An Alternating Heuristic for Medianoid and Centroid Problems in the Plane

By J. Bhadury, H.A. Eiselt, and J.H. Jaramillo

J. Bhadury, H.A. Eiselt, J.H. Jaramillo (2003). An alternating heuristic for medianoid and centroid problems in the plane. *Computers & Operations Research* 30, 553–565. DOI: 10.1016/S0305-0548(02)00024-2

Made available courtesy of Elsevier: http://www.elsevier.com/

*****Reprinted with permission. No further reproduction is authorized without written permission from Elsevier. This version of the document is not the version of record. Figures and/or pictures may be missing from this format of the document.*****

Abstract:
This paper develops two heuristics for solving the centroid problem on a plane with discrete demand points. The methods are based on the alternating step well known in location methods. Extensive computational testing with the heuristics reveals that they converge rapidly, giving good solutions to problems that are up to twice as large as those reported in the literature. The testing also provides some managerial insight into the problem and its solution.

Keywords: Competitive location model; Medianoid; Centroid; Stackelberg solution; Heuristic methods

Article:

Scope and purpose

When dealing with competitive location models, one popular solution concept is the Stackelberg solution. It assumes that one (group of) firm(s) acts as leader, while the other(s) act(s) as follower. In the locational context, the follower takes the locations of the leader as given and optimizes on that basis, whereas the leader will exercise foresight and take into account that a follower will subsequently locate additional facilities. It is commonly assumed that the leader knows how many facilities the follower will locate. In this bilevel programming problem, the leader's problem is called a centroid problem, whereas the follower faces a so-called medianoid problem. In both cases, the objective of the facility planner is to capture as much of the market as possible. Since centroid problems are inherently difficult, it is necessary to devise heuristic methods for all but the smallest models. This paper presents two such heuristics for the planar version of the centroid problem that are based on the repeated application of a medianoid solution, a much simpler problem. Computational results attesting to their performance are also included.

1. Introduction

Competitive location models were first introduced in the late 1920s by the economist Hotelling [1]. Following the taxonomy suggested by Eiselt et al. [2], the main components of these models are the space firms and customers locate in, the number of firms involved in the planning, the pricing policy followed by the competitors, the objective and solution concept employed, and, finally, the behavior of the customers. Even the few components listed here suggest that a large number of models can be devised. Many of the results in the literature suggest that competitive location models are very sensitive with respect to the assumptions, so that special care has to be taken to avoid extrapolating available results to similar, yet distinct, situations.

Among the solution concepts in the above taxonomy, one of the main ideas was put forward by the economist von Stackelberg [3]. It involves two planners or groups of planners, one called leaders, and the other followers. In the locational context, the leaders will locate in the first stage, followed by the followers who locate in the second stage. Often, one of the two groups has an advantage. For instance, in marketing, a "first mover advantage" refers to a situation that benefits the leader. In the context of competitive location, a pertinent reference is Ghosh and Buchanan [4]. If the leader has an advantage in a specific situation, then each firm would prefer to be a leader. However, in case of a leader advantage, a firm must be in a position to become a leader as far as its availability of resources is concerned. Once a firm has decided to become a leader, it will have to deal with a number of competitors who do not have the resources required to be a leader, and have resigned to act as followers. Given perfect information, a leader will be aware of the subsequent actions of a follower and consequently incorporate this information in his planning by solving an appropriate conditional location problem. On the other hand, a follower will take the actions of the leader(s) as a given and plan accordingly. The case of a follower (or second mover) advantage is much more problematic: in such a case, it does not benefit anybody to make the first move, and, if no incentive is introduced exogenously, none of the facilities will take any action.

The typical solution procedure for Stackelberg games is a backward recursion: first a reaction function is determined that indicates the optimal course of action of a follower, given all possible actions of the leader. The leader will then take this reaction function into account and determine his own optimal decision based on the follower's reaction. Given such a solution strategy, it is apparent that the leader's problem is considerably more difficult to solve than the follower's problem. The number of contributions in the literature regarding the two problems confirms this; see Eiselt and Laporte [5] for a comprehensive survey of sequential location models.

