On the Continuous Weber and k-Median Problems

(Extended Abstract)

Sándor P. Fekete* Joseph S. B. Mitchell† Karin Weinbrecht‡

Abstract

We give the first *exact* algorithmic study of facility location problems that deal with finding a median for a *continuum* of demand points. In particular, we consider versions of the "continuous k-median (Weber) problem" where the goal is to select one or more center points that minimize the average distance to a set of points in a demand *region*. In such problems, the average is computed as an integral over the relevant region, versus the usual discrete sum of distances. The resulting facility location problems are inherently geometric, requiring analysis techniques of computational geometry. We provide polynomial-time algorithms for various versions of the L_1 1-median (Weber) problem. We also consider the multiple-center version of the L_1 k-median problem, which we prove is NP-hard for large k.

1 Introduction

> *"There are three important factors that determine the value of real estate – location, location, and location."*

There has been considerable study of facility location problems in the field of combinatorial optimization. In general, the input to these problems includes a weighted set D of demand locations (with weight distribution δ and total weight A), a set F of feasible facility locations, and a distance function d that measures cost between a pair of locations. In one important class of questions, the problem is to determine one or more feasible *median* locations $c \in F$ in order to minimize the average cost from the demand locations, $p \in D$, to the corresponding central points c_p that are nearest to p:

$$\min_{C \subset F} \frac{1}{A} \int_{p \in D} \delta(p) d(p, C)\, dp.$$

If there is one median point to be placed, the problem is known as the classical *Weber problem*; it was first discussed in Weber's 1909 book on the pure theory of location for industries [52] (see [54] for a modern survey). More generally, for a given number $k \geq 1$ of facilities, the problem is known as the *k-median problem.* A problem of similar type with a different objective function is the so-called *k-center problem*, where the goal is to find a set of k center locations such that the maximum distance of the demand set from the nearest center location is minimized.

With many practical motivations, geometric instances of facility location problems have attracted a major portion of the research to date. In these instances, the sets D of demand locations and F of feasible placements are modeled as points in some geometric space, typically $\Re^2$, with distances measured according to the Euclidean (L_2) or Manhattan (L_1) metric. In these geometric scenarios, it is natural to consider not only finite (discrete) sets F of feasible locations, but also (continuous) sets having positive area. For the classical Weber problem, the set F is the entire plane $\Re^2$, while D is some finite set of demand points.

Location theory distinguishes between discrete and continuous location theory (see [22]). However, for median problems, this distinction has mostly been applied to the set of feasible placements, distinguishing between discrete and continuous sets F. It is remarkable that, so far, *continuous* location theory of median problems has almost entirely treated *discrete* demand sets D [22, 45]. We should note that there are several studies in the literature that deal with k-center problems with continuous demand, e.g, see [36, 51], where demand arises from the continuous point sets along the edges in a graph. See [50] for results on the placement of k capacitated facilities serving a continuous demand on a one-dimensional interval. Also, k-center problems have been studied extensively in a geometric setting, see e.g. [1, 15, 23, 25, 26, 27, 28, 29, 30, 35, 37, 48, 49]. However, designing discrete algorithms for k-center problems can generally be expected to be more immediate than for k-median problems, since the set of demand points that determine a critical center location will usually form just a finite set of $d+1$ points in d-dimensional space.

*fekete@math.tu-berlin.de; Fachbereich Mathematik, Technische Universität Berlin, 10623 Berlin, Germany.

†jsbm@ams.sunysb.edu; http://www.ams.sunysb.edu/~jsbm. Department of Applied Mathematics and Statistics, State University of New York, Stony Brook, NY 11794-3600. Partially supported by NSF grant CCR-9732220, and by grants from Bridgeport Machines, Hughes Research Labs, NASA, Sandia National Labs, Seagull Technology, and Sun Microsystems.

‡Zentrum für Angewandte Informatik, Universität zu Köln, 50923 Köln, Germany.

Computational Geometry 2000 Hong Kong China

Continuous demand for k-median problems is also missing from the classification in [5]. We contend that the practical and geometric motivations of the problem make it very natural to consider exact algorithms for dealing with a continuous demand distribution for k-median problems: if a demand occurs at some position $p \in D$, according to some given probability density $\delta(p)$, then we may be interested in minimizing the *expected* distance $\int_{p\in D} d(c,p)\delta(p)dp$ for a feasible center location $c \in F$.

To the best of our knowledge, there are only few references that discuss k-median problems with continuous demand: See the papers [44, 58] for a discussion of continuous demand that arises probabilistically by considering a discrete demand in an unbounded environment with a large number of demand points, leading to a heuristic for optimal placement of many center points. Drezner [17] describes in Chapter 2 of his book that normally a continuous demand is replaced by a discrete one, for which the error is "quite pronounced for some problems". (See his chapter for some discussion of the resulting error.) Wesolowsky and Love [55] (and also in their book [34] with Morris) and Drezner and Wesolowsky [18] consider the problem of continuous demand for rectilinear distances. Practical motivations include the modeling of postal districts and facility design. They compute the optimal solution for one specific example, but fail to give a general algorithm. More recently, Carrizosa, Muñoz-Márquez, and Puerto [6, 7] use convexity properties for problems of this type to deal with the error resulting from nonlinear numerical methods for approximating solutions. It should be noted that the objective function is no longer convex when distances are computed in the presence of obstacles.

In this paper, we study the k-median problem, and its specialization to the Weber problem ($k = 1$), in the case of continuous demand sets. Another way to state our continuous Weber (1-median) problem is as follows: In a geometric domain (e.g., cluttered with obstacles), determine the ideal "meeting point" c^* that minimizes the average time that it takes an individual, initially located at a random point in D, to reach c^*. Another application comes from the problem of locating a fire station in order to minimize the average distance to points in a neighborhood, where we consider the potential emergencies (demands) to occur at points in a continuum (the region defining the neighborhood D). As we noted above, this objective function is different from the situation in which we want to minimize the *maximum* distance instead, a problem that has been studied extensively in the context of discrete algorithms.

Choice of Metric. Many papers on geometric location theory have dealt with continuous sets F of feasible placements, including [2, 4, 11, 12, 14, 19, 31, 32, 33, 55, 56, 57]. In the majority of these papers, distances are measured according to the L_1 metric. In fact, it was shown by Bajaj [3] that if L_2 distances are used, then even in the case of only five demand locations ($|D| = 5$), the problem cannot be solved using radicals; in particular, it cannot be solved by exact algorithmic methods that use only ruler and compass. (Chandrasekaran and Tamir [9] give a polynomial-time approximation scheme that uses the ellipsoid method.) In this paper, we too concentrate on the problem using the L_1 metric. While we can exactly solve some very simple special cases in the L_2 metric, in general the integrations that are required to solve the problem are likely to be just as intractable as the classic Weber problem. (Results that approximate the Euclidean cases are discussed more in the full paper.)

