.A., and Kay, M.G. (1996). Comparison of genetic algorithms, random restart and two-opt switching for solving large location-allocation problems. *Computers and Operations Research*, 23: 587-596.
50. Jacobsen, S.K. (1983). Heuristics for the capacitated plant location model. *European Journal of Operational Research*, 12: 253-261.
51. Khumawala, B. (1974). An efficient heuristic procedure for the capacitated warehouse location problem. *Naval Research Logistics Quarterly*, 21: 609-623.
52. Krau, S. (1997). *Extensions du probleme de Weber*, Ph D. Thèse, École Polytechnique de Montréal.
53. Kuenne, R.E. and Soland, R.M. (1972). Exact and approximate solutions to the multisource Weber problem. *Mathematical Programming*, 3: 193-209.
54. Kuhn, H.M. (1973). A note on Fermat's problem. *Mathematical Programming*, 4: 98-107.
55. Love, R.F. (1976). One dimensional facility location-allocation using dynamic programming. *Management Science*, 22: 614-617.
56. Love, R.F. and Juel, H. (1982). Properties and solution methods for large location-allocation problems. *Journal of the Operational Research Society*, 33: 443-452.
57. Love, R.F. and Morris, J.G. (1975). A computational procedure for the exact solution of location-allocation problems with rectangular distances. *Naval Research Logistics Quarterly*, 22: 441-453.

58. Love, R.F., Morris, J.G., and Wesolowsky, G.O. (1988). *Facilities Location: Models and Methods*, North Holland, New York.
59. Luis, M., Salhi., S., and Nagy, G. (2007). The capacitated continuous location problem. *EURO Working Group on Locational Analysis XVI*, Estoril, Portugal.
60. Megiddo, M. and Supowit, K.J. (1984). On the complexity of some common geometric location problems. *SIAM Journal on Computing*, 13: 182-196.
61. Du Merle, O., Villeneuve, D., Desrosiers, J. and Hansen, P. (1997). Stabilisation dans le cadre de la generation des colonnes. Les cahiers du GERAD, University of Montreal.
62. Mladenovic, N. and Brimberg, J., (1995). A descent-ascent technique for solving the multisource Weber problem. *Yugoslav Journal of Operations Research*, 5: 211-219.
63. Mladenović, N., Brimberg, J., Hansen, P. and Moreno-Perez, J., (2007). The *p*-median problem: A survey of metaheuristic approaches. *European Journal of Operational Research*, 179: 927-939.
64. Mladenović, N. and Hansen, P. (1997). Variable neighborhood search. *Computers and Operations Research*, 24: 1097-1100.
65. Moreno-Perez, J., Rodrigez, C., and Jimenez, N. (1991). Heuristic cluster algorithm for multiple facility location-allocation problem. *RAIRO - Recherche Operationnelle/Operations Research*, 25: 97-107.
66. Murtagh, B.A. and Niwattisyawong, S.R. (1982). An efficient method for the multi-depot location-allocation problem. *Journal of the Operational Research Society*, 33: 629-634.
67. Ostresh Jr., L.M. (1973). MULTI-exact solutions to the M-center location-allocation problem. In: G. Rushton., M.F. Goodchild, and L.M. Ostresh Jr. (Eds.), *Computer Programs for Location-Allocation Problems*. Monograph Number 6, Department of Geography, University of Iowa, Iowa City, IA.
68. Ostresh Jr., L.M., (1975). An efficient algorithm for solving the two center location-allocation problem, *Journal of Regional Science*, 15: 209-216.
69. Patel, M.H. and Chidambaram, A. (2002). A new method for minimax location on a sphere. *International Journal of Industrial Engineering*, 9: 96-102.
70. Rosen, J.B. and Xue, J.L. (1991). Computational comparison of two algorithms for the Euclidean single facility location problem. *ORSA Journal on Computing*, 3: 207-212.
71. Rosing, K.E. (1992). An optimal method for solving the (generalized) multi-Weber problem. *European Journal of Operational Research*, 58: 414-426.
72. Salhi, S. (2006). Heuristic search: The science of tomorrow. In: S. Salhi (Ed.), *OR48 Keynote Papers*, ORS, Bath, pp 39-58.
73. Salhi, S. and Gamal, M.D.H. (2003). A GA-based heuristic for the multi-Weber problem. *Annals of Operations Research*, 123: 203-222.
74. Scott, A.J. (1970). Location-allocation systems: A review. *Geographical Analysis*, 2: 95-119.
75. Sullivan, P.J. and Peters, N. (1980). A flexible user oriented location-allocation algorithm. *Journal of Environmental Management*, 10: 181-193.
76. Taillard, É.D. (2003). Heuristic methods for large centroid clustering problems. *Journal of Heuristics*, 9: 51-73.
77. Teitz, M.B. and Bart, P. (1968). Heuristic methods for estimating the generalized vertex median of a weighted graph. *Operations Research*, 16: 955-961.
78. Van Roy, T.L. (1986). A cross decomposition algorithm for capacitated facility location. *Operations Research*, 12: 253-261.
79. Vine, M.F., Degnan, D., and Hanchette, C. (1997). Geographic information systems: Their use in environmental epidemiologic research. *Environmental Health Perspectives*, 105: 598-605.
80. Weiszfeld, E. (1937). Sur le point pour lequel la somme des distances de *n* points donnes est minimum. *Tohoku Mathematical Journal*, 43: 355-386.
81. Welch, S., Salhi, S., and Drezner, Z. (2006). The multifacility maximin planar location problem with facility interaction. *IMA Journal of Management Mathematics*, 17: 397-412.
82. Wei, H., Murray, A.T., and Xiao, N. (2006). Solving the continuous space *p*-centre problem: Planning application issues. *IMA Journal of Management Mathematics*, 17: 413-425.
83. Wesolowsky, G.O. (1982). Location problems on a sphere. *Regional Science and Urban Economics*, 12: 495-508.
84. Wesolowsky, G.O. (1993). The Weber problem: History and perspectives. *Location Science,* 1: 5-23.
85. Whitaker, R. (1983). A fast algorithm for the greedy-interchange for large-scale clustering and median location problems. *Information Systems and Operational Research*, 21: 95-108.
86. Zainuddin, Z.M. (2004). *Constructive and Tabu Search Heuristic for the Capacitated Continuous Location-Allocation Problem*, PhD Thesis, School of Mathematics and Statistics, University of Birmingham, UK.
87. Zainuddin, Z.M. and Salhi, S. (2007). A perturbation heuristic for the capacitated continuous location problem. *European Journal of Operational Research*, 179: 1194-1207.