In the locational context, Hakimi [6] coined the phrases medianoid for the set of optimal locations of the followers and *centroid* for the set of optimal locations of the leader. He also provides a short taxonomy for these problems. In this scheme, an $(r|X_p)$ medianoid is defined as the problem of a follower who intends to locate r facilities of his own, given that the leader has located p facilities, where the set of leader's locations is denoted by X_p. Similarly, an $(r|\ p)$ centroid is defined as the problem of a leader who wants to locate p facilities, knowing that the follower will react and optimally locate r facilities of his own. As a medianoid problem is a difficult location problem in its own right and its solution determines the reaction function used by the leader to determine its centroid, it is obvious that in all reasonably realistic cases, no closed form solutions will exist for the centroid problem. For medianoid and centroid locations on networks, Hakimi demonstrated that the $(r|X_1)$ medianoid on general graphs is NP-hard, and so is the $(1|\ p)$ centroid. Furthermore, some nice properties that hold in standard location models such as the node (or Hakimi) property are lost even on tree networks. A few heuristic methods have been suggested, most prominently by Benati and Laporte [7] and Serra and ReVelle [8]. However, the largest reported data set to which the suggested algorithms have been applied is the famous 55-node problem by Swain [9] used in [8], and the largest values of p and r reported on are nine.

For problems in the Euclidean plane, even fewer results are available in the literature. Drezner [10] described an exact algorithm for the $(1|X_p)$ medianoid problem that runs in $O(n^2 \log n)$ time, where n denotes the number of customers in the problem. For problems with $r = 2$ and 3 and arbitrary values of p, Infante-Macias and Muñoz-Perez [11] provide efficient algorithms for medianoid problems, given that rectilinear distances are employed. However, not much is known about $(r|X_p)$ medianoids and $(r|p)$ centroids in the plane with general values of r and p; a gap that this paper attempts to fill.

The remainder of this paper is organized as follows. The next section introduces the model, Section 3 discusses the heuristic models investigated in the paper, Section 4 describes the design of the series of computational tests, its results and implications, and the last section summarizes the contribution of the paper, provides an outlook and suggests future research.

2. The model

Consider a two-dimensional plane in which n customers are located at points n_i, $i = 1,\ldots,n$, so that $N = \bigcup_{i=1}^{n} n_i$. The demand of a customer at n_i is assumed to be fixed at w_i, known as its "weight". The structure of a fixed demand suggests that the homogeneous good under consideration is essential. The facilities all employ mill pricing, and prices are agreed upon or legislated, i.e., fixed and equal. Given this pricing structure, customers, who are assumed to have perfect information and behave rationally in the sense of minimizing their cost, will patronize the closest source. This assumption is justified in case of truly homogeneous goods as shown by Eiselt and Laporte [12].

Suppose now that two firms locate facilities that can be thought of as individual stores or facilities. This is done in a sequential fashion, with the location leader locating p facilities, and the location follower locating r facilities. Given the customer behavior discussed above, a facility will capture a customer if and only if it is closer to that customer than any other facility. As Hakimi [13] suggested, ties are broken in favor of the location leader, an assumption justified by customers who show loyalty to the incumbent facility.

Suppose now that the distances are measured according to the Euclidean metric and let $d(x, y)$ denote the distance between two points x and y. Also, given a set Z of points, the distance between some point x and the set Z is the shortest distance between x and any point in the set Z, i.e., $d(x, Z) = min_{z\in Z}\ \{d(x, z)\}$. Let now $X = \bigcup_{i=1}^{p} x_i$ denote a set of given points, sometimes referred to as "seeds". The *Voronoi set* $V(x_i)$ associated with the seed x_i is then the set of points $V(x_i) = \{z: d(x_i, z) \leq d(x_k, z)\ \forall x_k \neq x_i\}$. The *Voronoi diagram* is then the collection of Voronoi sets $V(X) = \bigcup_{i=1}^{p} V(x_i)$; see O'Rourke [14] for a more detailed coverage of Voronoi diagrams. If the seeds represent facilities, and customers patronize the closest facility to them, then $V(x_i)$ includes all customers that patronize a store at x_i.

The definition of Voronoi diagrams allows us now to formalize the concept of medianoids and centroids. Following convention, the location leader will locate p facilities at points $X_p = \{x_1, x_2, \dots, x_p\}$, and the follower will locate r facilities at $(Y_r|X_p) = \{y_1, y_2, \dots, y_r\}$, the conditional notation indicating that the follower takes the locations of the leader's facilities at X_p into account.

Given now the sets X_p and $(Y_r|X_p)$, the follower will *capture*, a term coined by ReVelle [15], all customers that are closer to any of his own facilities than to the closest of his competitor's facilities. Formally, the follower captures all customers in the set

$$V(Y_r|X_p) = \left\{n_i: d\left(n_1, (Y_r|X_p)\right) < d(n_i, X_p)\right\},$$

so that the demand or weight captured by the follower is

$$W(Y_r|X_p) = \sum_{n_i \in v(Y_r|X_p)} w_i\ .$$

The follower will now maximize his capture, i.e., find a set of locations $(Y_r^*|X_p)$, so that

$$W(Y_r^*|X_p) \geq W(Y_r|X_p)$$

for all feasible sets of locations $(Y_r|X_p)$. This describes the medianoid.