Summary of Results. In this paper, we give the first exact algorithmic results for location problems that are continuous on both counts, in the set D as well as the set F. In our model D and F are each given by polygonal domains. Our goal is to compute a set of k ($k \geq 1$) optimal centers in the feasible set F that minimize the average distance from a demand point of D to the nearest center point. Our results include:

(1) A linear-time ($O(n)$) algorithm for computing an optimal solution to the 1-median (Weber) problem when $D = F = P$, a simple polygon having n vertices, and distance is taken to be L_1 geodesic distance inside P.

(2) An $O(n^2)$ algorithm for computing an optimal 1-median for the case that $D = F = P$, a polygon with holes, and distance is taken to be (straight-line) L_1 distance.

(3) An $O(I + n \log n)$ algorithm (where $I = O(n^4)$ is the complexity of a certain arrangement) for computing an optimal 1-median for the case that $D = F = P$, a polygon with holes, and distance is taken to be L_1 geodesic distance inside P.

(4) A proof of NP-hardness for the k-median problem when the number of centers, k, is part of the input, and $D = F = P$ is a polygon with holes. This adds specific meaning to the statement by Wesolowsky and Love [55] that computing the optimal position of several locations "is obviously very tedious when (the number of locations) is very large".

(5) Generalizations of our results to the following cases: $D \neq F$; non-uniform probability densities over the demand set D; fixed-orientation metrics (generalization of L_1), which can be used to approximate the Euclidean metric; and, higher dimensions.

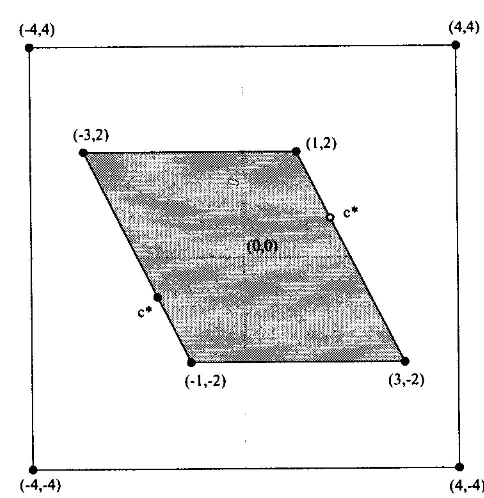

Figure 1: A simple example: The region $D = F = P$ is shown as a polygon with a single (parallelogram) hole, shown shaded. There are two points (marked "c^*") that minimize the average straight-line L_1 distance. (In this particular example, the optima lie on the boundary of P.)

Related Work. There is a vast literature on location theory; for a survey, see the book of Drezner [16], with its over 1200 citations that not only include papers dealing with mathematical aspects of optimization and algorithms, but also various applications and heuristics. A good overview of research with a mathematical programming perspective is given in the book of Mirchandani and Francis [38].

There has been considerable activity in the computational geometry community on facility location problems that involve computing geometric "centers" and medians of various types. The problem of determining a 1-center, point c, to minimize the maximum distance from c to a discrete set D of points is the familiar minimum enclosing disk problem, which has linear-time algorithms based, e.g., on the methods of Megiddo. The geodesic 1-center of simple polygons has an $O(n \log n)$ algorithm [46]. Exciting recent results of Sharir et al. [48, 20, 8] have yielded nearly-linear-time algorithms for the planar *two*-center problem. The more general p-center problem has been studied recently by [49].

The results outlined in this abstract constitute much of the PhD thesis of Weinbrecht [53]; most of the details necessarily omitted here are presented in-depth in the thesis.

2 Preliminaries

We will let $Z = (x, y)$ denote a candidate center point in $F \subseteq \Re^2$. (We concentrate on two-dimensional problems until Section 8, where we discuss extensions to higher dimensions.) We defer discussion of multiple center points ($k > 1$) to Section 7; for now, $k = 1$ and we consider the Weber (1-median) problem.

We let P denote a polygonal domain, possibly with holes. For purposes of our discussions, we focus on the case in which $D = F = P$: we restrict Z to P, which also equals the demand set. Our results will apply more generally to cases in which $D \neq F$, but we simplify our discussion here to the case $D = F$.

Furthermore, we restrict our discussion in this abstract to the case in which the demand is uniformly distributed over the set $D = P$, so our goal is to compute average distance.

Thus, the value of our objective function, denoted $f(Z)$, is given by the integral

$$f(Z) = f(x, y) = \int\int_{(u,v)\in P} d(Z, (u, v))\,du\,dv,$$

where $d(\cdot, \cdot)$ denotes either (straight-line) L_1 distance or geodesic L_1 distance within P. (We abuse notation slightly by writing $f(Z) = f((x, y)) = f(x, y)$.)

In order to analyze the k-median problem with respect to geodesic distances, we will utilize several definitions and results from the theory of geometric shortest paths among obstacles. The *shortest path map*, SPM(s), with respect to source point s, is a decomposition of P into cells according to the "combinatorial structure" (sequence of obstacle vertices along the path) of shortest paths from s to points in the cells. In particular, the *last* obstacle vertex along a shortest s-t path is the *root* of the cell containing t. The root of a cell lies on its boundary and can reach each point of the cell directly. We store with each vertex, v, of P the geodesic distance, $d_G(s, v)$, from s to v, as well as a pointer to the *predecessor* of v, which is the vertex (possibly s) preceding v in a shortest path from s to v. The predecessor pointers provide an encoding of the *shortest path tree*, SPT(s). The boundaries of cells consist of portions of obstacle edges and *bisector curves*. The bisector curves are the locus of points p that are (geodesically) equidistant from two roots, u and v: they satisfy $d_G(s, u) + d(u, p) = d_G(s, v) + d(v, p)$, where $d_G(\cdot, \cdot)$ denotes the shortest path (geodesic) distance function, and $d(\cdot, \cdot)$ denotes our underlying distance function (L_1, L_2, etc.). Figure 2 shows the types of bisectors that can arise in the L_1 metric; Figure 3 shows an example of an L_1 shortest path map. If t lies in the cell rooted at r, the geodesic distance to t is given by $d_G(s, t) = d_G(s, r) + d(r, t)$. Shortest path maps can be computed in optimal time $O(n \log n)$, both in the Euclidean metric and in the L_1 metric [24, 39, 41]. For more information, see the survey chapters by Mitchell [42, 43].

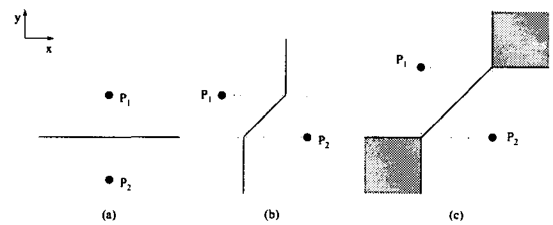

Figure 2: L_1 bisectors.

