

Planning a Purely Translational Motion for a Convex Object in Two-Dimensional Space Using Generalized Voronoi Diagrams*

Daniel Leven and Micha Sharir

School of Mathematical Sciences, Tel Aviv University, Tel Aviv, Israel

Abstract. An $O(n \log n)$ algorithm for planning a purely translational motion for a simple convex object amidst polygonal barriers in two-dimensional space is given. The algorithm is based on a new generalization of Voronoi diagrams (similar to that proposed by Chew and Drysdale [1] and by Fortune [2]), and adapts and uses a recent technique of Yap for the efficient construction of these diagrams.

1. Introduction

In this paper we consider the following problem. Given a convex object B of relatively simple shape and a set of polygonal obstacles, which are composed of a set of n line segments (“walls”), plan a purely translational motion of B from a given initial position to a desired final position during which B avoids collision with the obstacles, or determine that no such motion exists. A recent paper by Kedem and Sharir [5] (see also [4]) has presented an $O(n \log^2 n)$ time algorithm for solving this problem, based on calculating the vector differences $A_i - B$ of each of the obstacles A_i and B , and then on calculating the union of these differences, obtaining an “obstacle space” K for a reference point on B , and thus reducing the motion-planning problem to the calculation of the connected components of the complement of K . Another related work is that of Ó’Dúnlaing and Yap [10] for the specific case where B is a disc. They showed that if the free space available for B is bounded, a motion of B between two positions Z, Z' exists if and only if there exists a motion of (the center of) B along the Voronoi diagram of the obstacles from some initial position W to some final position W' on that Voronoi diagram, where W, W' are suitably defined retractions of Z, Z'

* Work on this paper by the second author has been supported in part by a grant from the U.S.-Israeli Binational Science Foundation.

into the Voronoi diagram. The complexity of the algorithm in [10] is $O(n \log n)$, which is the time required to construct the Voronoi diagram for a set of n line segments [14].

The idea behind the present paper is similar to that of [10]. The reason that the Voronoi diagram is a suitable tool for the case where B is a disc is that given a position Z of (the center of) B which is not on the Voronoi diagram of the obstacles, we can move B in a canonical manner to a position W on the Voronoi diagram in such a way that the distance of B from the set of obstacles when positioned at W is greater than the distance of B from the obstacles when positioned at Z . Hence no harm is done if we move B from Z to W . Our idea is to define a generalized Voronoi diagram which has a similar property for a given convex object B of relatively simple shape (e.g., a polygon with a small number of sides) which, however, is allowed only to translate and not to rotate (for a disc this further constraint is of course redundant). As it turns out, this generalized diagram has similar properties to those of more conventional Voronoi diagrams; it has $O(n+N)$ edges and vertices, and can be constructed using similar techniques in $O(n \log N)$ time, where N is the number of obstacles and n is the total number of obstacle corners, and where the complexity of B is assumed to be fixed. This in turn also yields an $O(n \log N)$ time algorithm for planning a purely translational motion of arbitrary convex objects of simple shape.

The idea of using generalized Voronoi diagrams of the sort proposed here has recently and independently been suggested by Chew and Drysdale [1] and by Fortune [2]. However, they had different motivations for introducing such a structure; for example, Chew and Drysdale were interested in the problem of determining the largest convex shape homothetic to a given shape, which will fit in the plane without touching any of n given points. Their analysis thus corresponds to assuming that the given obstacles are all singleton points. The present paper contains a detailed analysis of the resulting diagram, which is more general than the one studied in [1], presents a different method for the calculation of the diagram, and discusses applications of the diagram to the translational motion planning problem mentioned above.

This paper is organized as follows. In Section 2 we define the appropriate generalization of Voronoi diagrams which we call a B -Voronoi diagram, and state and prove its main properties. In Section 3 we present an algorithm for the calculation of this diagram in $O(n \log N)$ time and in Section 4 we show how to use this diagram for the solution of the above-mentioned translational motion planning problem.

2. The B -Voronoi Diagram for a Set of Convex Polygonal Objects

Let B be a compact convex object of some simple shape (the term simple will be defined more precisely in the next section), which we assume to have nonempty interior, and let O be an arbitrarily chosen fixed interior point of B . Let S be a set of N closed interior-disjoint polygonal objects ("obstacles") composed of a total of n line segments ("walls") which form the boundary of the free space V

available for B . We assume that each object in S is convex and allow objects in S to share a common wall (in this way nonconvex obstacles may be represented as the union of several interior-disjoint convex subparts). To avoid certain technical difficulties, we will assume that the obstacles in S and B are placed in “general position” and thus exclude degenerate configurations of these objects. These assumptions are required either to prove the properties which we state in this section or for the algorithm which we describe in Section 3, and are as follows:

Assumptions on general position:

- (a) No boundary edge of B or of an obstacle in S is vertical.
- (b) No boundary edge of B is parallel to the boundary edge of an obstacle or to a line joining a pair of boundary corners of these obstacles.
- (c) No more than three objects in S can be tangent to any homothetic copy of B and lie in its exterior.
- (d) No homothetic copy of B is such that two adjacent boundary segments of it simultaneously touch two wall corners and their common endpoint touches a wall.

Remarks. (1) Assumption (a) is required only for the construction of the diagram in Section 3, but not for proving its properties, and can always be ensured by an appropriate rotation of the coordinate frame.

(2) Assumption (b) is required to ensure that the B -Voronoi diagram will be one-dimensional. Degenerate configurations which occur when this assumption is not satisfied can be handled by standard techniques; for example, one can use the method of infinitesimal perturbations mentioned in [13].

(3) Assumption (c) resembles the assumption, usually made in the analysis of the standard Voronoi diagram for a set of points, that no more than three points in the set may be cocircular (see [11]). If this assumption is violated the diagram may have vertices of degree greater than 3. This will cause some problems in the contour tracing procedure described in Section 3 but, as in the calculation of standard diagrams, the algorithm can be slightly modified to handle such a degeneracy.

(4) Assumption (d) is required to prove property (9) below which states roughly that the diagram does not contain “degenerate” vertices. In practice a slight modification of our algorithm will suffice to make it work even if this assumption is not satisfied.