In addition to being the medianoid, the function $(Y_r^*|X_p)$ is also the reaction function taken into consideration by the location leader. Given this, the leader will determine a set X_p^*, such that

$$W(Y_r^*|X_p^*) \leq W(Y_r^*|X_p)$$

for all feasible sets of locations X_p, where X_p is the solution to the leader's $(r|p)$ centroid problem.

3. Heuristic methods for medianoid and centroid problems

This section will develop heuristic methods to solve medianoid and centroid problems in the plane. We first describe two approaches to the medianoid problem, which are subsequently applied to solve the centroid problem as well.

3.1. A greedy heuristic

The first heuristic to solve medianoid problems is based on Drezner's [10] exact method for $(1|X_p)$ medianoids. In order to develop our heuristic, we first describe Drezner's method. It begins by finding a Voronoi diagram $V(X_p)$ for the existing facilities at X_p. For each customer location at $n_i \in V(x_k)$, a disk D_i is determined, so that n_i is at the center of D_i and x_k is located on the circumference of D_i. Define now I_ℓ as the resulting intersection of any set of two or more such disks and define S_ℓ as the set of customers whose disks generate the intersection I_ℓ. It is then apparent that any new facility located in the interior of I_ℓ is closer to all customers in S_ℓ , so that the new facility will capture all of them. Now define $w(S_\ell) = \sum_{i:n_i \in S_\ell} w_i$ as the demand captured by a new facility if it were to locate in I_ℓ. The intersections can now be ordered, so that $w(I_1) \geq w(I_2) \geq \cdots$, where ties are broken arbitrarily. The 1-medianoid problem is then solved by locating the new facility anywhere in I_1. Drezner shows that this task can be accomplished in O(n log n) time.

An example may illustrate the idea behind the intersections of the disks. In Fig. 1, four customers n_1, n_2, n_3, and n_4 have demands of $w_1 = 7, w_2 = 5, w_3 = 3$, and $w_4 = 1$. The intersections $I_1, \ldots, I_{10}$ are shown in the figure, where $S_1 = \{n_1, n_2, n_3, n_4\}$ with $w(S_1) = 16$, $S_2 = \{n_1, n_2, n_3\}$ with $w(S_2) = 15$, $S_3 = \{n_1, n_2, n_4\}$ with $w(S_3) = 13$, $S_4 = \{n_1, n_2\}$ with $w(S_4) = 12$, $S_5 = \{n_1, n_4\}$ with $w(S_5) = 8$, $S_6 = \{n_2, n_3\}$ with $w(S_6) = 8$, $S_7 = \{n_1\}$ with $w(S_7) = 7$, $S_8 = \{n_2\}$ with $w(S_8) = 5$, $S_9 = \{n_3\}$ with $w(S_9) = 3$, and $S_{10} = \{n_4\}$ with $w(S_{10}) = 1$. The maximum number of such intersections is 2^n.

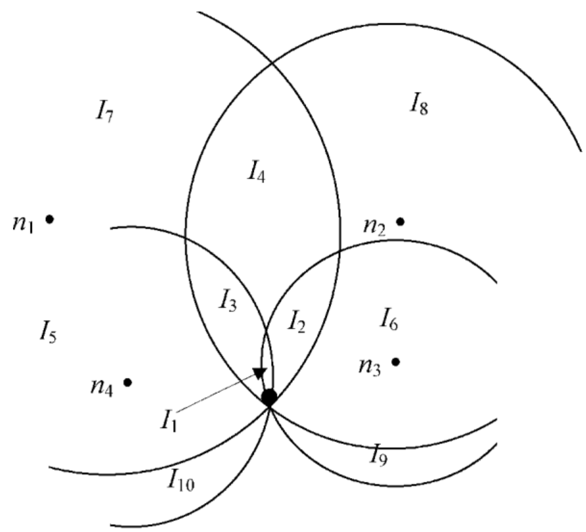

Fig. 1.

As the problem is NP-hard and the number of intersections is exponential, we employ a greedy technique to approximately solve the $(R|X_p)$ medianoid problem; this greedy heuristic is simply a sequential repetition of Drezner's heuristic. In particular, we locate a new facility in the intersection with the largest weight, delete all customers that are covered by the new facility, and repeat the process until all r new facilities are located. The procedure can formally be described as follows:

Greedy heuristic
Let S:=N = $\bigcup_{i=1}^{n}\{n_i\}$.
For q = 1–r do
  If S = ∅,
    locate the qth facility anywhere,
else
    compute the sets I_1;I_2,… for S,
    locate the qth facility anywhere in I_1, and
    set S:=$S \setminus \bigcup_{n_i \in I_1} n_i$.

Given that this greedy heuristic applies Drezner's $O(n^2 \log n)$ method r times, its complexity is $O(rn^2 \log n)$.

