Guaranteed Voronoi Diagrams of Uncertain Sites

William Evans* and Jeff Sember†

Abstract

In this paper we investigate the Voronoi diagram that is induced by a set of sites in the plane, where each site's precise location is uncertain but is known to be within a particular region, and the cells of this diagram contain those points guaranteed to be closest to a particular site. We then examine the diagram for sites with disc-shaped regions of uncertainty, prove that it has linear complexity, and provide an optimal $O(n \log n)$ algorithm for its construction. We also examine the diagram for polygonal regions of uncertainty, and prove that it has linear complexity as well. We then describe a generalization of these diagrams, in which each Voronoi cell is associated with a subset of the sites, and each point in a cell is guaranteed to be closest to some site in the subset associated with the cell.

1 Introduction

Suppose we do not know the precise locations of n sites (n points in the plane) and yet we would like to determine, for every point in the plane, the closest site to that point. If we know the approximate location of each site, say, that the ith site lies in a subset D_i of the plane, then we might be able to answer this question perhaps not for every point but for many points in the plane. Our goal is to find, for each site i, the set of points that are guaranteed to be closer to that site than to any other. In other words, no matter where each site lies (as long as the jth site is in D_j for every j) the closest site to the point is always site i. For some points, we cannot guarantee a closest site. These points form a subset of the plane that we call the 'neutral zone'.

In this paper, we first formally define the partition of the plane into cells of guaranteed closest points and the neutral zone and state some properties of this partition. We then consider the special case when the uncertain regions (i.e. the subsets D_i) are discs and show that the complexity of the partition in this case is linear in the number, n, of sites, and that it can be calculated in $O(n \log n)$ time.

We also consider the case where each D_i is a polygon, and show that the complexity of the resulting partition is linear in the total number of polygon edges.

We then consider a finer partition of the neutral zone into regions of points that we can guarantee are closest to some site in a set of sites. For example, points that may be closest to sites 1 or 2 form the region for the set $\{1, 2\}$. We show that the complexity of this finer partition is at most $O(n^4)$ for uncertain discs.

2 Related work

Voronoi diagrams are a fundamental data structure in computational geometry; see [2] for a survey. Voronoi diagrams involving uncertain sites were investigated with respect to the probabilistic concepts of *expected closest site* and *probably closest site* in [3].

The guaranteed Voronoi diagram of a set of uncertain regions is closely related to the standard Voronoi diagram of those regions. Thus our results rely heavily on properties of standard Voronoi diagrams such as diagrams for circles [7] and diagrams for segments [5].

One of the biggest differences between the guaranteed Voronoi diagram and traditional variants of Voronoi diagrams is that the union of the regions associated with uncertain sites does not cover the plane. The guaranteed Voronoi diagram contains a neutral region that contains those points that are not guaranteed to be closest to any particular site. Zone diagrams also have this property. In zone diagrams, for a point to be in a site's region, it must be closer to the site than to any point in any other site's region. The recursive nature of this definition raises the question of the uniqueness and existence of zone diagrams; a question that Asano et al. [1] answered (positively).

3 Properties

We are given a set of regions in the plane $\mathcal{D} = \{D_1, \ldots, D_n\}$, called *uncertain regions*, each containing a site. Let $H(i, j)$ be the set of points in the plane that are at least as close to site i as site j. That is,

$$H(i, j) = \{p \mid \forall x \in D_i \ \forall y \in D_j \ d(p, x) \leq d(p, y)\}$$

where $d(\cdot, \cdot)$ is Euclidean distance. We denote the boundary of $H(i, j)$ by $\langle i, j \rangle$; formally,

$$\langle i, j \rangle = \{p \mid \max_{x \in D_i} d(p, x) = \min_{y \in D_j} d(p, y)\} \ .$$

*UBC Computer Science,Vancouver, B.C., Canada, V6T 1Z4; will@cs.ubc.ca

†UBC Computer Science,Vancouver, B.C., Canada, V6T 1Z4; jpsember@cs.ubc.ca; research is supported by NSERC

The *cell* for site i, denoted $R_{\langle i \rangle}$, is

$$R_{\langle i \rangle} = \bigcap_{j \neq i} H(i,j) \ . \tag{1}$$

An *edge of* $\langle i,j \rangle$ in $V(\mathcal{D})$ is a maximal connected set of points $p \in \langle i,j \rangle$ that lie on the boundary of cell $R_{\langle i \rangle}$.