We will refer to the position of B in the plane in which O lies at the origin as the *standard position* of B and denote by B_0 the set of points in B when given that standard position. Let p, q be two points in V . We define the B -distance from p to q (or equivalently, the B -closeness of p to q), denoted by $d_B(p, q)$, as

$$d_B(p, q) = \inf\{\lambda : q \in p + \lambda B_0\}.$$

Note that since O is an interior point of B , the B -distance function is always finite and continuous, obeys the triangle inequality but need not necessarily be

symmetric, and is thus not induced by a metric in general (such a notion was introduced by Minkowski in 1911 and was called a *convex distance function* (see [7])). The B -distance from a point $p \in V$ to a given object $s_i \in S$ is defined in a natural way as

$$d_B(p, s_i) = \inf\{\lambda : p + \lambda B_0 \cap s_i \neq \emptyset\}.$$

(Note that under our assumptions that B and S are in general position, there exists a unique $y \in s_i$ at which the B -distance from p to s_i is attained.)

Using d_B we define the B -Voronoi diagram $\text{Vor}_B(S)$ associated with the set S and B as follows. For each $i \neq j$ define

$$H(i, j) = \{y \in E^2 : d_B(y, s_i) \leq d_B(y, s_j)\},$$

i.e., the set of all points whose B -distance to s_i is no greater than their B -distance to s_j . Then define the (closed) B -Voronoi cell $C_B(s_i)$ associated with s_i to be

$$C_B(s_i) = \bigcap_{j \neq i} H(i, j),$$

i.e. the set of all points y whose B -closeness to s_i is not greater than y 's B -closeness to any other element of S . Finally, the B -Voronoi diagram $\text{Vor}_B(S)$ is defined to be the set of points which belong to more than one B -Voronoi cell. The edges of this diagram are maximal connected arcs which belong to exactly two B -Voronoi cells and the vertices are those points which belong to more than two cells. To simplify the following analysis we shall assume that for each pair of objects s_1, s_2 in S which share a common boundary segment I , the exterior angles formed between s_1, s_2 at the endpoints of I are both less than or equal to 180° , and thus the locus of points equidistant from any given pair of objects in S is always a one-dimensional set (if that angle is 180° , this property follows from our assumption that B and S are in general position), so that in this case the B -Voronoi diagram is indeed one-dimensional. If this assumption is violated at some point x common to the boundary of two obstacles s_1, s_2 (i.e., the exterior angle between s_1 and s_2 at x is greater than 180°), then we can treat x as a singleton obstacle, replace s_1, s_2 by $s_1 - \{x\}, s_2 - \{x\}$, respectively, and redefine B -closeness similar to the definition of closeness in [6] and [14]. However, for the sake of simplicity we will continue to make this assumption, which can always be enforced when partitioning a nonconvex obstacle into a collection of openly disjoint convex parts. By our assumptions that B and S are in general position no point on $\text{Vor}_B(S)$ lies in more than three B -Voronoi cells. Figure 2.1 shows an example of such a B -Voronoi diagram. Note that the B -Voronoi diagram for the case where B is a disc is exactly the standard Voronoi diagram of the set S of obstacles [10].

The B -Voronoi diagram just defined has the following properties.

- (1) The collection of B -Voronoi cells covers the whole plane.

Proof. Follows immediately from the definition of d_B and the fact that O is an interior point of B . \square




Fig. 2.1. The B -Voronoi diagram.

(2) Let y be a point in some cell $C_B(s_i)$ and let y' be the B -closest point to it on s_i . Then

- (a) The line segment $I = yy'$ is wholly contained in $C_B(s_i)$ and in addition, if there is a point $z \in I$ which is contained also in another cell $C_B(s_j)$ then the whole segment yz is also contained in $C_B(s_j)$ (we capture this property by calling $C_B(s_i)$ *star-shaped* with respect to s_i).
- (b) If p is any point such that $y'yp$ is a straight line, then the B -closest point on s_i to p is y' .

Proof. (a) Suppose that there is an interior point $z \in I$ which belongs to a second cell $C_B(s_j)$ and let z' be the B -closest point to z on s_j (see Fig. 2.2(a)). We have by our assumption

$$d_B(z, s_j) \leq d_B(z, s_i).$$

Let $d_B(y, s_i) = \lambda$. If the inequality were strict (i.e., z did not belong to $C_B(s_i)$) then z' would belong to the interior of the triangle $yy'y''$ where yy'' is a line segment parallel to zz' such that $d_B(y, y'') = \lambda$. But since the triangle $yy'y''$ is wholly contained in $y + \lambda B_0$, this would imply $d_B(y, s_j) < d_B(y, s_i)$ contrary to our assumptions. Therefore z' must lie on the segment $y'y''$ and also be a boundary point of $y + \lambda B_0$ which, however, implies that every point on I between y and z is equally close to s_i, s_j as asserted.

(b) Let λ be the scalar multiple of B such that $\lambda B_0 + y$ touches s_i at y' . Since y' is the B -closest point to y on s_i and since B and s_i are convex there is a line T passing through y' which separates them (see Fig. 2.2(b)). Let $\lambda' > \lambda$ be the multiple of B such that $\lambda' B_0 + p$ also touches s_i at y' . It is easily checked that T also separates $\lambda' B_0 + p$ and s_i . But then since B and s_i are convex it follows that the B -distance from p to any point on T (and hence also to any point in s_i) is greater than or equal to λ' , and this proves our claim. \square




Fig. 2.2

(3) Suppose I_1, I_2 are two line segments in $C_B(s_i)$ which connect two points x_1, x_2 in $C_B(s_i)$ to their respective B -closest points y_1, y_2 on s_i . Then either I_1, I_2 do not intersect except possibly at their endpoints or else one of these segments is wholly contained in the other.

Proof (the technique used here resembles that of [5]). Suppose I_1, I_2 intersect at a single point and let $d_B(x_1, s_i) = \lambda$ and $d_B(x_2, s_i) = \lambda'$. Let z_1 (resp. z_2) be a point at a B -distance λ (resp. λ') from x_1 (resp. x_2) in the direction $x_2 y_2$ (resp. $x_1 y_1$) (see Fig. 2.3). Clearly the triangle $x_1 z_1 y_1$ is contained in $x_1 + \lambda B_0$ and the triangle $x_2 z_2 y_2$ is contained in $x_2 + \lambda' B_0$ and therefore both these triangles do not contain a point (in particular the points y_1, y_2) of s_i is their interior. But since the two triangles $x_1 z_1 y_1, x_2 z_2 y_2$ are similar and have pairwise parallel sides, and their two sides I_1, I_2 intersect one another, we must have either $y_1 \in \Delta x_2 z_2 y_2$ or $y_2 \in \Delta x_1 z_1 y_1$. Therefore $z_1 y_2 y_1 z_2$ must be a straight line segment and $y_1 y_2$ must therefore be a subsegment of the boundary of s_i . But this means that the boundary of B contains a line segment which is parallel to a boundary segment of s_i which contradicts our assumptions of general position. \square

(4) Each edge of the B -Voronoi diagram is simple (i.e., the removal of a single point on such an edge divides it into at most two connected components).