3.2. Minimum differentiation heuristic

Our second heuristic method for the $(r|X_p)$ medianoid problem is based on Hotelling's observation that at equilibrium, the facilities of duopolists tend to cluster together at some central point of the

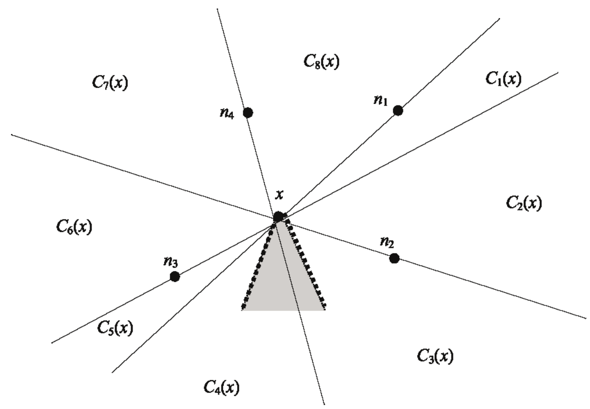

Fig. 2.

market. While d'Aspremont et al. [16] have shown that this observation does not hold in the case of variable prices, it does hold in case of fixed and equal prices. Moreover, the "pairing" of multiple facilities was confirmed by Eaton and Lipsey [17]. Observations in practice reveal that some classes of facilities, e.g., fast food chains, exhibit a strong tendency to cluster. Such background evidence constitutes the rationale of our *Minimum Differentiation Heuristic*. Its idea is to position a new facility y at an arbitrarily small distance $\varepsilon > 0$ away from an existing facility. The capture of the new facility can then be determined as follows. Construct a line that is equidistant to the existing and the new facility, and let H denote the halfplane given by that line, which does not include x. The new facility y will then capture all customers in $H \cap V(x)$. Next, we have to determine in which direction from the existing facility at x the new facility should be located.

For this purpose, consider the following procedure. Draw a straight line through each pair (x, n_i); $n_i \in V(x)$ for some existing facility x. This generates $2|V(x)|$ cones $C_v(x)$, $v = 1,\ldots,2|V(x)|$. Starting anywhere, number these cones in, say, a clockwise direction. Each cone $C_v(x)$, $v = 1,\ldots, |V(x)|$, has an *opposing cone* $C_{v+|V(x)|}(x)$. Similarly, the opposing cone for $C_v(x)$, $v = |V(x)| + 1,\ldots,2|V(x)|$ is $C_{v-|V(x)|}(x)$. Any hyperplane that contains the existing facility x and some point in $C_v(x)$ will divide the sets of customers $V(x)$ into two subsets $V_v^+(x)$ and $V_v^-(x)$ with $V_v^+(x) \cup V_v^-(x) = V(x)$. Define now the weight of a set $V_v^+(x)$ as $w(V_v^+(x)) = \sum_{i:n_i \in V_v^+(x)} w_i$ and similar for $V_v^-(x)$. Furthermore, let $V^*(x)$ be a set for which $w(V^*(x)) = \max_v \{w(V_v^+(x)); w(V_v^-(x))\}$. Let $C^*(x)$ now be a cone that is generated by the union of all rays rooted at x that are perpendicular to hyperplanes through x and any point in $V^*(x)$, so that a new facility in the open set $C^*(x)$ captures $V^*(x)$. Note that $C^*(x)$ is not one of the $C_v(x)$ cones.

As an illustration, consider Fig. 2, which includes an existing facility x and four customers n_1, n_2, n_3, and n_4 whose demands are w_1=1, w_2=5, w_3=7, and w_4=3. The cones are $C_1(x),\ldots,C_8(x)$ and the sets of customers captured are $V_1^+(x)=\{n_2\}$ and $V_1^-(x)$ $\{n_3,n_4,n_1\}$ with weights $w(V_1^+(x))=5$ and $w(V_1^-(x)) = 11$ for $C_1(x)$ and $C_5(x)$; $V_2^+(x) = \{n_2, n_3\}$ and $V_2^-(x) = \{n_4, n_1\}$ with weights $w(V_2^+(x)) = 12$ and $w(V_2^-(x)) = 4$ for $C_2(x)$ and $C_6(x)$; $V_3^+(x) = \{n_3\}$ and $V_3^-(x) = \{n_4, n_1, n_2\}$ with weights $w(V_3^+(x)) = 7$ and $w(V_3^-(x)) = 9$ for $C_3(x)$ and $C_7(x)$; and $V_4^+(x) = \{n_3,n_4\}$ and $V_4^-(x) = \{n_1,n_2\}$ with weights $w(V_4^+(x)) = 10$ and $w(V_4^-(x)) = 6$ for $C_4(x)$ and $C_8(x)$. The largest weight is achieved for the set $V^*(x) = \{n_2, n_3\}$, so that the optimal medianoid location is at a distance of ε from x in the cone $C^*(x)$ shown as the shaded area and bordered by the broken lines.