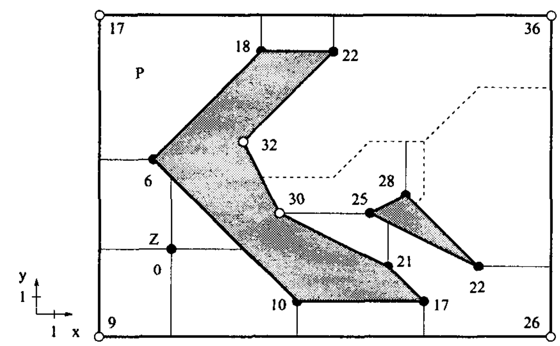

Figure 3: SPM(Z) for a polygon P having two holes. Vertices are labeled with their L_1 geodesic distances from Z.

3 Local Optimality Conditions

For any given center location Z, the objective function value $f(Z)$ that gives the average distance from Z to points of P can be evaluated by decomposing P into a set of "simple" pieces, computing the average distance for each piece, and then obtaining the total average distance as a weighted sum of the average distances for the pieces. In the case of straight-line L_1 distance, we simply use a trapezoidization (or triangulation) of P to determine the pieces; this can be done in linear time if P is simple, or in $O(n \log n)$ time if P has holes. In the case of geodesic distance, the shortest path map, SPM(Z), gives a decomposition of P into cells (each of which can be refined to yield a decomposition into $O(1)$-size pieces), each having a corresponding root vertex on its boundary. By computing the average distance from points of a cell to the cell's root r, and then adding this average to the distance $d_G(Z, r)$ and summing over all cells, we obtain the average geodesic distance, $f(Z)$.

The average distance associated with a single simple piece is given by the following result, whose straightforward proof is omitted here.

Lemma 1 *For a triangle τ with vertices A, B, and C, let a be the length of $\overline{AB}$, let c be the length of the altitude from C, and let b be the distance from A to the foot of the altitude from C. Then the average L_1 distance of points in τ from vertex A is $\frac{1}{3}(a+b+c)$.*

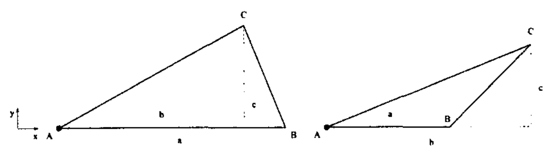

Figure 4: Notation used in computing the average L_1 distance of $\triangle ABC$ from A.

The objective function, $f(Z)$, is a continuous function of the location of the center Z. Since the set $F = P$ of feasible placements is a compact domain, it follows that there is an optimum, which must, necessarily, be locally optimal. The following lemma characterizes local optimality:

Proposition 2 *If $Z \in P$ is a local optimum of f, then there cannot be a feasible direction $h = (h_x, h_y)$, i. e., $Z + \varepsilon h \in P$ for sufficiently small ε), such that $f(Z) > f(Z + \varepsilon h)$. If the gradient ∇f exists in some neighborhood of Z, then this implies that for any feasible direction h, we have $\langle \nabla f, h \rangle \geq 0$.*

In particular, for interior points that are locally optimal, the gradient must be zero; for points in the interior of boundary edges, the gradient must be orthogonal to the boundary.

For any position Z of a center, the region P is subdivided into two pieces: the set $W(Z)$ of all points for which a shortest path to Z reaches Z in a ("western") direction with nonnegative x-coordinate, while $E(Z)$ is the set of all points for which a shortest path to Z reaches Z in an ("eastern") direction with non-positive x-coordinate. Similarly, we define $N(Z)$ and $S(Z)$. (Note that this partition looks different for straight-line and for geodesic distances; see Figure 5 (right) for the case of geodesic distances.)

Then we have the following:

Lemma 3 *Consider the objective function f for average straight-line or geodesic L_1 distance in a region P. Let $Z = (x_z, y_z)$ be a point in P, with neither x nor y coinciding with the x- or y-coordinate of a local minimum or maximum of the boundary of P, and let A be the area of P. Then the first partial derivatives of f are well-defined and given by:*

$$\begin{aligned} f_x(Z) &= \tfrac{1}{A}\left(W(Z) - E(Z)\right), \\ f_y(Z) &= \tfrac{1}{A}\left(S(Z) - N(Z)\right). \end{aligned} \qquad (1)$$

Proof.

We discuss the x-coordinates for straight-line L_1 distances, see Figure 5 (left). Consider $Z' = Z + (h, 0)$ for some sufficiently small h. Then we get

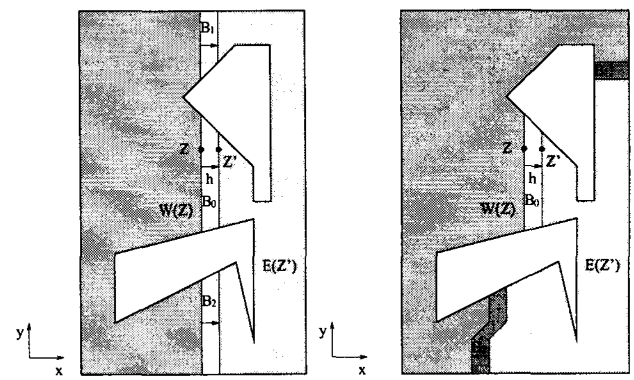

Figure 5: Computing the partial derivative in x-direction for (left) straight-line L_1 distances; (right) geodesic L_1 distances.

$$\begin{aligned} f_x(x_z, y_z) &= \lim_{h \to 0} \frac{f(x_z + h, y_z) - f(x_z, y_z)}{h} \\ &= \lim_{h \to 0} \left(\frac{\frac{1}{A}\left(hW(x_z, y_z) - hE(x_z + h, y_z) + O(h^2)\right)}{h} \right. \\ &\quad \left. + \frac{f(x_z, y_z) - f(x_z, y_z)}{h} \right) \\ &= \frac{1}{A} \lim_{h \to 0} \frac{hW(x_z, y_z) - hE(x_z + h, y_z) + O(h^2)}{h} \\ &= \frac{1}{A}\left(W(x_z, y_z) - E(x_z, y_z)\right) = \frac{1}{A}\left(W(Z) - E(Z)\right). \end{aligned}$$

For geodesic distances, we use a similar argument in the full paper (see [53]), but the situation looks as in Figure 5 (right).

□

It is readily seen that the function $f_x(Z) = \frac{1}{A}(W(Z) - E(Z))$ is not continuous at points where the boundary of P has a local minimum or maximum of its x-coordinates; we say that a chord through such a point is "critical."

However, $f_x(Z)$ is lower semi-continuous and monotonic, so there is a well-defined *vertical median chord* c_x at x-coordinate x_m such that $f_x(Z_1) < 0$ for all $Z_1 = (x_1, y_1)$ with $x_1 < x_m$, and $f_x(Z_2) > 0$ for all $Z_2 = (x_2, y_2)$ with $x_2 > x_m$. Similarly, there is a unique *horizontal median chord* c_y at y-coordinate y_m. We call $Z_m = (x_m, y_m)$ the *L_1 origin* of P.