The boundaries of all such cells $R_{\langle i \rangle}$ form the *guaranteed Voronoi diagram* for the set $\mathcal{D}$, and we denote it by $V(\mathcal{D})$. Some properties of $V(\mathcal{D})$ are easy to show.

If every uncertain region is a single point, $V(\mathcal{D})$ is the standard nearest-point Voronoi diagram for the regions, which we denote by $V^\circ(\mathcal{D})$. In general, for arbitrary uncertain regions, every cell $R_{\langle i \rangle}$ of $V(\mathcal{D})$ is a subset of the corresponding cell $R_{\langle\langle i \rangle\rangle}$ of $V^\circ(\mathcal{D})$.

It is possible for a cell boundary to not be a one-dimensional curve. Consider $D_i = \{(x,0) \mid x \in [0,2]\}$ and $D_j = \{(x,0) \mid x \in [2,4]\}$. In this case, $\langle i,j \rangle$ is the halfplane $\{(x,y) \mid x \leq 1\}$, and $R_{\langle i \rangle} = \langle i,j \rangle$. To generalize, if D_j intersects CH_i, the convex hull of D_i, then $H(i,j) = \langle i,j \rangle$; and if this intersection is not confined to vertices of CH_i, then $H(i,j) = \langle i,j \rangle = \emptyset$. From this point on, we assume that any nonempty intersection of two regions D_i and D_j is not confined to vertices of CH_i.

A site whose cell is empty can still influence the cell of another site. For example, if the interiors of D_i and D_j intersect, then $R_{\langle i \rangle} = \emptyset$, yet an edge of $\langle k,i \rangle$ for some other site k can still appear in $V(\mathcal{D})$.

A connected subset of the plane S is *inside-tangent* to another such subset C if $S \subseteq C$ and the boundary of C intersects S; and S is *outside-tangent* to C if $S \cap C$ is a non-empty subset of the boundary of C.

Lemma 1 *Every point p on an edge of $\langle i,j \rangle$ in $V(\mathcal{D})$ is the center of a unique disc C_p that has inside-tangent D_i, outside-tangent D_j, and intersects the interior of no $D_k \in \mathcal{D}$ for $k \notin \{i,j\}$.*

Proof. This follows immediately from the definition of an edge of $\langle i,j \rangle$. $\square$

Consider a point p on an edge of $\langle i,j \rangle$ in $V(\mathcal{D})$, and its disc C_p from Lemma 1. Let b be a point of tangency of C_p with D_j. Define $\delta(p)$ to be the (unique) point on segment $\overline{pb}$ that is the center of a disc C° that has outside tangent D_j (at the point b) and outside tangent D_i. Since $C^\circ \subseteq C_p$, C° also intersects the interior of no $D_k \in \mathcal{D}$ for $k \notin \{i,j\}$. Thus $\delta(p)$ lies on an edge of the bisector $\langle\langle i,j \rangle\rangle$ between D_i and D_j that is part of the standard Voronoi diagram $V^\circ(\mathcal{D})$ for the regions $\mathcal{D}$. In fact, $\delta(p)$ is on the boundary of region $R_{\langle\langle i \rangle\rangle}$.

Note that if more than one region D_j is outside-tangent to C_p (or if D_j is tangent to C_p at more than one point), then there is more than one candidate point of tangency b. To make $\delta(p)$ well-defined, we select a b according to some total order on possible b's.

We now show that the ordering of points p on the boundary of a cell $R_{\langle i \rangle}$ in $V(\mathcal{D})$ agrees with that of points $\delta(p)$ on the boundary of $R_{\langle\langle i \rangle\rangle}$ in $V^\circ(\mathcal{D})$. To do this, we will need the following lemma.