Similar to the greedy heuristic described above, the minimum differentiation heuristic first determines the potential captures of a new facility next to each of the existing facilities. It then ranks the potential market captures from largest to smallest, allocates the first new facility to the location that allows the maximal capture, deletes all customers captured in the process, and repeats the procedure until all new facilities have been located. Barring degenerate cases in which two customers are on the same straight line as an existing facility, it can be demonstrated that given p existing facilities, no more than $2p$ new facilities are required to ensure that the entire demand is captured by the follower. This procedure can be algorithmically described as follows:

Minimum differentiation heuristic
Let $S := N = \bigcup_{i=1}^{n}\{n_i\}$.
For $q = 1$–r do
    If S = ∅,
        locate the qth facility anywhere,
    else
        compute the sets $V^*(x_k)$ and their values $w(V^*(x_k))$ for $k = 1,\ldots,p$, and
        reorder them, so that $w(V^*(x_1)) \geq w(V^*(x_2)) \geq \cdots$,
        locate the qth facility anywhere in $C^*(x_1)$, and
        set $S := S \setminus \bigcup_{n_i \in V^*(x_1)} n_i$.

Having developed two heuristic algorithms for medianoid problems in the first part of this section, we are now able to develop a procedure that generates solutions for the centroid problem. The basic idea is as follows. Given the set of locations of the leader X_p, the $(r|X_p)$ medianoid is to optimally locate the r facilities of the follower given the leader's locations. Once that is done, the follower's facilities are located at Y_r and now the original leader may tentatively assume the role of the follower and reoptimize the set of its facilities by solving a $(p|Y_r)$ problem. This process is then repeated until some stop criterion is satisfied. In other words, the two players alternately solve medianoid problems. For additional results on the convergence of similar alternating procedures (for equilibrium problems), readers are referred to Sherali and Soyster [18].

This process is based on some results in the literature. A repeated alternate optimization process may lead to a Nash equilibrium, if one exists. Consider, for example, a line segment along which customers are uniformly distributed and on which two decision makers attempt to locate a single facility each. This is a simplified version of Hotelling's original model with fixed and equal costs. Given an arbitrary location of one facility, its opponent will solve the medianoid problem that is a point next to the existing facility, arbitrarily close, on the "longer" side, i.e., the side with the higher demand. If now the other decision maker relocates, he will do so again arbitrarily close to its opponent on its longer side. This process will converge to a locational pattern in which both facilities are located at or near the center of the line segment; this is Hotelling's original "minimum differentiation". This locational pattern is an equilibrium, and it is easy to demonstrate that it is also a pattern in which one location is the centroid, and the other the medianoid. The most ambitious use of this principle was made by Okabe and Suzuki [19] who repeatedly reoptimized one of (up to) 256 facilities in a unit square in order to determine the resulting locational pattern.

A similar result is known about trees. Slater [20] and later Hakimi [6] demonstrated that the (1|1) centroid in a tree is located at a median of the given tree. Bhadury and Eiselt [21] have shown that two competing facilities who engage in repeated optimization will—provided that an appropriate tie-breaking rule is employed—converge to an equilibrium that exists at a median of the tree and an adjacent node that is located in the largest subtree spanned by the median (assuming that co-location is prohibited). Given Slater's result, this is also a centroid. Whether or not such repeated optimization always converges to an equilibrium is not known. Furthermore, whether or not an equilibrium —provided one exists—always coincides with the locations of a centroid and a medianoid is also unknown.

The computational complexity of the minimum differentiation heuristic can be established as follows. Computing the Voronoi diagram in the first step of the "else" loop is a task that can be accomplished in O(p log p) time using techniques from computational geometry; for details see, e.g., O'Rourke [14]. Computing $w(V^*(x_k))$ will require $O(n^2)$ time, and the sorting in the next step requires O(p log p) time. As a result, the time complexity of each of the outer "for" loops is $O(n^2)$. Since the outer loop itself has to be executed r times, this leads to a total time complexity of $O(n^2r)$ for the minimum differentiation heuristic.

4. Computational results

Each of the two heuristic methods discussed in the previous section, the greedy algorithm and the minimum differentiation heuristic, were coded in FORTRAN 90. Each of codes is about 700 lines in length. The computational equipment was an IBM PC with a Pentium II processor and 64 megabytes of RAM.