In some situations, we make use of properties of higher-order derivatives of f. In particular, we use the following lemma:

Lemma 4 *Consider the objective function f for average straight-line or geodesic L_1 distance in a region P. Let $Z = (x_z, y_z)$ be a point in P, with neither x nor y coinciding with the x- or y-coordinate of a local minimum or maximum of the boundary of P, and no point of the boundary of $W(Z)$ or $E(Z)$ coinciding with a vertex of P. Then there is a neighborhood of Z for which f is a cubic function in x and y.*

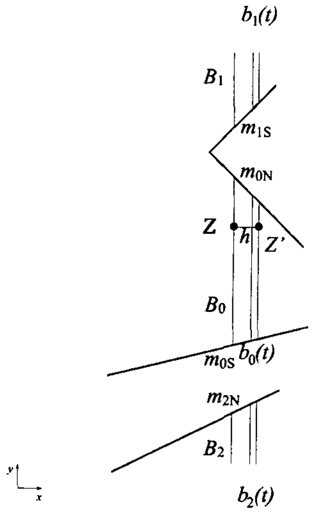

Figure 6: Trapezoids and slopes for higher order derivatives

Proof.
See Figure 6 for the case of straight-line distances, and consider changing Z to $Z' = Z + (h, 0)$. Then

$$W(x+h,y) = W(x,y) + B_0 + B_1 + \ldots + B_k \qquad (2)$$

and

$$E(x+h,y) = E(x,y) - B_0 - B_1 - \ldots - B_k, \qquad (3)$$

with trapezoids $B_0, \ldots, B_k$ as shown in the figure. Let $b_i(t)$ be the width of trapezoid B_i at x-coordinate t, and m_{iN} and m_{iS} be the slopes of the upper and lower line segments bounding B_i. Then

$$B_i = hb_i(x) + \frac{1}{2}h^2(m_{iN} - m_{iS}).$$

Then we get

$$f_{xxx}(x,y) = 2\sum_{i=0}^{k}(m_{iN} - m_{iS}).$$

This is a constant expression, so f is cubic in x.

For geodesic distances, the claim follows in a similar manner. Instead of being trapezoids bounded by two vertical chords, however, the areas B_i are bounded by two bisectors of the SPM. It follows from elementary properties of L_1 bisectors that these bisectors move in a parallel fashion, provided that during the move from Z to Z', no polygon vertex is hit by the bisector. Thus, the B_i's are pseudo-trapezoids, and the area of each B_i is still a quadratic function of h. (Details are contained in the full version of the paper.)
□

4 Simple Polygons

In this section, we show how to compute in optimal ($O(n)$) time a point that minimizes the average geodesic L_1 distance for a simple polygon.

It follows from our observations from the previous section that a point has locally optimal x-coordinate if and only if its vertical chord passes through the L_1 origin of the region. Let $Z_m = (x_m, y_m)$ be this origin.

In the following, we use the structure of simple polygons to show that Z_m has to lie within the feasible region P.

Theorem 5 *The point Z_m is feasible and thus a unique global optimum.*

Proof. Assume to the contrary that Z_m is infeasible. Then the chord c_x subdivides P into two or more pieces: let E denote the part to the right ("east") of c_x, and W the part to the left ("west") of c_x. Similarly, the chord c_y subdivides P into two or more pieces, with N ("north") denoting the part above c_y, and S ("south") denoting the part below c_y.

We distinguish the following cases, illustrated in Figure 7:

Case 0: Neither c_x nor c_y are critical. If $Z_m \notin P$, then, without loss of generality, the situation is as shown in Figure 7 (top).

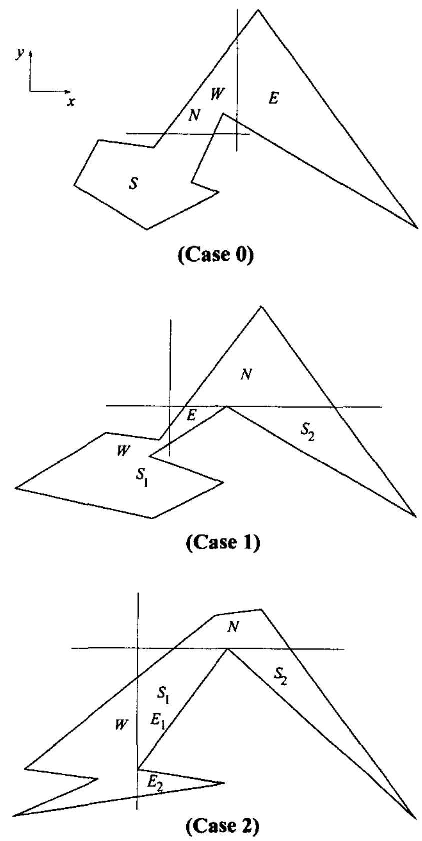

Figure 7: Intersection of median chords.

Then the two chords subdivide P into three pieces, E, $W \cap N$, and S. By assumption, we have $A(E) > 0$, $A(S) > 0$, $A(W \cap N) > 0$, and $A(E) + A(S) + A(W \cap N) = A(P)$. Furthermore, the local optimality assumption on x_m implies $A(E) = \frac{1}{2}A(P)$, and the local optimality assumption on y_m implies $A(S) = \frac{1}{2}A(P)$. This implies $A(W \cap N) = 0$, a contradiction.

Case 1: Precisely one of c_x and c_y is critical. If $Z_m \notin P$ then, without loss of generality, the situation is as shown in Figure 7 (center), with c_y being critical.

As in Case 0, the vertical chord c_x partitions P into two pieces, "eastern" E and "western" W, and, by local optimality, $A(E) = A(W) = \frac{1}{2}A(P)$. On the other hand, c_y partitions P into an "upper" piece N, a "lower" piece S_1, and a second "lower" piece S_2, all with positive area. (Note that the interior of S_2 may consist of several connected components, if the chord c_y meets several local maxima of the boundary of P.) By local optimality, we know that $A(N) \leq \frac{1}{2}A(P)$, $A(N)+A(S_2) \geq \frac{1}{2}A(P)$. Since $N \cup S_2 \subsetneq E$, we have a contradiction.

Case 2: Both c_x and c_y are critical. If $Z_m \notin P$ then, without loss of generality, the situation is as shown in Figure 7 (bottom).

As in Case 1, the horizontal chord c_y subdivides P into a "northern" piece N, a "southern" piece S_1, and a second "southern" piece S_2, all with positive area. Similarly, the vertical chord c_x subdivides P into a "western" piece W, an "eastern" piece E_1, and a second "eastern" piece E_2, all with positive area. By local optimality, we know that $A(N) \leq \frac{1}{2}A(P)$, $A(N) + A(S_2) \geq \frac{1}{2}A(P)$, and similarly, $A(W) \leq \frac{1}{2}A(P)$, $A(W) + A(E_2) \geq \frac{1}{2}A(P)$. Since $N \cup S_2 \subsetneq E$, we have a contradiction.
□

Theorem 6 *The point Z_m can be computed in linear time.*

Proof. We describe how to compute the x-coordinate x_m of Z_m; the y-coordinate is found in a similar manner.