Lemma 2 *If b is a point on the boundary of disc P centered at p, and d a point on the boundary of disc Q centered at q, and line segments $\overline{pb}$ and $\overline{qd}$ intersect at a single point, interior to both, then either b is in the interior of Q or d is in the interior of P.*

Proof. Assume such an intersection point w exists. Without loss of generality, assume $d(w,b) \leq d(w,d)$. By the triangle inequality,

$$\begin{aligned} d(q,b) &< d(q,w) + d(w,b) \\ &\leq d(q,w) + d(w,d) \end{aligned}$$

which implies b is in the interior of Q. $\square$

We denote a point p being encountered before point q as we traverse a simple closed curve counter-clockwise (ccw) from starting point s by $p \prec_s q$.

Lemma 3 *If p, q, and s are points on the boundary of cell $R_{\langle i \rangle}$ (with nonempty interior), and $p \prec_s q$, then $\delta(p) \prec_{\delta(s)} \delta(q)$.*

Proof. Each point p on the boundary of cell $R_{\langle i \rangle}$ is mapped to a point $\delta(p)$ on the boundary of cell $R_{\langle\langle i \rangle\rangle}$. Note that segment $\overline{p\delta(p)}$ does not intersect the interior of $R_{\langle i \rangle}$, since for every point p' on this segment, the disc with center p' that has outside-tangent D_j does not contain all of D_i except when $p' = p$. Note also that this disc does intersect D_i (and no other D_k), thus $\overline{p\delta(p)}$ is within $R_{\langle\langle i \rangle\rangle}$. Therefore if $p \prec_s q$ and $\delta(p) \succeq_{\delta(s)} \delta(q)$ then some two of the segments $\{\overline{s\delta(s)}, \overline{p\delta(p)}, \overline{q\delta(q)}\}$ intersect. Without loss of generality, assume $\overline{p\delta(p)}$ intersects $\overline{q\delta(q)}$.

By Lemma 1, disc C_p exists which has outside-tangent some $D_{j \neq i}$ at b, such that $\delta(p) \in \overline{pb}$. Similarly, disc C_q exists which has outside-tangent some $D_{k \neq i}$ at d where $\delta(q) \in \overline{qd}$. This implies $\overline{pb}$ intersects $\overline{qd}$ (since $\overline{p\delta(p)} \subset \overline{pb}$ and $\overline{q\delta(q)} \subset \overline{qd}$). The intersection is a single interior point since $p \notin \overline{qd}$ and $q \notin \overline{pb}$ (otherwise C_p or C_q would not contain D_i), and $\delta(p) \neq b$ and $\delta(q) \neq d$ (otherwise D_i intersects D_j or D_k, and $R_{\langle i \rangle}$ has an empty interior). By Lemma 2, either b is in the interior of C_q or d is in the interior of C_p, which is a contradiction. $\square$

4 Uncertain discs

We now consider the case where the uncertain regions are discs; see Figure 1.

Each disc has a nonnegative radius r_i, and a center S_i. Each $p \in \langle i,j \rangle$ satisfies

$$d(p,S_i) + r_i = d(p,S_j) - r_j \ . \tag{2}$$

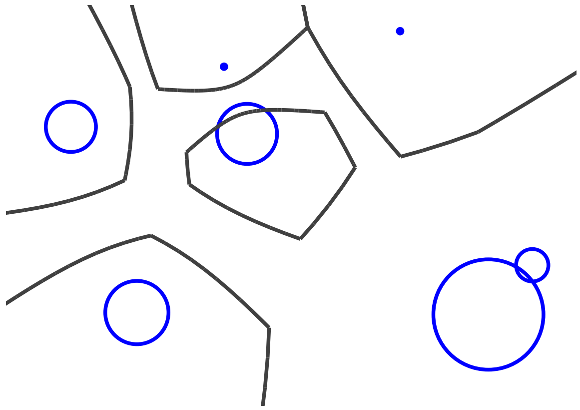

Figure 1: Guaranteed Voronoi diagram

Since r_i and r_j are constants, the points p which satisfy (2) lie on an arm of a hyperbola with foci at S_i and S_j. If the discs' radii are both zero, this is the perpendicular bisector of $\overline{S_i S_j}$; otherwise, it is the hyperbolic arm closest to S_i.

Some properties of $V(\mathcal{D})$ include the following.

Each cell of $V(\mathcal{D})$ is convex, since the cells are intersections of convex halfplanes bounded by hyperbolic arms.