The test problems were then generated as follows. Customers were randomly generated, following a uniform distribution on a 50 x 50 grid. The demand w_i of customer n_i was randomly generated, following a uniform distribution on the integers between 1 and 200. Throughout the test series, we have assumed that $r = p$, i.e., the leader and the follower have an equal number of facilities to locate. The reason for this somewhat restrictive assumption is that with it, leader and follower have the same initial "strength" in terms of the number of facilities. It then becomes possible to draw conclusions regarding the leader and follower positions. A total of 500 different problems were generated randomly.

As expected for the medium-sized problems in our test series, both heuristic methods required only moderate amounts of computing time. An example is provided in Fig. 3, where the number of customers varies from 40 to 100 with a fixed number of p = 20 facilities to be located by each competitor. Here, the time required by the minimum differentiation heuristic is almost constant at less than 10 s, whereas the time needed by the greedy heuristic to solve the problem increases linearly with the number of customers. All problems were solved within 5 min.

The computation times with the number of customers fixed at n =100 and the number of facilities as variable are similar. See, for example, Fig. 4 which shows the results for n = 100. As shown in the figure, the time required by the minimum differentiation heuristic is nearly constant, whereas the computation time of the greedy heuristic is quite high for a small number, e.g., 5, of facilities, and then drops off substantially as p increases. An explanation for this behavior suggests itself. As the minimum differentiation heuristic places competing facilities adjacent to existing facilities, its time requirement is likely to be small and dependent on p. On the other hand, the greedy heuristic must compute intersections I_ℓ and then locate facilities in them. As the number of facilities increases, the distances between customers and their closest facility decrease, so that the number of intersections I_ℓ decreases as well, hence requiring less computational effort.

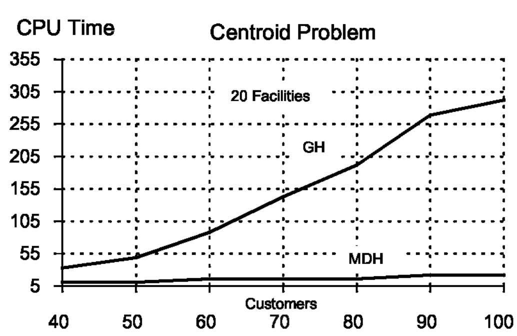

Fig. 3. CPU time in seconds for the two heuristic methods.

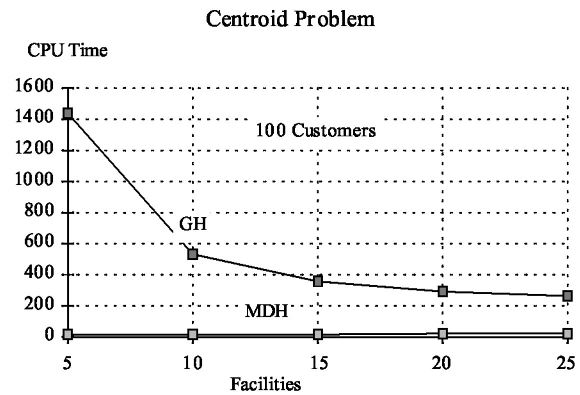

Fig. 4. CPU time in seconds.

As far as convergence is concerned, both heuristics exhibit a similar behavior that is shared by many other exact and heuristic methods: initial improvements are great, but level off after about 40 iterations, after which only marginal improvements are obtained before convergence. This behavior is shown in Fig. 5, where the number of iterations are plotted on the abscissa, whereas the ordinate indicates the market share lost when the opponent of the planning facility solves its medianoid problem and relocates.

The robustness of the methods with respect to different initial solutions was tested on a small set of ten problems and the results are summarized in Fig. 6. It turned out that while most solutions found by different initial solutions also differed, the market captures did not. Typically, the differences between best and worst solutions was within a few percentage points of the total market.

As far as the quality of the solutions is concerned, the test series reveals that the solutions found by the minimum differentiation heuristic are consistently better by about 6 percentage points as opposed to the solutions determined by the greedy method in which the planning facility loses an average of 64% of its market share when its opponent relocates to the medianoid solution. However, it may be suspected that the solutions are sensitive with respect to the number of facilities to be located. More specifically, if $r \neq p$, the minimum differentiation heuristic may no longer perform as well as it does in this series, particularly if $r > p$. This is an issue that would be worthwhile addressing in future research.