In linear time (using Chazelle's algorithm [10]), we build the vertical trapezoidization of P, which is defined by drawing vertical chords through every vertex of P. Each piece (cell), τ_i, of the map is either a vertical-walled trapezoid or a triangle (a degenerate trapezoid). Consider the adjacency graph $\mathcal{G}$ of these pieces τ_i (i.e., the planar dual of the trapezoidization); since P is a simple polygon, $\mathcal{G}$ is a tree.

The x-coordinate x_m of the optimal point Z_m has to intersect some piece τ_m. Consider the partition of $\mathcal{G}$ induced by removing τ_m; let $C_{\max}$ be the heaviest connected component in this subdivision, and let $\tau_{\max}$ be the unique node of $C_{\max}$ adjacent to τ_m. The weight of $C_{\max}$ cannot exceed $A(P)/2$, or the local optimality of Z_m would be violated: moving Z_m from τ_m by an infinitesimal ε into $\tau_{\max}$ would reduce the distance to Z_m by ε for a set of points of a total area more than $A(P)/2$, while increasing it by at most ε for a set of points of total area less than $A(P)/2$. Thus, τ_m is a median in the weighted tree $\mathcal{G}$.

A median in a weighted tree can be computed in linear time (e.g., see Goldman [21]). This allows us to compute in linear time a piece τ_m that contains the critical coordinate x_m.

Once τ_m has been identified, it is easy to compute x_m: assuming that x_m is not given by one of the two vertical chords, we only have to identify the (unique) x-coordinate that splits P into two pieces of equal area. Since the slope of the boundary segments of τ_m does not change, this computation can be carried out by solving one quadratic equation.
□

5 Straight-Line L_1 Distance

For straight-line distances, a finite average is guaranteed even for disconnected regions P, as long as they are compact. As in the case of geodesic distances, we can consider local optimality for finding a global optimum of f. The example in Figure 8 shows that even for the special case of a simple polygon P, the L_1 origin of P may not be a feasible point.

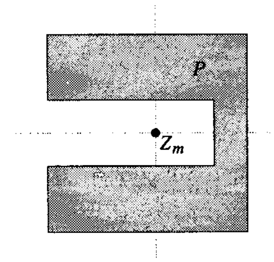

Figure 8: For straight-line distances, the L_1 origin of a simple polygon P may be infeasible

This makes it more involved to compute all local optima. In the following, we describe how to evaluate them in $O(n^2)$ time.

Theorem 7 *For straight-line L_1 distances, a point $Z^* = (x^*, y^*)$ in a polygonal region P that minimizes the average distance f to all points in P can be found in time $O(n^2)$.*

Proof. We apply the local optimality conditions. We start by computing in time $O(n \log n)$ the L_1 origin Z_m of P; if Z_m is feasible (i.e., $Z_m \in P$), we are done. If no interior point of P is a local optimum, then we have to consider the boundary of P. This yields a set E_d of $O(n)$ line segments and a set V_d of $O(n)$ vertices that we examine for local optimality.

We overlay the set of vertical and horizontal lines through all vertices of P with E_d, subdividing each segment in E_d into $O(n)$ pieces, bounded by a total of $O(n^2)$ "overlay" vertices V_o. Let E_o be the resulting set of $O(n^2)$ subsegments. See Figure 9.

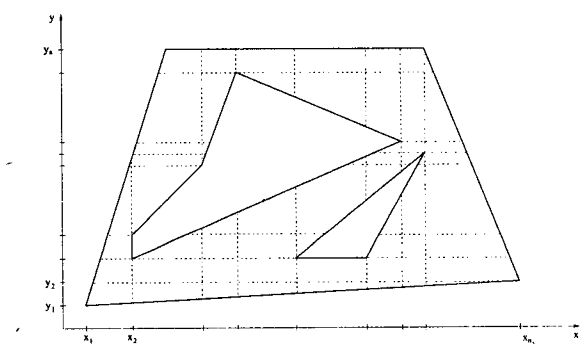

Figure 9: Subdivision of the polygon into cells.

Now we can examine the interior points $p_t = (t, y(t))$ of each edge $e_j \in E_o$ for local optimality. Let s_j be a vector parallel to e_j. By construction of E_o, the vertical and horizontal lines through p_t cannot encounter a vertex of P as p_t slides along one subsegment of E_o. Thus, it follows from Lemma 4 that $f_x(p_t)$ is a quadratic function in t, then so is $f_y(p_t)$. Therefore, considering the local optimality condition

$$\langle \nabla f, s_j \rangle = 0$$

for all $O(n^2)$ subsegments $e_j \in E_o$ yields a set of $O(n^2)$ quadratic equations in t. These can be solved in amortized time $O(n^2)$, since we can obtain the coefficients of

each quadratic equation in amortized constant time by advancing from cell to cell in the overlay arrangement. This gives, for each subsegment e_j, at most two local optima, $q_{j,1}$ and $q_{j,2}$. Let V_ℓ be the union of E_o and all $q_{j,1}$ and $q_{j,2}$. By construction, V_ℓ contains $O(n^2)$ elements, and all local optima of f occur at points of V_ℓ. Thus, our goal is to evaluate the objective function at each of these points and to select the best one. This is simply done in amortized time $O(1)$ per candidate, by walking over the overlay arrangement and incrementally updating the value of the objective function.
□

In many cases, the following property of straight-line medians can be applied for a reduction of the set of boundary segment that we need to consider. (See Figure 10 for an illustration.) If $Z_m = (x_m, y_m)$ is the L_1 origin of P and $p_1 = (x_1, y_1)$ and $p_2 = (x_2, y_2)$ are points in P, we say that p_1 *dominates* p_2, if p_1 lies in the rectangle spanned by Z_m and p_2.

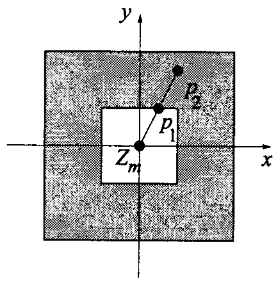

Figure 10: p_2 is dominated by p_1 and cannot be a local optimum

Lemma 8 *Let $p_1 = (x_1, y_1)$ and $p_2 = (x_2, y_2)$ be points in P. If p_1 dominates p_2, then $f(p_1) \le f(p_2)$.*

Proof. Suppose that $x_2 > x_1 \ge x_m$ and $y_2 \ge y_1 \ge y_m$. Then $W(p_2) \ge W(p_1) \ge \frac{A}{2}$ and $S(p_2) \ge S(p_1) \ge \frac{A}{2}$, so moving a center from p_2 to p_1 cannot increase the objective value.
□

Using a plane-sweep algorithm, we can to identify the non-dominated pieces of the boundary in time $O(n \log n)$. If this set has complexity $o(n)$, then we get a reduction of the overall complexity.

An easy instance of the problem solved in this section is shown in Figure 11 (based on the simple example of Figure 1); this example shows the necessity of solving quadratic equations in computing the solution. We get the non-dominated points p_1, p_2, p_3 and the edge $e_{2,3} = \overline{p_2 p_3}$.