It is possible that more than one edge of $\langle i, j\rangle$ appears in $V(\mathcal{D})$.

We will now show that the number of edges in a guaranteed Voronoi diagram of n discs is $O(n)$. We will do this by showing that for each cell $R_{\langle i\rangle} \in V(\mathcal{D})$, there is a mapping from each edge in $R_{\langle i\rangle}$ to a distinct edge in the corresponding cell of $V^\circ(\mathcal{D})$, which is known to have $O(n)$ edges.

Theorem 4 *The number of edges in a guaranteed Voronoi diagram of n uncertain discs is $O(n)$.*

Proof. We will show that the number of edges in $V(\mathcal{D})$ is at most twice the number of edges in $V^\circ(\mathcal{D})$. The theorem then follows from the fact that $V^\circ(\mathcal{D})$ has $O(n)$ edges (property (7) of [7]).

Consider the edges around $R_{\langle i\rangle}$ in ccw order. We charge each edge E of $\langle i, j\rangle \in V(\mathcal{D})$ to the edge F of $\langle\langle i, j\rangle\rangle$ on which $\delta(p)$ lies, for p the ccw-first point of E (or any interior point p if E is ccw-infinite). Suppose two distinct edges E_1 and E_2 of $\langle i, j\rangle$ map to the same edge F of $\langle\langle i, j\rangle\rangle$ in $R_{\langle\langle i\rangle\rangle}$. Since E_1 and E_2 are distinct but both of $\langle i, j\rangle$, there must exist an edge E' of $\langle i, k\rangle$ $(k \neq j)$ between them in the ccw traversal of $R_{\langle i\rangle}$ that maps to some other edge F' of $\langle\langle i, k\rangle\rangle$ in $R_{\langle\langle i\rangle\rangle}$. This contradicts Lemma 3 since all points of F either precede or follow the points of F' in ccw-order.

Thus each edge in $V^\circ(\mathcal{D})$ of $\langle\langle i, j\rangle\rangle$ is charged at most twice: once by an edge of $\langle i, j\rangle$, and once by an edge of $\langle j, i\rangle$. Hence $V(\mathcal{D})$, like $V^\circ(\mathcal{D})$, has $O(n)$ edges. $\square$

We now show how $V(\mathcal{D})$ for a set of discs can be constructed by first constructing $V^\circ(\mathcal{D})$ for the discs, then performing a linear-time transformation from $V^\circ(\mathcal{D})$ to $V(\mathcal{D})$.

We can construct $I_{\langle i\rangle}$, a sequence of neighboring sites to cell $R_{\langle i\rangle}$, by starting from an edge containing some point p on the boundary of $R_{\langle i\rangle}$ and traversing the boundary edges in ccw order. We construct $I_{\langle\langle i\rangle\rangle}$, the sequence of neighboring sites to cell $R_{\langle\langle i\rangle\rangle}$, by a similar ccw traversal, starting from the edge containing $\delta(p)$.

Lemma 5 *For every cell $R_{\langle i\rangle} \in V(\mathcal{D})$, $I_{\langle i\rangle}$ is a subsequence of $I_{\langle\langle i\rangle\rangle}$.*

Proof. If site j is in $I_{\langle i\rangle}$ then, since $\delta(\cdot)$ maps points on edges of $\langle i, j\rangle$ to points on edges of $\langle\langle i, j\rangle\rangle$, j is in $I_{\langle\langle i\rangle\rangle}$. Furthermore, the order of sites in $I_{\langle i\rangle}$ is preserved in $I_{\langle\langle i\rangle\rangle}$ since $\delta(\cdot)$ preserves this order by Lemma 3. $\square$

Theorem 6 *$V(\mathcal{D})$ for n sites can be constructed in $O(n \log n)$ time, and this running time is optimal.*

Proof. The running time of any algorithm to construct $V(\mathcal{D})$ is $\Omega(n \log n)$, since if the site radii are all zero, $V(\mathcal{D})$ is the standard Voronoi diagram of n points.