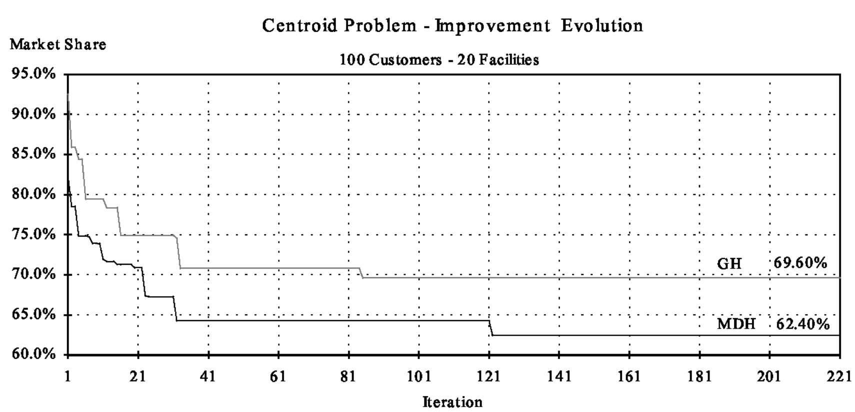

Fig. 5. Improvement evolution of the heuristics.

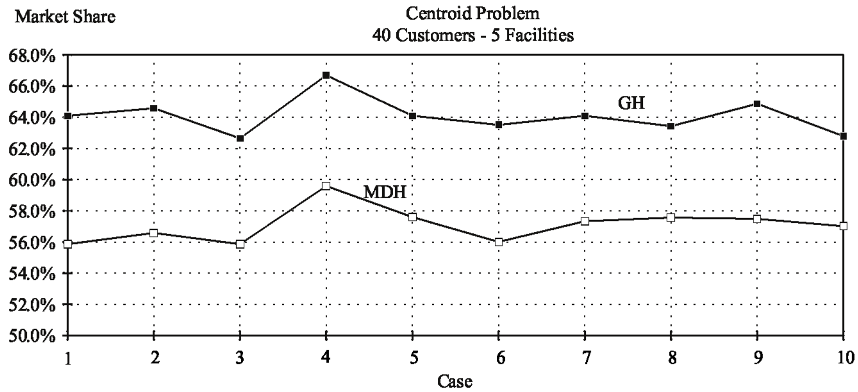

Fig. 6. Sensitivity of the heuristic solutions to initial locations.

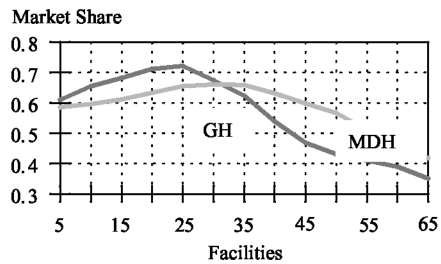

Fig. 7. Market share as a function of $p = r$.

The solutions in the test series also reveal another interesting phenomenon shown in Fig. 7. As far as the best known solutions are concerned, if the values of *r* and *p* are small, then the leader ends up with about 40% of the total market. As the number of facilities is increased, the leader's capture decreases at first, until it begins to increase again.

The reason is that when the number of facilities exceeds 40–45% of the number of customer sites, it is optimal for the leader to locate some of his facilities on the customer sites themselves. Once this is done, the follower cannot capture this site anymore even if he would locate directly at it, as by assumption, ties are broken in favor of the leader. As a result, once the number of facilities exceeds 40–45% of the number of customers, the follower ends up with less than 50% of the market.

5. Summary and conclusions

In this paper, we have designed and tested two heuristic algorithms for the centroid location problem. The basic idea was to use a method designed for medianoids, and repeatedly apply it, alternately designating each of the two competitors as leader and follower. The results of the computational testing revealed fast convergence of both heuristics with the minimum differentiation heuristic having the edge. Also, the minimum differentiation heuristic appears to be the more robust of the two. It also consistently outperforms the greedy heuristic. As in all competitive location problems, the solutions appear to be very sensitive with respect to the assumptions of the model and the relocation process. Changing, for instance, the tie-breaking rule could be expected to have a profound effect on the solutions.

Future research could go into a variety of directions. The most obvious direction points to the investigation of other solution methods. Such techniques could be based on metaheuristics such as tabu search (suitably modified for centroid problems), or adaptations of methods developed for *p*-median problems, such as Rosing's [22] heuristic concentration.

Other possible directions include the investigation of different allocation rules, i.e., different principles according to which customers are assigned to (or freely choose) facilities. Whether or not different allocation or choice functions change the solution or its objective (as Serra et al. [23] demonstrate they do not for medianoid problems) is an open question.

Other challenging problems relate to the question whether or not repeated re-optimization leads to an equilibrium. A pertinent result was established for a competitive location model on trees by Bhadury and Eiselt in [20] but it is not known if their result carries over to more general problems. Similarly, it is still an open problem if an equilibrium, provided it exists, coincides with the locations of centroids and medianoids.

Acknowledgements
J. Bhadury was supported in part by a grant from the Office of Research and Sponsored Programs at California State University-Hayward. H.A. Eiselt was supported in part by a grant from the Natural Sciences and Engineering Research Council of Canada under Grant No. OGP00009160. All support is gratefully acknowledged. We thank the anonymous referees for their comments that helped to clarify a number of issues.