6 Geodesic Distances in Polygons with Holes

Now we discuss an even more complicated case, which arises when considering geodesic L_1 distances in polygonal regions P that may have holes. Again, we analyze the set of locally optimal points: as long as a potential center can be moved in some axis-parallel fashion that lowers the average L_1 geodesic distance to all the points, it cannot be optimal.

The local optimality of a point Z is closely connected to the shortest path subdivision that it induces: for local optimality in the x-direction, the subdivision into $W(Z)$ and $E(Z)$ needs to be balanced; for local optimality in the y-direction, the subdivision into $N(Z)$ and $S(Z)$ needs to be balanced. The boundary between $W(Z)$ and $E(Z)$ is formed by bisectors for position Z. It follows from basic properties of shortest path maps that the total complexity of this boundary is $O(n)$. (See, e.g., [40].)

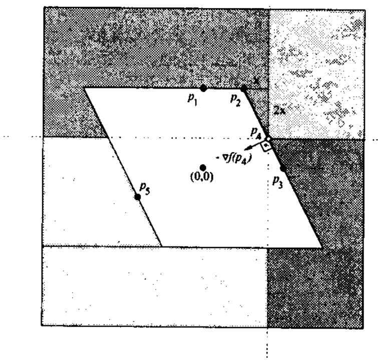

Figure 11: The example from Figure 1 is analyzed: point $p_4 = (2\sqrt{7} - \frac{9}{2}, 11 - 4\sqrt{7})$ and its mirror image, p_5, are the two optima.

As we showed in Lemma 4, there is a neighborhood for each point $Z \in P$ where the objective function f is cubic, provided that no bisector for Z meets a boundary vertex. This motivates the following lemma:

Lemma 9 *There is a subdivision of P of worst-case complexity $I = \Theta(n^4)$, such that f is a cubic function within each piece of the subdivision.*

Proof. Lemma 4 implies that we are done if we can compute a subdivision of the claimed complexity such that we can move continuously between any two points in the interior of a connected cell of the subdivision, without any bisector encountering a vertex of the polygon during this motion. Provided that there is a position Z for which a bisector encounters a vertex v of the polygon, this vertex v is contained in $W(Z)$ as well as in $E(Z)$. Thus, there are two topologically different paths from Z to v, one fully contained in $W(Z)$, the other contained in $E(Z)$. This implies that there are two topologically different paths from v to Z, i. e., Z must lie on a bisector of v. Therefore, the required subdivision is obtained by considering the $O(n)$ bisectors of the $O(n)$ polygon vertices. Each bisector has a complexity of $O(n)$, so the subdivision is defined by the overlay of $O(n^2)$ line segments, yielding an arrangement of worst-case complexity $I = O(n^4)$. Examples exist (see the full paper) to show that this bound on I is tight in the worst case. (Chiang and Mitchell [13] have studied similar arrangements that arise in overlaying shortest path maps in the Euclidean shortest path metric.)
□

Considering the local optima on each cell of the arrangement allows us to obtain the following:

Theorem 10 *For geodesic L_1 distances, a feasible point $Z^* = (x^*, y^*)$ in a polygonal region P with holes that minimizes the average distance f to all points in P can be found in worst-case time $O(I + n \log n)$.*

Proof. (sketch) The idea is to reduce the problem to a set of $O(I)$ candidates, with $O(1)$ such candidates in the interior or on the boundary of each cell. The function parameters for f within each cell can be determined in total time $O(I + n \log n)$ by traversing the arrangement and doing updates when changing from one cell to its neighbor. After determining the $O(I)$ candidate locations, we can determine a best among them by computing their objective values, again in total time $O(I + n \log n)$ by being careful how to update the objective function values.

Thus, consider an arbitrary cell of the decomposition. If there is a local minimum interior to the cell, the gradient ∇f must vanish. Since f is cubic within the cell, this means that we get a system of two quadratic equations (both components of the gradient must be zero) with two variables (x and y). Such a system can be solved in constant time using radicals.

Similarly, we can determine the local optima with respect to variation along a boundary segment of a cell. For each segment, the gradient needs to be orthogonal to the segment. As in the straight-line case, this yields a quadratic equation that can be solved in constant time.

Finally, there are $O(I)$ vertices in the arrangement, each of which we consider as candidates.
□

7 Many Centers

We now consider the k-median problem of placing k centers into a polygonal region P, such that the overall average distance of all points $p \in P$ to their respective closest centers is minimized. We show that this problem is NP-hard for polygons with holes.

Here, we give only an outline of the full proof. Our construction uses a reduction from PLANAR 3SAT.

First, the planar graph corresponding to an instance I of PLANAR 3SAT is represented in the plane as a *planar rectilinear layout*, with each vertex corresponding to a horizontal line segment, and each edge corresponding to a vertical line segment that intersects precisely the line segments corresponding to the two incident vertices. There are well-known algorithms (e.g., [47]) that can achieve such a layout in linear time and linear space. See Figure 12.

Next, the layout is modified such that the line segments corresponding to a vertex and all edges incident to it are replaced by a loop – see Figure 13 (top). At each vertex corresponding to a clause, three of these loops (corresponding to the respective literals) meet. Finally, the edges of all loops are replaced by a sequence of small squares that are interconnected by narrow corridors; also, each vertex for a clause is replaced by a single small square and linked to the adjacent variable loops by three narrow corridors. See Figure 13 (bottom) for the overall picture.

Let $3k$ be the total number of squares in all variable loops. It can be checked that there is a placement of k variables with a low overall average distance of points to their closest centers (i.e., a low objective function value), if and only if there is a placement such that each square has a center or is next to a square with a center. This means that each clause square must be next to a square with a

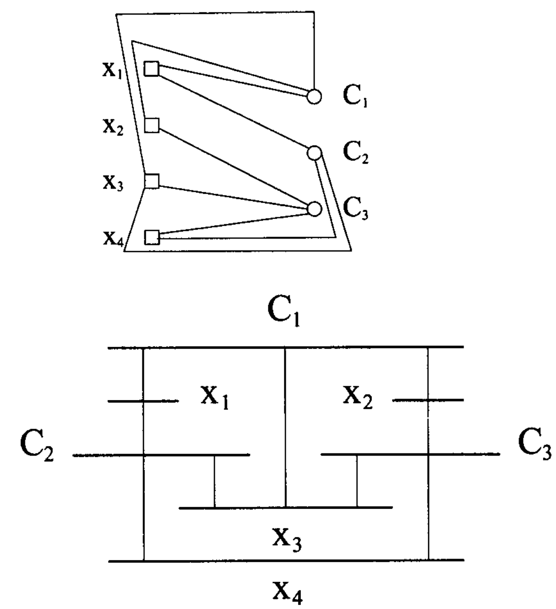

Figure 12: The graph G_I for the Planar 3SAT instance $I = (x_1 \vee x_2 \vee x_3) \wedge (\bar{x}_1 \vee \bar{x}_3 \vee x_4) \wedge (\bar{x}_2 \vee x_3 \vee \bar{x}_4)$, and its geometric representation.

center. We prove that this is possible if and only if there is a satisfying truth assignment for the PLANAR 3SAT instance I. Details are contained in the full version of the paper.