Fig. 2.3

Proof. Let p be a point on an edge e on the boundary of $C_B(s_1), C_B(s_2)$ such that the removal of p divides e into more than two disconnected components, and let C_1, C_2, C_3 be any three such components in some small neighborhood of p . Define $\phi_i(x), i=1, 2$, for $x \in e$ to be the B -closest point on s_i to x . Denote $\phi_i(p)$ by p_i . Plainly ϕ_1 is continuous and thus $\phi_1(C_k), k=1, 2, 3$, are connected portions on the boundary of s_1 whose closure contains p_1 . The star-shapedness of $C_B(s_1)$ and the fact that the B -diagram is one-dimensional easily imply that p_1 is an endpoint of all three projections $\phi_1(C_k)$ and that these projections are all disjoint. Thus one of them, say $\phi_1(C_1)$, must be the singleton point p_1 . But then property (2b) implies that C_1 must be a line segment all of whose points are equally B -closest to p_1 and to s_2 . Arguing as in the proof of (2), it follows that $p_1 p_2$ must be parallel to a side of B ; thus our assumptions on general position of B and S imply that p_2 must be a corner of s_2 (see Fig. 2.4). However, a symmetrical argument, using ϕ_2 instead of ϕ_1 , would imply that p_1 must be a corner of s_1 , thereby contradicting our assumptions of general position, namely that no side of B is parallel to the line joining two corners of obstacles in S . \square




Fig. 2.4

(5) Let a be a point on the locus C of points which are equally B -close to two distinct objects s_i, s_j and let y_i (resp. y_j) be the (unique) B -closest point to a on s_i (resp. s_j). Denote the B -distance from a to y_i and y_j by λ . Let T_i (resp. T_j) be a common tangent at y_i (resp. y_j) to s_i (resp. s_j) and $a + \lambda B_0$ and suppose that these tangents are unique. Then the (unique) tangent to C at a is the line T which passes through a and the intersection point z of T_i and T_j .

Proof. Let a' be a point on T sufficiently near a , let y'_i be a point on T_i such that $a'y'_i$ is parallel to ay_i , and let y'_j be a point on T_j defined in an analogous manner. Let

$$\lambda' = \lambda \cdot \frac{|a'y'_i|}{|ay_i|} = \lambda \cdot \frac{|a'y'_j|}{|ay_j|}.$$

Then it is easily seen that $a' + \lambda' B_0$ contains y'_i, y'_j and is fully contained in the wedge between T_i and T_j . Since at least one of the two objects s_i (resp. s_j) and $a + \lambda B_0$ is smooth at their point of contact y_i (resp. y_j), it follows that the B -distances of a' from both s_i and s_j are equal to λ' up to first-order approximation, proving that T is indeed the tangent to C at a . This argument can also be used to show that T is the unique such tangent. \square

- (6) **Corollary.** (a) *If the boundary of B is smooth (i.e., B has a unique tangent at each point on its boundary) then each B -Voronoi edge is smooth.*
 (b) *If B is a polygonal object then each edge of the B -Voronoi diagram is a polygonal arc. Furthermore, if B has k sides then the total number of corners in the edges of $\text{Vor}_B(S)$ is at most $kN + n$.*

Proof. (a) is trivial, because in this case the conditions of (5) hold at each interior point of any of the edges of the diagram.

(b) Let e be a B -Voronoi edge separating $C_B(s_i)$ from $C_B(s_j)$. Clearly, e can be broken into portions at each of which the two contacts between the appropriate expansion of B and s_i, s_j are made between fixed corners or segments of B and fixed segments or corners of s_i and s_j . As long as these type of contact are fixed, it is easily seen that e traces a straight line segment (this follows, e.g., from (5)). Moreover, this straight segment is terminated when contact is made between a corner of B and a corner of either s_i or s_j . It can be shown that if O moves around the boundary of $C_B(s_i)$ with B expanding or shrinking so as to maintain contact with s_i then the total number of points at which a corner of B makes contact with a corner of s_i during that traversal is at most $k + n_i$, where n_i is the number of sides of s_i (see, e.g., [3]). Hence, summing over all B -cells of the diagram, we obtain the asserted bound. \square

Remark. The same arguments used to prove (6b), combined with the observations in [3], can be used to show that, under the assumptions in (6b), the locus of points B -equidistant to two objects s_i, s_j having n_i, n_j sides, respectively, can be calculated in time $O(k + n_i + n_j)$. This observation will be significant in the algorithm for constructing the diagram, given in Section 3 below.




Fig. 2.5

(7) $\text{Vor}_B(S)$ need not be connected. In fact it can have up to $O(n)$ connected components. However, every connected component of $\text{Vor}_B(S)$ is unbounded.

Proof. Since the B -Voronoi diagram for B a disc is the standard Voronoi diagram of S , the first claim follows from the fact that the standard Voronoi diagram may have $O(n)$ connected components (see [12]). The diagram can also become unbounded due to nonsmooth corners of B . Figure 2.5 shows examples of both these types of unboundedness. Let K be a bounded connected component of $\text{Vor}_B(S)$. Since the set of B -closeness cells covers the whole plane there is an object s_i whose cell surrounds K . Obviously s_i lies outside K . Let y be any point in K and let x be the B -closest point to it on s_i . Continue the segment xy past y and away from s_i until it reaches an interior point p in $C_B(s_i)$ (such a point must exist since $C_B(s_i)$ surrounds K). But then by property (2b) the closest point to p is x and this contradicts the fact that $C_B(s_i)$ is star-shaped. \square

(8) **Corollary.** There are no isolated points in $\text{Vor}_B(S)$.

(9) Let p be a vertex of $\text{Vor}_B(S)$ on the boundary of three cells $C_B(s_i)$, $C_B(s_j)$, $C_B(s_k)$. Then every neighborhood of p intersects the interiors of all three cells.