Constructing $V^\circ(\mathcal{D})$ for the disc sites $\mathcal{D}$ takes $O(n \log n)$ time [4]. We generate the sequence $I_{\langle\langle i\rangle\rangle}$ of sites comprising the boundary of cell $R_{\langle\langle i\rangle\rangle}$ in $V^\circ(\mathcal{D})$ for $i = 1, 2, \ldots, n$ from this diagram in linear time by a simple traversal. From $I_{\langle\langle i\rangle\rangle}$ we construct the boundary of $R_{\langle i\rangle}$ by generating and intersecting the sequence of hyperbolic arcs it specifies. Lemma 5 ensures that we consider a correctly ordered super-sequence of the arcs bounding $R_{\langle i\rangle}$. This suffices to construct the boundary of $R_{\langle i\rangle}$ in time proportional to the length of $I_{\langle\langle i\rangle\rangle}$.

Since each of the $O(n)$ edges of $V^\circ(\mathcal{D})$ appears in two cell boundaries, the running time for the construction of the edges of all cells of $V(\mathcal{D})$ is $O(n)$. The time to construct $V(\mathcal{D})$ is thus dominated by the time to construct $V^\circ(\mathcal{D})$. $\square$

5 Uncertain polygons

We now turn our attention to the case where the region of uncertainty for each site is a polygon. In this case, each $\langle i, j\rangle$ consists of some number of (possibly unbounded) parabolic arcs, each induced by a vertex u of D_i and a vertex[1] or open edge v of D_j. We denote such a parabolic arc by $\langle i^u, j^v\rangle$, and define an edge of $\langle i^u, j^v\rangle$ to be a maximal connected set of points $p \in \langle i^u, j^v\rangle$ that lie on the boundary of cell $R_{\langle i\rangle}$. We define $\langle\langle i^u, j^v\rangle\rangle$ for $V^\circ(\mathcal{D})$ analogously.

Theorem 7 *The number of edges in the guaranteed Voronoi diagram of $\mathcal{D}$, a set of n polygons with m total edges, is $O(m)$.*

[1] In this case, the induced parabola degenerates to a line.

Proof. We show that the number of edges in $V(\mathcal{D})$ is at most twice the number of edges in the standard Voronoi diagram $V^{\triangle}(\mathcal{D})$ of the polygons $\mathcal{D}$ plus twice the complexity of the furthest point Voronoi diagram of the vertices in D_i summed over all i. The theorem then follows from the fact that $V^{\triangle}(\mathcal{D})$ has $O(m)$ complexity [5] and that the total complexity of the furthest point Voronoi diagrams is $O(m)$ [6].

Let E be an edge of $\langle i^u, j^v\rangle$ on the boundary of $R_{\langle i\rangle}$ and let p be an interior point of E. Section 3 showed there must exist a point $\delta(p)$ on an edge of $\langle\!\langle i^w, j^v\rangle\!\rangle$ where w is a vertex or edge of D_i.

Consider the edges around $R_{\langle i\rangle}$ in ccw order. We charge each edge E of $\langle i^u, j^v\rangle$ to the edge F of $\langle\!\langle i^w, j^v\rangle\!\rangle$ on which $\delta(p)$ lies, for p the ccw-first point of E (or any interior point p if E is ccw-infinite). Now it may happen that a consecutive sequence of edges around $R_{\langle i\rangle}$ all map to F. (By Lemma 3, the edges must be consecutive if they map to the same F.) Let E_1 of $\langle i^{u_1}, j^v\rangle$ and E_2 of $\langle i^{u_2}, j^v\rangle$ be two successive (adjacent) edges in this ccw sequence. The point p shared by E_1 and E_2 lies on an edge of the furthest-point Voronoi diagram of the vertices of D_i that separates the furthest-point regions for u_1 and u_2. We charge the edge E_2 to this edge T of the furthest-point Voronoi diagram. We now show that at most two edges are charged to each T. Every such p intersecting T is the center of a disc C_p that has inside-tangent D_i (at the two farthest vertices u_1, u_2 associated with T) and outside-tangent D_j. Assume by way of contradiction that there are three such points, p_1, p_2, p_3 in order along T. Observe that C_{p_2} is contained within $C_{p_1} \cup C_{p_3}$; thus D_j must be outside-tangent to C_{p_2} at either u_1 or u_2 to avoid intersecting the interior of the other two discs. But then $D_i \cap D_j$ is a nonempty subset of $\{u_1, u_2\}$, both vertices of CH_i, a contradiction.