8 Conclusions

In this paper, we have given the first exact algorithmic results for the Weber problem for a continuous set of demand locations. We have shown that for L_1 distances in the plane, we can determine an optimum center in polynomial time, with the complexity ranging from $O(n)$ for the case of geodesic distances in simple polygons, to $O(n^2)$ for straight-line distances in general polygonal regions, and $O(n^4)$ for geodesic distances in polygons with holes.

Our results rely on a careful understanding of the local optimality criteria. Our local optimality conditions generalize to the case of more general (non-uniform) nonnegative demand densities $\delta(p)$ by using the following observation. Regardless of the demand density function, any center location $Z \in P$ induces a subdivision of P into $E(Z)$ and $W(Z)$, and $N(Z)$ and $S(Z)$ by shortest-path bisectors. Then the local optimality condition on Z requires that $E(Z)$ and $W(Z)$, and $S(Z)$ and $N(Z)$ are balanced in the following sense: instead of requiring that $E(Z)$ and $W(Z)$, and $N(Z)$ and $S(Z)$ have the same area, we consider the points Z where the integrals $\int_{p \in W(Z)} \delta(p)dp$ and $\int_{p \in E(Z)} \delta(p)dp$, and $\int_{p \in N(Z)} \delta(p)dp$ and $\int_{p \in S(Z)} \delta(p)dp$ are the same. Points with these properties are called *δ-medians.* Similar ideas can be used for describing boundary points. If, for a particular δ, there is a limited number of δ-medians, they can be computed in polynomial time, and it is possible to compare objective values in polynomial time, then we can determine a δ-center for the given region. This includes the case in which the demand function is given by point weights in combination with a uniform demand distribution over P, which is a problem

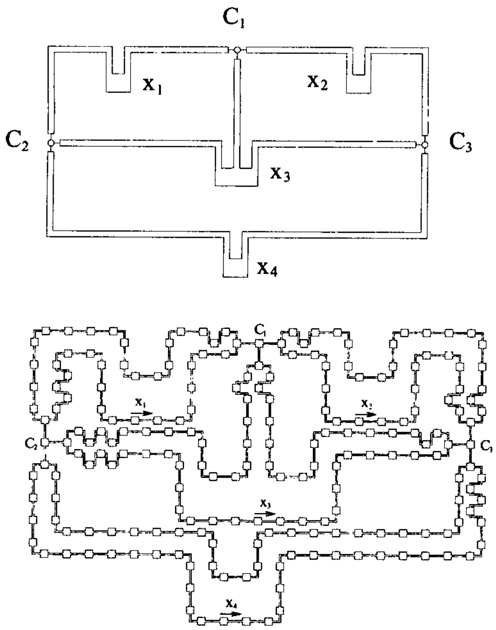

Figure 13: Replacing variables by loops (left); final polygon (right)

formulated by Wesolowsky and Love [55]. It is also easy to see that the above methods can be applied for the case in which $F \neq D$.

Approximation algorithms for the Euclidean case can be devised based on generalizing the L_1 metric results to fixed orientation metrics. The local optimality conditions become more complex; however, the inherent algebraic complexity remains the same, for any metric whose disks are convex polygons.

Our methods can also be applied in higher dimensions, by generalizing the local optimality conditions and carrying through the analysis in a very similar manner to the two-dimensional case. Figure 14 shows that a generalization of Lemma 6, however, does not hold in three-dimensional space (the main reason being that any axis-parallel plane cuts the region into not more than two pieces), so we cannot use the same idea to exploit simplicity for achieving a better complexity than in the case with holes. However, we can still use a subdivision into cells and study the objective function within each cell. As in the two-dimensional case, the objective function is cubic for each coordinate, if P is a polyhedral region.

Acknowledgements

We would like to thank Arie Tamir, Horst Hamacher, and Justo Puerto for pointing out relevant references, and Iris Weber for a heroic literature effort.

References

[1] P. K. Agarwal, M. Sharir, and E. Welzl. The discrete 2-center problem. In *Proc. 13th Annu. ACM Sympos. Comput. Geom.*, pages 147–155, 1997.

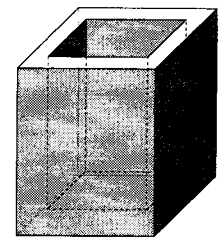

Figure 14: In three-dimensional space, there may not be a feasible point that is a median of P in all coordinates.

[2] Y. P. Aneja and M. Parlar. Algorithms for Weber facility location in the presence of forbidden regions and/or barriers to travel. *Transportation Science*, 28(1):70–76, 1994.

[3] C. Bajaj. The algebraic degree of geometric optimization problems. *Discrete and Computational Geometry*, 3:177–191, 1988.

[4] R. Batta, A. Ghose, and U. S. Palekar. Locating facilities on the Manhattan metric with arbitrarily shaped barriers and convex forbidden regions. *Transportation Science*, 23(1):26–36, 1989.

[5] M. L. Brandeau and S. S. Chiu. An overview of representative problems in location research. *Management Science*, 35(6):645–674, 1989.

[6] E. Carrizosa, M. Muñoz-Márquez, and J. Puerto. Location and shape of a rectangular facility in $\Re^2$. convexity properties. *Mathematical Programming*, 83:277–290, 1998.

[7] E. Carrizosa, M. Muñoz-Márquez, and J. Puerto. The Weber problem with regional demand. *European Journal of Operational Research*, 104:358–365, 1998.

[8] T. M. Chan. More planar two-center algorithms. *Comput. Geom. Theory Appl.*, 13:189–198, 1999.

[9] R. Chandrasekaran and A. Tamir. Algebraic optimization: the Fermat-Weber location problem. *Math. Program.*, 46(2):219–224, 1990.

[10] B. Chazelle. Triangulating a simple polygon in linear time. *Discrete and Computational Geometry*, 6:485–524, 1991.

[11] R. Chen and G. Y. Handler. The conditional p-center problem in the Plane. *Naval Research Logistics*, 40:117–127, 1993.

[12] V. Chepoi. A multifacility location problem on median spaces. *Discrete Applied Mathematics*, 64:1–29, 1996.

[13] Y.-J. Chiang and J. S. B. Mitchell. Two-point Euclidean shortest path queries in the plane. In *Proc. 10th ACM-SIAM Sympos. Discrete Algorithms*, pages 215–224, 1999.

[14] J. Choi, C.-S. Shin, and K. Kim. Computing weighted rectilinear median and center set in the presence of obstacles. In *Ninth Annual International Symposium on Algorithms and Computation*, 1998.

[15] Z. Drezner. On the rectangular p-center problem. *Naval Res. Logist. Q.*, 34:229–234, 1987.

[16] Z. Drezner. *Facility Location: A Survey of Applications and Methods.* Springer Series in Operations Research. Springer–Verlag, New York, 1995.