Proof. Suppose to the contrary that there is a point p , whose B -distances to s_i, s_j, s_k are all equal to some λ , such that there exists a small neighborhood N of p which contains no interior point of $C_B(s_i)$. We may assume that N contains no point of a fourth cell, since if this were the case for all sufficiently small such neighborhoods N , then p would be equally B -close to four objects and we have ruled out such cases by our assumptions on general position. Let the B -closest point to p on s_i be y_i and consider the intersection I of the line segment py_i with N . By property (2) py_i is contained in $C_B(s_i)$ and therefore it cannot intersect interior points of any other cell. Also, by our assumption, there are no interior points of $C_B(s_i)$ on either side of a portion of I near p , and therefore that portion of I lies entirely in the interior of the union of $C_B(s_j)$ and $C_B(s_k)$. But (as in the

two such positions, as follows from the argument given in [5]. Thus the properties of B -Voronoi diagrams provide a different proof of this property. (A more general proof which uses a somewhat different technique has recently been given in [9] and [4].)

3. Efficient Construction of the B -Voronoi Diagram

In this section we present an efficient algorithm for constructing $\text{Vor}_B(S)$ for simple convex objects. By simple objects we mean objects for which certain operations, which are used during the construction of the B -Voronoi diagram as described below, can be accomplished in constant time. Such operations are: (a) computing the B -distance between any two points in V or from a given point in V to a straight line; (b) calculating the intersection points of the locus of points which are B -equidistant to two sides or corners of objects in S and another such curve or a straight line segment. For more complex objects the actual complexity of such operations should be taken into account in estimating the complexity of the algorithm.

As will be seen, the properties of $\text{Vor}_B(S)$ which have been proved in the previous section, and which are quite similar to the properties of more conventional Voronoi diagrams, enable us to construct $\text{Vor}_B(S)$ by much the same technique as that described by Yap [14] for the construction of (standard) Voronoi diagrams for a set of simple curves.

Our algorithm proceeds as follows. Let P be the set of the $2N$ leftmost and rightmost points on objects in S (recall that these points are unique by assumption (a) of general position). Following [14] we divide the plane by vertical lines into $2N$ slabs such that (the interior of) each slab contains exactly one point of P . The algorithm runs in stages. In each stage it partitions the set of slabs into disjoint pairs of adjacent slabs and combines each such pair into a new, larger, slab. That is, in the first stage the $2N$ slabs are combined in disjoint pairs into N new slabs, in the second stage these N slabs are combined in pairs into $N/2$ slabs and so on. After $O(\log N)$ stages only a single slab remains and the algorithm terminates.

With each slab the algorithm associates a partial B -Voronoi diagram for an appropriate set of objects which intersect that slab. These diagrams are merged together whenever two adjacent slabs are combined to form a new slab. However, to keep down the cost of creating and merging those diagrams, we will maintain for each slab a partial diagram which covers only those portions of the slab that contain points in P (see below for a more detailed description of this process; note that slab portions for which no diagram is yet maintained are simple "corridors" extending through the entire width of a slab and containing no inner obstacles).

The general merging step of the algorithm thus combines partial B -Voronoi diagrams of two adjacent slabs into a (possibly partial) B -Voronoi diagram of the combined slab. This is done as follows. Consider a slab Q divided by a




Fig. 3.1

vertical line into two slabs; a left slab Q_L and a right slab Q_R . The intersection of the free space V with Q is divided into several components by the boundary of objects which cross Q through its entire width from left to right (i.e., objects whose leftmost points lie to the left of Q and whose rightmost points lie to the right of Q ; see Fig. 3.1). Each such component is called a *window*; note that our definition of windows is somewhat similar to that of *quads* in [14].

With each window W we associate a B -Voronoi diagram, defined over the entire plane, which is induced by the portions of objects contained in W , as follows. For each object s which intersects (the interior of) W , let s^W denote that intersection. In addition let U_W (resp. D_W) denote the upper (resp. lower) portion of the slab Q bounded below (resp. above) by the portion of the object which forms the upper (resp. lower) boundary of W . Then the B -Voronoi diagram of W is the B -Voronoi diagram for the set $S(W)$ consisting of U_W , D_W , and all the nonempty s^W , for $s \in S$ (see Fig. 3.2). Note that U_W , D_W may not be convex, in which case they are split into convex subparts at their concave corners, each of which is regarded as a corner of P (see the comment at the beginning of Section 2). It is clear that the B -Voronoi diagrams of different windows in the same slab Q can overlap one another only outside Q (in fact, as will be noted




Fig. 3.2

below, such overlapping is always of a very trivial nature), and our algorithm will therefore calculate and maintain each of these partial diagrams separately. As already noted, we will maintain B -Voronoi diagrams only for those windows which contain at least one of the pints in P . This ensures that at each stage of the algorithm there are at most $2N$ windows for which the diagram need be constructed (a fact which is crucial for the linearity of each stage), and also, of course, that upon termination the diagram of the last remaining slab will indeed be the complete B -Voronoi diagram of S .

The main step that needs to be described is thus the step which computes the B -Voronoi diagram of a given window W of Q . W is divided by the vertical line separating Q_L , Q_R into two sets of windows, those in the slab Q_L (which we call the L -windows) and those in the slab Q_R (which we call the R -windows). Let L (resp. R) be the set of intersections of the objects in S with the L -windows (resp. R -windows). The construction of the B -Voronoi diagram of W is accomplished by the following three steps. First we construct the partial diagram for each L -window (resp. R -window) that contains at most one point of P (and whose diagram has not yet been constructed). Next we combine the B -Voronoi diagrams of all the L -windows (resp. R -windows) into a full diagram $\text{Vor}_B(L)$ (resp. $\text{Vor}_B(R)$) of the left (resp. right) part of W . Finally, we merge the diagrams of the left and right parts of W into the B -Voronoi diagram of W . We now describe each of these steps in detail.