Thus the number of edges on the boundary of region $R_{\langle i\rangle}$ is at most the number of edges on the boundary of region $R_{\langle\!\langle i\rangle\!\rangle}$ plus twice the number of edges in the furthest-point Voronoi diagram for the vertices of D_i. The theorem then follows since each edge of $V^{\triangle}(\mathcal{D})$ bounds two regions $R_{\langle\!\langle i\rangle\!\rangle}$ and $R_{\langle\!\langle j\rangle\!\rangle}$. $\square$

6 Extension to subsets of closest points

In this section, we look at an extension of the Voronoi diagram which assigns every point in the plane to a cell, including points in the neutral zone.

Equation (1) can be generalized so that each point in a cell is guaranteed to be at least as close to a site in a particular subset of sites than to any other site. For a set $S \subseteq \{1 \dots n\}$, we define the cell for S (denoted $R_{\langle S\rangle}$) as

$$R_{\langle S\rangle} = \bigcup_{i\in S}\Big[\bigcap_{j\notin S} H(i,j)\Big] - \bigcup_{S'\subset S} R_{\langle S'\rangle}$$

where $R_{\langle\emptyset\rangle} = \emptyset$. See Figure 2 for an example of such a *guaranteed subset Voronoi diagram*, which we denote by $V^{\{\}}(\mathcal{D})$.

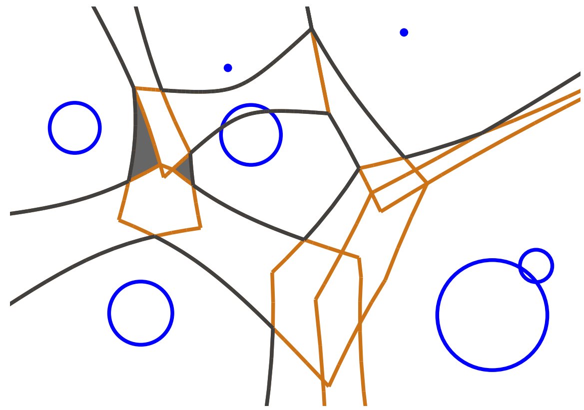

Figure 2: Guaranteed subset Voronoi diagram

The cells of $V^{\{\}}(\mathcal{D})$ are not necessarily connected. In Figure 2, for instance, the two shaded regions belong to the same cell.

Lemma 8 *The number of edges in a guaranteed subset Voronoi diagram of n uncertain discs is $O(n^4)$.*

Proof. The proof follows immediately from the fact that each edge in $V^{\{\}}(\mathcal{D})$ is an edge in the arrangement of the $2 \cdot \binom{n}{2}$ possible hyperbolic arcs $\langle i, j\rangle$. $\square$

References

[1] T. Asano, J. Matoušek, and T. Tokuyama. Zone diagrams: Existence, uniqueness, and algorithmic challenge. *SIAM J. Comput.*, 37(4):1182–1198, 2007.

[2] F. Aurenhammer. Voronoi diagrams—a survey of a fundamental geometric data structure. *ACM Comput. Surv.*, 23(3):345–405, 1991.

[3] F. Aurenhammer, G. Stöckl, and E. Welzl. The post office problem for fuzzy point sets. In *CG '91: Proc. International Workshop on Computational Geometry - Methods, Algorithms and Applications*, pages 1–11, London, UK, 1991. Springer-Verlag.

[4] S. Fortune. A sweepline algorithm for Voronoi diagrams. In *SCG '86: Proc. 2nd Annual Symposium on Computational Geometry*, pages 313–322, New York, NY, USA, 1986. ACM.

[5] D. G. Kirkpatrick. Efficient computation of continuous skeletons. In *Proc. 20th Annual Symposium on Foundations of Computer Science*, pages 18–27, 1979.

[6] D. T. Lee. On k-nearest neighbor Voronoi diagrams in the plane. *IEEE Trans. Comput.*, 31(6):478–487, 1982.

[7] M. Sharir. Intersection and closest-pair problems for a set of planar discs. *SIAM J. Comput.*, 14(2):448–468, 1985.