[17] Z. Drezner. Replacing discrete demand with continuous demand. In Z. Drezner, editor, *Facility Location: A Survey of Applications and Methods*, Springer Series in

Operations Research, chapter 2. Springer–Verlag, New York, 1995.

[18] Z. Drezner and G. O. Wesolowsky. Optimal location of a facility relative to area demands. *Naval Research Logistics Quarterly*, 27:199–206, 1980.

[19] R. Durier and C. Michelot. On the set of optimal points to the Weber problem: further results. *Transportation Science*, 28(2):141–149, 1994.

[20] D. Eppstein. Faster construction of planar two-centers. In *Proc. 8th ACM-SIAM Sympos. Discrete Algorithms*, 1997.

[21] A. J. Goldman. Optimal center location in simple networks. *Transp. Sci.*, 5:240–255, 1971.

[22] H. W. Hamacher and S. Nickel. Classification of location problems. *Location Science*, 6:229–242, 1998.

[23] J. Hershberger. A faster algorithm for the two-center decision problem. *Inform. Process. Lett.*, 47:23–29, 1993.

[24] J. Hershberger and S. Suri. An optimal algorithm for Euclidean shortest paths in the plane. *SIAM J. Comput., to appear.*

[25] D. S. Hochbaum and D. Shmoys. A best possible heuristic for the k-center problem. *Math. Oper. Res.*, 10:180–184, 1985.

[26] R. Z. Hwang, R. C. T. Lee, and R. C. Chang. The slab dividing approach to solve the Euclidean p-center problem. *Algorithmica*, 9:1–22, 1993.

[27] O. Kariv and S. L. Hakimi. An algorithmic approach to network location problems. i: The p-centers. *SIAM J. Appl. Math.*, 37:513–538, 1979.

[28] S. Khuller and Y. J. Sussmann. The capacitated k-center problem. In *Proc. 4th Annu. European Sympos. Algorithms*, volume 1136 of *Lecture Notes Comput. Sci.*, pages 152–166. Springer-Verlag, 1996.

[29] M. T. Ko and Y. T. Ching. Linear time algorithms for the weighted tailored 2-partition problem and the weighted rectilinear 2-center problem under L_∞-distance. *Discrete Appl. Math.*, 40:397–410, 1992.

[30] M. T. Ko, R. C. T. Lee, and J. S. Chang. An optimal approximation algorithm for the rectilinear m-center problem. *Algorithmica*, 5:341–352, 1990.

[31] A. Kolen. Equivalence between the direct search approach and the cut approach to the rectilinear distance location problem. *Operations Research*, 29(3):616–620, 1981.

[32] Y. Kusakari and T. Nishizeki. An algorithm for finding a region with the minimum total L_1-distance from prescribed terminals. In *Eigth Annual International Symposium on Algorithms and Computation*, pages 324–333. Springer–Verlag, 1997.

[33] R. C. Larson and G. Sadiq. Facility locations with the Manhattan metric in the presence of barriers to travel. *Operations Research*, 31(4):652–669, 1983.

[34] R. F. Love, J. G. Morris, and G. O. Wesolowsky. *Facilities Location: Models & Methods.* North Hollandde Gruyter, New York, 1988.

[35] N. Megiddo. The weighted Euclidean 1-center problem. *Math. Oper. Res.*, 8(4):498–504, 1983.

[36] N. Megiddo and A. Tamir. New results on the complexity of p-center problems. *SIAM J. Comput.*, 12:751–758, 1983.

[37] N. Megiddo and E. Zemel. A randomized $O(n \log n)$ algorithm for the weighted Euclidean 1-center problem. *J. Algorithms*, 7:358–368, 1986.

[38] P. B. Mirchandani and R. L. Francis, editors. *Discrete Location Theory.* Wiley, New Yorkl, 1990.

[39] J. S. B. Mitchell. An optimal algorithm for shortest rectilinear paths among obstacles. In *Abstracts 1st Canad. Conf. Comput. Geom.*, page 22, 1989.

[40] J. S. B. Mitchell. A new algorithm for shortest paths among obstacles in the plane. *Ann. Math. Artif. Intell.*, 3:83–106, 1991.

[41] J. S. B. Mitchell. L_1 shortest paths among polygonal obstacles in the plane. *Algorithmica*, 8:55–88, 1992.

[42] J. S. B. Mitchell. Shortest paths and networks. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 24, pages 445-466. CRC Press LLC, Boca Raton, FL, 1997.

[43] J. S. B. Mitchell. Geometric shortest paths and network optimization. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, page to appear. Elsevier Science Publishers B.V. North-Holland, Amsterdam, 1999.

[44] C. Papadimitriou. Worst case and probabilistic analysis of a geometric location problem. *SIAM J. Computing*, 3:542–557, 1981.

[45] F. Plastria. Continuous location problems. In Z. Drezner, editor, *Facility Location: A Survey of Applications and Methods*, Springer Series in Operations Research, chapter 11. Springer–Verlag, New York, 1995.

[46] R. Pollack, M. Sharir, and G. Rote. Computing of the geodesic center of a simple polygon. *Discrete Comput. Geom.*, 4:611–626, 1989.

[47] P. Rosenstiehl and R. E. Tarjan. Rectilinear planar layouts and bipolar orientations of planar graphs. *Discrete and Computational Geometry*, 1:343–353, 1986.

[48] M. Sharir. A near-linear algorithm for the planar 2-center problem. *Discrete Comput. Geom.*, 18:125–134, 1997.

[49] M. Sharir and E. Welzl. Rectilinear and polygonal p-piercing and p-center problems. In *Proc. 12th Annu. ACM Sympos. Comput. Geom.*, pages 122–132, 1996.

[50] H. D. Sherali and F. L. Nordai. NP-hard, capacitated, balanced p-median problems on a chain graph with a continuum of link demands. *Mathematics of Operations Research*, 13:32–49, 1988.

[51] A. Tamir. On the solution value of the continuous p-center location problem on a graph. *Mathematics of Operations Research*, 12:340–349, 1987.

[52] A. Weber. *Über den Standort der Industrien, 1. Teil: Reine Theorie des Standortes.* Tübingen, Germany, 1909.

[53] K. Weinbrecht. *Kontinuierliche Standortprobleme in Polygonen.* PhD thesis, Universität zu Köln, 1999.

[54] G. Wesolowsky. The Weber problem: History and perspective. *Location Science*, 1:5–23, 1993.

[55] G. O. Wesolowsky and R. F. Love. Location of facilities with rectangular distances among point and area destinations. *Naval Research Logistics Quarterly*, 18:83–90, 1971.

[56] G. O. Wesolowsky and R. F. Love. The optimal location of new facilities using rectangular distances. *Operations Research*, 19:124–130, 1971.

[57] G. O. Wesolowsky and R. F. Love. A nonlinear approximation method for solving a generalized rectangular distance Weber problem. *Management Science*, 11:656–663, 1972.

[58] E. Zemel. Probabilistic analysis of geometric location problems. *SIAM J. Alg. and Discrete Methods*, 6:189–200, 1985.