Step 1. Construction of the B -Voronoi Diagram of an L -Window or an R -Window

If the B -Voronoi diagram of an L -window or an R -window has not yet been constructed, then either we are in the first (innermost) stage of the algorithm and the window W^* in question contains just one point in P , or W^* contains no point of P . Therefore at most three objects s_1, s_2, s_3 can intersect W^* in one of the following situations; either (a) U_{W^*} (resp. D_{W^*}) consists of (at most) two objects (it may consist of two objects if it is a part of an original nonconvex object which has been split into convex subparts as discussed at the beginning of Section 2), D_{W^*} (resp. U_{W^*}) consists of one object and there is no interior object in W^* , or (b) each of U_{W^*}, D_{W^*} consists of one object and there is one interior object in W^* . Hence we can compute the B -Voronoi diagram of W^* by first computing the locus of points which are equally B -close to s_1, s_2 (by Corollary (6b) and the remark following it can be done in time $O(m_1 + m_2)$ if we assume B to be simple or $O(k + m_1 + m_2)$ if we assume B to be a k -gon, where m_1 (resp. m_2) is the number of segments and points composing s_1 (resp. s_2)), and then by adding to the diagram s_3 . By a careful implementation, which resembles the contour tracing procedure which is outlined below, this stage can be done in time $O(m_1 + m_2 + m_3)$ (or $O(k + m_1 + m_2 + m_3)$) where m_3 is the number of segments composing the third object s_3 (see also [8] for a description of the contour tracing technique). Thus the construction of the B -Voronoi diagram of W^* requires $O(k + m)$ time, where m is the number of corners of objects in W^* .

**Step 2. Vertical Merge of the Diagrams of the L -Windows
(resp. R -Windows)**

Let X, Y be two adjacent L -windows which are separated by a portion of an object A and assume that X lies above Y . We claim that the only cells of the B -diagrams of X, Y which can overlap are a cell of the upper portion of Y and a cell of the lower portion of X . This follows from the following argument. First observe that these two diagrams can overlap only outside Q_L . Let z be a point lying (outside Q_L) in a cell $C_B(s_1)$ of $\text{Vor}_B(L)$ where s_1 is in X and is not contained in D_X . Then the object B -closest to z in Y must be (part of) U_Y . Indeed, suppose to the contrary that this object s_2 is not (part of) U_Y . Let $d_B(z, s_1) = \lambda_1$, $d_B(z, s_2) = \lambda_2$, and $p_i \in s_i$ be the B -closest point to z on s_i , $i = 1, 2$. Then it is easily seen that $\lambda_2 \geq \lambda_1$ (since $z \in C_B(s_1)$ in $\text{Vor}_B(L)$) and that both p_1 and p_2 belong to $z + \lambda_2 B_0$. But then the segment $p_1 p_2$ is contained in $z + \lambda_2 B_0$ and it follows that $z + \lambda_2 B_0$ must intersect the *interior* of U_Y (because B_0 has no vertical boundary edge), contradicting the assumption that $d_B(z, s_2) < d_B(z, U_Y)$.

Thus to merge the B -diagrams of two adjacent L -windows X, Y as above, all we have to do is to intersect the cell of D_X in the diagram for X with the cell of U_Y in the diagram of Y to obtain the cell of the object $A = D_X \cap U_Y$ in the combined diagram, and leave all other cells in both diagrams unchanged. This can plainly be done in time which is linear in the number of edges of the two cells $C_B(D_X), C_B(U_Y)$ and thus $\text{Vor}_B(L)$ can be constructed from the partial diagrams of the L -windows in linear time. Similarly, $\text{Vor}_B(R)$ can also be constructed in a completely analogous manner from the partial diagrams of the R -windows in linear time.

Step 3. "Horizontal" merge of $\text{Vor}_B(L), \text{Vor}_B(R)$

The final step merges $\text{Vor}_B(L)$ and $\text{Vor}_B(R)$ to obtain the B -Voronoi diagram of $S(W)$. Note that the objects in $S(W)$ which intersect both Q_L and Q_R are split by the vertical line separating Q_L and Q_R into two parts, one in L and the other in R . The present merge first calculates the B -diagram for $L \cup R$ (i.e., with each such object remaining split), and only then patches up the diagram to obtain $\text{Vor}_B(S(W))$.

To merge the diagrams $\text{Vor}_B(L), \text{Vor}_B(R)$, we have to compute the locus C of points which are equally B -closest to an object in L and to an object in R . Following Kirkpatrick [6] we call C the *contour* separating L from R . It is easy to see that C is a simple, topologically closed curve since each vertex of the combined diagram has degree 3 (for details concerning this standard argument see, e.g., [8]). Hence, C partitions the plane into disjoint open connected regions each of which consists either entirely of points which are B -closer to objects in L than to objects in R (in which case we call such a region, as in [8], an L -region) or entirely of points which are B -closer to objects in R than to objects in L (in which case we call such a region an R -region). Since by our construction each object in L lies to the left of each object in R we have the following lemma.




Fig. 3.3. The setting of Lemma 3.1.

Lemma 3.1 (Fig. 3.3). *Let c_1, \dots, c_{2t} be the common endpoints of segments of objects in L and in R on the vertical line which separates Q_L and Q_R , and suppose that these points are given in increasing order of their y -coordinates. Then the contour C consists of t portions C_1, \dots, C_t such that C_i starts at c_{2i-1} and ends at c_{2i} . If the leftmost and rightmost points of B are smooth then C_1, \dots, C_t are all simple connected arcs; otherwise each of these portions C_i is still simple but not necessarily connected. In this case each C_i is the limit (in the Hausdorff topology of closed sets) of a sequence of simple connected arcs connecting the corresponding points c_{2i-1} and c_{2i} ; thus every connected portion of C_i is unbounded in both directions (except at the endpoints c_{2i-1}, c_{2i}).*

Proof. We will first show that if the leftmost and rightmost points of B are both smooth then each pair of objects in L can be connected by an arc which does not intersect the contour and (by a symmetric argument) that each pair of objects in R can also be connected by such an arc. This will imply that if the leftmost and rightmost points of B are both smooth there is only one L -region and only one R -region in Q and hence the contour must have the form asserted by the Lemma, since it obviously passes through the points c_1, \dots, c_{2t} . For objects B whose leftmost or rightmost point is not smooth we will obtain the result by approximating B by the vector sum B_ϵ of B and a disc with a sufficiently small radius $\epsilon > 0$.