References:

1. Hotelling H. Stability in competition. Economic Journal 1929;39:41–57.
2. Eiselt HA, Laporte G, Thisse J-F. Competitive location models: a framework and bibliography. Transportation Science 1993;27:44–54.
3. Stackelberg H von. Grundlagen der theoretischen Volkswirtschaftslehre, 1943. Translated by Peacock AT. The Theory of the Market Economy. London-Edinburgh-Glasgow: W. Hodge, 1952.
4. Ghosh A, Buchanan B. Multiple outlets in a duopoly: a first entry paradox. Geographical Analysis 1988;20:111–21.
5. Eiselt HA, Laporte G. Review of sequential location problems. European Journal of Operational Research 1996;96:217–31.
6. Hakimi SL. On locating new facilities in a competitive environment. European Journal of Operational Research 1983;12:29–35.
7. Benati S, Laporte G. Tabu search for the (r|Xp)-medianoid and (r|p)-centroid problems. Location Science 1994;2:193–204.
8. Serra D, ReVelle C. Market capture by two competitors: the preemptive location problem. Journal of Regional Science 1994;34:549–61.
9. Swain R. A parametric decomposition algorithm for the solution of uncapacitated location problems. Management Science 1974;21:189–98.
10. Drezner Z. Competitive location strategies for two facilities. Regional Science and Urban Economics 1982;12: 485–93.
11. Infante-Macias R, Mu˜noz-Perez J. Competitive location with rectilinear distances. European Journal of Operational Research 1995;80:77–85.
12. Eiselt HA, Laporte G. Demand allocation functions. Location Science 1998;6:175–87.
13. Hakimi SL. Locations with spatial interactions: competitive locations and games. In: Mirchandani PB, Francis RL, editors. Discrete location theory. New York: Wiley-Interscience, 1990. p. 439–78.
14. O'Rourke J. Computational geometry in C. Cambridge, UK: Cambridge University Press, 1998.
15. ReVelle C. The maximum capture or 'sphere of influence' location problem: Hotelling revisited on a network. Journal of Regional Science 1986;26:343–58.
16. d'Aspremont J, Gabszewicz JJ, Thisse J-F. On Hotelling's 'stability in competition'. Econometrica 1979;49:1145–50.
17. Eaton BC, Lipsey RG. The principle of minimum differentiation reconsidered: some new developments in the theory of spatial competition. Review of Economic Studies 1975;42:27–49.
18. Sherali HD, Soyster AL. Convergence analysis and algorithmic implications of two dynamic processes toward an oligopoly—competitive fringe equilibrium set. Computers & Operations Research 1988;15:69–81.

19. Okabe A, Suzuki A. Stability of spatial competition for a large number of firms on a bounded two-dimensional space. Environment and Planning 1987;A19:1067–82.
20. Slater P. Maximin facility location. Journal of Research of the National Bureau of Standards. Section B: Mathematical Sciences 1975;79:107–15.
21. Bhadury J, Eiselt HA. Reachability of locational Nash equilibria. Operations Research Spectrum 1998;20:101–7.
22. Rosing KE. Heuristic concentration: a study of stage one. Environment and Planning B 2000;27:137–50.
23. Serra D, Eiselt HA, Laporte G, ReVelle CS. Market capture models under various customer choice rules. Environment and Planning B 1999;26:741–50.

J. Bhadury (Ph.D., University of Texas at Dallas, 1991) is an Associate Professor of Management at CSU-Hayward. He has taught previously at SUNY-Buffalo and University of New Brunswick, Canada. His primary area of research is logistics (especially location theory) and management science. His publications have appeared in INFORMS Journal of Computing, the Journal of Regional Science, the European Journal of Operational Research, Location Science=Computers & Operations Research, OMEGA, and Social Choice and Welfare. Currently, he also serves as the Chair of the Section on Location Analysis of INFORMS.

H.A. Eiselt received his doctorate from the University of G   ottingen, Germany. He has taught at Concordia University in Montreal, the Chinese University of Hong Kong, Bilkent University in Ankara, and he is currently employed at the University of New Brunswick in Fredericton, Canada. His papers have appeared in Mathematical Programming, Operations Research, Networks, the Journal of Regional Science, the European Journal of Operational Research, Computers & Operations Research, the Journal of the Operational Research Society, and many others. He was formerly the editor of Location Science, and presently serves on a number of editorial boards.

Jorge Jaramillo works as an engineer in Tampa Airlines based in Medellin, Colombia. He obtained his Master's degree from the Department of Industrial Engineering in SUNY-Buffalo where this work was initiated. His areas of interests include development of heuristics for location theory problems and he has published in Location Science= Computers & Operations Research.