Details are as follows. Let $s_i, s_j \in L$ and let y_i (resp. y_j) be the leftmost point on s_i (resp. s_j). For each $\lambda > 0$ let $a_i(\lambda)$ (resp. $a_j(\lambda)$) be the leftmost point whose B -distance to s_i (resp. s_j) is λ . Since we assume that B has no vertical boundary edges, it follows that the locus A_i (resp. A_j) of all points $a_i(\lambda)$ (resp. $a_j(\lambda)$), $\lambda \geq 0$, is an infinite ray emanating leftward from y_i (resp. y_j), and that A_i and A_j are parallel. Let e_λ be the line segment connecting $a_i(\lambda)$ and $a_j(\lambda)$ and e_0 the line

segment connecting y_i and y_j . We claim that if the rightmost point of B is smooth then for a large enough λ every point $p \in e_\lambda$ lies in an L -region, so that s_i and s_j are connected by the polygonal arc $y_i a_i(\lambda) a_j(\lambda) y_j$ which is completely contained in an L -region, as asserted. It is easily checked that for each $\delta > 0$ there exists $\lambda_0(\delta) > 0$ such that for all $\lambda > \lambda_0(\delta)$ and $p \in e_\lambda$, the B -distance of p to y_i, y_j is less than $\lambda + \delta$ (here we use the fact that the rightmost point on B is smooth). Choose δ to be the smallest B -distance from points in e_0 to objects in R . Then the above observation implies that the B -distance from any point in e_λ to s_i, s_j is less than its B -distance to any object in R , as asserted. \square

If the rightmost point on B is not smooth consider B_ϵ as defined above for some $\epsilon > 0$. We have the following lemma.

Lemma 3.2. (a) For each $\epsilon > 0$ and $p, q \in \mathbf{R}^2$ we have

$$d_{B_\epsilon}(p, q) \leq d_B(p, q) \leq (1 + \epsilon'/K) d_{B_\epsilon}(p, q),$$

where $\epsilon' \rightarrow 0$ as $\epsilon \rightarrow 0$ and $K > 0$ is the Euclidean distance from O to the boundary of B .

(b) Let $s_1 \in L, s_2 \in R$ be two obstacles. For each $\epsilon > 0$ let Γ_ϵ be the locus of points x satisfying

$$d_{B_\epsilon}(x, s_1) = d_{B_\epsilon}(x, s_2)$$

and similarly let Γ be the locus of points x satisfying

$$d_B(x, s_1) = d_B(x, s_2).$$

Then Γ is the limit set (in the Hausdorff topology of closed sets) of Γ_ϵ as $\epsilon \rightarrow 0$ (Fig. 3.4).




Fig. 3.4

Proof. (a) Since $B_\varepsilon \supseteq B$ the left inequality is obvious. As for the other inequality, it is easy to see that for each $\varepsilon' > 0$ there exists $\delta > 0$ such that for each $\varepsilon < \delta$ we can write the vector from O to any point on the boundary of B_ε as $x + z$, where x and z are collinear, $x \in B_0$, and $\|z\| < \varepsilon'$. Suppose $q \in p + \lambda B_\varepsilon$. Then we can write

$$q - p = \lambda(x + z),$$

where x, z are as above. But then $(K/\varepsilon')z \in B$ and we can write

$$\lambda(x + z) = \lambda(1 + \varepsilon'/K)x',$$

where

$$x' = \frac{x + (\varepsilon'/K) \cdot (K/\varepsilon')z}{1 + \varepsilon'/K} \in B.$$

Hence $d_B(p, q) \leq \lambda(1 + \varepsilon'/K)$ and (a) is thus established.

(b) First we show that if $\varepsilon_n \rightarrow 0$ and $x_n \in \Gamma_{\varepsilon_n}$ converge to some x , then $x \in \Gamma$. We have

$$d_{B_{\varepsilon_n}}(x_n, s_1) = d_{B_{\varepsilon_n}}(x_n, s_2).$$

By (a) we then have, for an appropriate corresponding sequence $\varepsilon'_n \rightarrow 0$,

$$(1 + \varepsilon'_n/K)^{-1} d_B(x_n, s_1) \leq d_{B_{\varepsilon_n}}(x_n, s_1) = d_{B_{\varepsilon_n}}(x_n, s_2) \leq d_B(x_n, s_2)$$

so that by continuity of d_B

$$d_B(x, s_1) \leq d_B(x, s_2)$$

and by symmetry we obtain equality, i.e., $x \in \Gamma$.

Next let $x \in \Gamma$, that is $d_B(x, s_1) = d_B(x, s_2)$, and let $\delta > 0$. By properties (4) and (9) of Section 2 it follows that the disc D_δ of radius δ about x must contain two points y, z such that

$$d_B(y, s_1) > d_B(y, s_2)$$

and

$$d_B(z, s_1) < d_B(z, s_2).$$

Hence if ε is sufficiently small we will also have

$$d_{B_\varepsilon}(y, s_1) - d_{B_\varepsilon}(y, s_2) > 0 > d_{B_\varepsilon}(z, s_1) - d_{B_\varepsilon}(z, s_2).$$

This and the continuity of d_{B_ε} implies that $D_\delta \cap \Gamma_\varepsilon \neq \emptyset$ and this proves (b). \square

Lemma 3.2 now implies that the contour C is the limit set of the contours C_ε , defined in a similar manner for B_ε , as $\varepsilon \rightarrow 0$. This, however, shows that for every sufficiently large disc D about the origin $C_\varepsilon \cap D \rightarrow C \cap D$ as $\varepsilon \rightarrow 0$ and thus C has the form asserted by the lemma. \square

Using Lemma 3.1 we can compute the contour in the following way. Following Kirkpatrick [6] we divide each cell of $\text{Vor}_B(L), \text{Vor}_B(R)$ into subcells by adding line segments which join each vertex ν of the diagram to the three points on the three corresponding objects in L or in R to which ν is B -closest. We further refine the subcells by adding new vertices to $\text{Vor}_B(L)$ (resp. $\text{Vor}_B(R)$) and joining these vertices by straight line segments to their B -closest points on L (resp. R) so that each subcell will contain points which are B -closest to just one obstacle segment or corner. These line segments are called (as in [6]) *spokes*. The boundary of each resulting subcell thus consists of two spokes, an edge on the B -Voronoi diagram (or possibly two such edges if the subcell is unbounded), and a (possibly empty) subsegment of an obstacle edge. Clearly, the number of subcells and spokes is linear in the number of vertices of the B -Voronoi diagram and the number of obstacle corners (i.e., $O(N+n)$).

The contour is now traced, starting at the point c_1 defined in Lemma 3.1, in the following way, which resembles similar contour-tracing techniques as described in [6], [12], and [8]. Assume that we are at a point z on the contour which is an interior point of a subcell U of $C_B(s_i)$ in $\text{Vor}_B(L)$ and of a subcell V of $C_B(s_j)$ in $\text{Vor}_B(R)$, and let w_i (resp. w_j) be the side (or corner) on s_i (resp. s_j) to which z is B -closest (Fig. 3.5(a)). We trace the contour portion containing z , which is the locus of points which are equally B -close to w_i and w_j , until its first intersection q with the boundary of either U or V (note that the way in which we have partitioned the diagram into subcells and the assumptions that we have made on the shape of B enable us to find this intersection point in constant time). Suppose that the contour intersects the boundary of V first and enters into a new subcell V' of, say, $C_B(s_k)$ in $\text{Vor}_B(R)$ and let w_k be the side (or corner) to which all points in V' are B -closest among all objects in R . Then after the intersection q the contour continues along a new subarc consisting of points equally B -close to w_i and w_k . In some cases (but only if the leftmost or rightmost point of B is not smooth) both subcells U and V may be unbounded and the contour portion being traced may extend to infinity (Fig. 3.5(b)). In such cases we know from Lemmas 3.1 and 3.2 that the contour “returns from infinity” from the same direction and we can proceed as follows. Suppose that the contour portion being traced extends to infinity to the left (resp. right) of the vertical line separating L and R . We then replace the subcell U (resp. V) by the unbounded subcell U' (resp. V') which is adjacent and lies clockwise (resp. counterclockwise) to it in $\text{Vor}_B(L)$ (resp. $\text{Vor}_B(R)$) and proceed as before to obtain the next (also unbounded) contour portion.

We continue tracing the contour in this manner until the point c_2 defined in Lemma 3.1 is reached. In the same manner we trace all the other portions of the contour. To complete the construction of the diagram we discard those portions




Fig. 3.5. Contour-tracing procedure.

of the contour which form the boundary of objects which were split by the vertical line separating L and R . The complexity of this tracing procedure is plainly linear in the number of intersections of the contour with edges and spokes of both $Vor_B(L)$, $Vor_B(R)$. The following Lemma proves that the number of such intersection points is linear in the number of objects in L and R .

Lemma 3.3. (a) *The contour may either intersect a spoke of $Vor_B(L)$ or of $Vor_B(R)$ once or contain just one subsegment of such a spoke.*

(b) *The contour intersects edges of $Vor_B(L)$ or $Vor_B(R)$ at most $O(m)$ times where m is the number of objects in L and R .*

Proof. (a) Follows from the fact that each cell of $Vor_B(S(W))$ is star-shaped.

(b) Follows from the fact that each intersection of the contour with an edge of $Vor_B(L)$ or of $Vor_B(R)$ is a vertex of $Vor_B(S(W))$ and there are only $O(m)$ such vertices. □

These observations yield the following main theorem.

Theorem 3.1. *The B -Voronoi diagram for a set S of N polygonal convex obstacles consisting altogether of n sides can be computed in time $O(n \log N)$.*

Proof. The algorithm consists of $O(\log N)$ stages. At each stage we compute B -Voronoi diagrams only for windows which contain at least one point of P and hence there are only $O(N)$ such “active” windows at each stage. The number of objects in each window W (recall that these objects are either U_w, D_w , or intersections of objects in S with W) is at most two plus the number of points of P within W . Summing over all windows for which B -Voronoi diagrams are being computed in the current stage, we see that the total number of objects appearing in those windows is at most the number of points in P plus twice the number of active windows. Since there are only $O(N)$ such active windows we conclude that the number of “active” objects at each stage of the algorithm is $O(N)$. Therefore the total number of wall segments and corners in the “active” windows at each stage is $O(N + n)$ and hence from Lemma 3.3 and the argument preceding it, it follows that the computation done at each stage is also $O(N + n)$. \square

Note. We have assumed that the object B is of some simple form so that we can locate the intersections of the contour with the boundary of each subcell in constant time. If, however, B is a polygonal object with k sides then we can divide each cell into a more redefined set of subcells in such a way that the B -Voronoi edge on the boundary of each subcell is a straight line segment. The number of subcells will then be $O(kN + n)$ and the algorithm will work in time $O((kN + n) \log N)$ (see property (6) of Section 2).

4. Translational-Motion Planning Using B -Voronoi Diagrams

In this final section we show how the translational-motion planning problem described in Section 1 can be solved using B -Voronoi diagrams. Let $S = \{s_1, \dots, s_N\}$ be the set of convex obstacles amidst which a given convex body B is allowed to translate, and let $\text{Vor}_B(S)$ be the B -Voronoi diagram of S . In this section we will impose the additional requirement that the space V bounded by the objects in S is a bounded set. This is no real loss of generality because, given S , it is easy to see that there exists a large enough square R (which can be efficiently computed from S) such that a free motion between two positions $p, q \in V \cap R$ of B exists within V if and only if it exists with $V \cap R$. Let us define, for each point p in which (the reference point O on) B can be placed, a set $r(p)$ of points on $\text{Vor}_B(S)$ as follows. Suppose $p \in V$ lies in the interior of some B -cell $C_B(s_i)$. Let x be the B -closest point to p on s_i . Then define $r(p)$ to be the intersection of the ray from x which passes through p with the boundary of the B -cell $C_B(s_i)$. If $p \in \text{Vor}_B(S)$ then we put $r(p) = \{p\}$. By the boundedness of V and by property (2) of Section 2, $r(p)$ is always nonempty, and is either a single point or a straight line segment (see Fig. 4.1(a)). Let $\text{Vor}'_B(S)$ be a subdiagram of $\text{Vor}_B(S)$ which is composed of all points of $\text{Vor}_B(S)$ whose B -closeness to S is greater than one. We can easily compute $\text{Vor}'_B(S)$ from $\text{Vor}_B(S)$ in linear time by considering each subcell of $\text{Vor}_B(S)$ separately. We have the following theorem.




Fig. 4.1

Theorem 4.1. *Let p, q be two positions of B with the same fixed orientation. Then there exists a continuous obstacle-avoiding purely translational motion of B from p to q if and only if $r(p)$ and $r(q)$ belong to the same connected component of $\text{Vor}'_B(S)$.*

Proof. Suppose first that $r(p)$ and $r(q)$ lie in the same connected component of $\text{Vor}'_B(S)$. Then the required motion of B from p to q is obtained by first moving (the reference point O on) B along the straight segment from p to a point in $r(p)$, then to a point in $r(q)$ along a path in $\text{Vor}'_B(S)$ which connects these two sets, and finally along the straight segment from $r(q)$ back to q . (Note that the initial and final portions of that motion are collision free because the B -distance to S keeps increasing as we move from p (resp. q) along the straight segment to $r(p)$ (resp. $r(q)$.) Conversely, let $\nu(t), 0 \leq t \leq 1$, be a continuous path in V that O traces as B is translated from $p = \nu(0)$ to $q = \nu(1)$ without colliding with any obstacle. It is clear that $d_B(\nu(t), S) > 1$ for every t since otherwise B would intersect an obstacle during that motion. By definition of the B -Voronoi diagram we have for each $t, d_B(r(\nu(t)), S) \geq d_B(\nu(t), S)$. Thus the union K of all the sets $r(\nu(t)), 0 \leq t \leq 1$, is contained in $\text{Vor}'_B(S)$. We claim that K is locally connected at a sufficiently small neighborhood of each $r(\nu(t))$ for $\nu(t) \notin \text{Vor}_B(S)$. We can also assume without loss of generality that $\nu(t) \in \text{Vor}_B(S)$ only finitely many times, and we will show how to complete K for each such t to make it fully connected while still being contained in $\text{Vor}'_B(S)$.

To show the above claim it suffices to show that if $p_n \rightarrow p$, with $p \notin \text{Vor}_B(S)$, then there exists a sequence $w_n \in r(p_n)$ which converges to some point in $r(p)$ (Fig. 4.1(a)). Let $p_n \rightarrow p$ and for each $n \geq 1$ let $w_n \in r(p_n)$ be an arbitrary point such that the sequence w_n converges to some point w . By passing to subsequences if necessary, we may assume that there exist two objects s_1, s_2 in S such that for each $n \geq 1, p, p_n \in C_B(s_1), d_B(w_n, s_1) = d_B(w_n, s_2) \leq d_B(w_n, s)$ for $s \neq s_1, s_2$, the straight segments $x_n p_n w_n$ are all contained in $C_B(s_1)$, where x_n is the point on s_1 which is B -closest to p_n (hence also to w_n ; see property (2) of Section 2), and x_n converge to some $x \in s_1$. By continuity these properties will also hold for x ,

p , and w . But then x must be the B -closest point to p and to w on s_1 , w must belong to $\text{Vor}_B(S)$, and the intersection of the ray from x through p with the boundary of $C_B(s_1)$ must contain w . By definition $w \in r(p)$ and the above claim is thus established. If $p \in \text{Vor}_B(S)$ (Fig. 4.1(b)), then the above arguments remain the same, except that w may not belong to $r(p)$. However, since both p and w belong to $C_B(s_1)$ and p lies on the segment from x to w , it follows from property (2) of Section 2 that the segment pw is contained in $\text{Vor}_B(S)$ (actually in $\text{Vor}'_B(S)$ as is easily checked). Hence if we add to K a finite number of such straight segments for points at which $\nu(t)$ intersects $\text{Vor}_B(S)$, we obtain a connected subset of $\text{Vor}'_B(S)$ containing both $r(p)$ and $r(q)$. This completes the proof of the theorem. \square

Remarks. (1) Note that $\text{Vor}'_B(S)$ may have vertices with only one incident edge; these are points p whose B -distance to S is exactly 1. In the terminology of [5], these points are precisely the points of local nonconvexity on the boundary of $K = \bigcup_{i=1}^N (s_i - B)$. Hence the B -Voronoi diagram provides us with an alternative (and more efficient) method of calculating these locally nonconvex corners of K , and thus leads to a more efficient calculation of K itself (see [5] and [4] for more detail).

(2) Another useful property of $\text{Vor}_B(S)$ is that it can be used for planning translational collision-free motion of any scalar multiple λB of B . For each given λ we only need to perform a linear time preprocessing of $\text{Vor}_B(S)$ to produce a truncated subdiagram similar to $\text{Vor}'_B(S)$ and then plan each required motion by an appropriate (linear-time) graph searching through that subdiagram (see [10] for a similar remark concerning motion planning for circles of arbitrary size).

References

1. L. P. Chew and R. L. Drysdale, Voronoi diagrams based on convex distance functions, *Proceedings of the ACM Symposium on Computational Geometry*, 235–244, 1985.
2. S. Fortune, A fast algorithm for polygon containment by translation, *Proceedings of the 12th International Colloquium on Automata, Language and Programming*, 189–198, 1985.
3. L. Guibas, L. Ramshaw, and J. Stofi, A kinetic approach to computational geometry, *Proceedings of the 24th IEEE Symposium on Foundations of Computer Science*, 100–111, 1983.
4. K. Kedem, R. Livne, J. Pach, and M. Sharir, On the union of Jordan regions, and collision-free translational motion amidst polygonal obstacles, *Discrete Comput. Geom.* 1 (1986), 59–71.
5. K. Kedem and M. Sharir, An efficient algorithm for planning collision-free translational motion of a convex polygonal object in 2-dimensional space amidst polygonal obstacles, *Proceedings of the ACM Symposium on Computational Geometry*, 75–80, 1985.
6. D. Kirkpatrick, Efficient computation of continuous skeletons, *Proceedings of the 20th Symposium on Foundations of Computer Science*, 18–27, 1979.
7. S. R. Lay, *Convex Sets and Their Applications*, Wiley, New York, 1972.
8. D. Leven and M. Sharir, Intersection and proximity problems and Voronoi diagrams, in *Advances in Robotics*, Vol. I (J. Schwartz and C. Yap, eds.), 1986.
9. R. Livne and M. Sharir, On intersection of planar Jordan curves, Technical Report 153, Computer Science Department, Courant Institute, New York University, 1985.
10. C. Ó'Dúnlaing and C. K. Yap, A “retraction” method for planning the motion of a disk, *J. Algorithms* 6 (1985) 104–111.

11. M. I. Shamos, Computational geometry, Ph.D. dissertation, Yale University, New Haven, CT, 1978.
12. M. Sharir, Intersections and closest pair problems for a set of circular discs, *SIAM J. Comput.* **14** (1985), 448-468.
13. J. T. Schwartz and M. Sharir, On the piano movers' problem: I. The case of a two-dimensional rigid polygonal body moving amidst polygonal barriers, *Comm. Pure Appl. Math.* **36** (1983), 345-398.
14. C. K. Yap, An $O(n \log n)$ algorithm for the Voronoi diagram of a set of simple curve segments, Technical Report 161, Computer Science Department, Courant Institute, New York University, 1985.

Received June 12, 1985, and in revised form June 19, 1986.