Optimal Voronoi diagram construction with n convex sites in three dimensions

Paul Harrington*
Imagine International Reinsurance Limited
43 St. Stephen's Green, Dublin 2, Ireland
Colm Ó Dúnlaing
Mathematics, Trinity College, Dublin 2, Ireland
Chee K. Yap
Courant Institute of Mathematical Sciences
251 Mercer Street, N.Y., N.Y. 10012, U.S.A.

December 10, 2004

Abstract

Lee and Drysdale [13] and Sharir [16] have given suboptimal $O(n \log^2 n)$ algorithms for computing the Voronoi diagrams of line-segments or discs in the plane.

This paper adapts their methods to constructing the Voronoi diagram for disjoint convex sites in three dimensions.

Let $C(n)$ be an upper bound on the number of features in any diagram involving n sites in three dimensions. Even for point sites, $C(n)$ is $\Omega(n^2)$, so it is reasonable to assume that $C(n)/n^2$ is monotonically nondecreasing. Under this assumption, our construction with n sites takes optimal time $O(C(n))$.

Results of Aurenhammer's [3] and Chazelle's [9] show that if the sites are spherical then the diagram can be constructed in time $O(n^2)$, but $C(n)$ has yet to be determined for more general classes of site.

The sites are assumed to be compact and semi-algebraic of bounded algebraic complexity. For the sake of simplicity we assume that the sites are strictly convex and rounded (having no sharp corners nor edges). Assuming strict convexity simplifies the analysis, as does the roundedness assumption, though the latter is neither necessary nor restrictive.

1 Introduction

Given a set S of n *sites*, nonempty subsets of $\mathbb{R}^3$, their *Voronoi diagram* is the set of points equidistant from two or more closest sites. This definition makes most sense when the sites are pairwise disjoint

*E-mail paul_j_harrington@hotmail.com, odunlain@maths.tcd.ie, yap@cs.nyu.edu.

A preliminary version of these results was presented at the third Irish conference on the Mathematical Foundations of Computer Science and Information Technology, Dublin, July 2004.

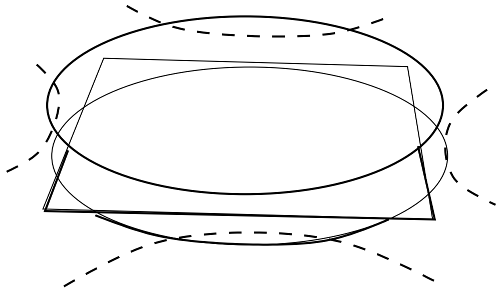

Figure 1: two discs and one regular k-gon, k Voronoi edges (Paragraph 1.1). The edges, which are unbounded, are indicated by broken curves.

closed, bounded, and convex [14]. For a general survey of Voronoi diagrams and their applications, see [4].

In two dimensions the complexity of constructing Voronoi diagrams is well-understood. Of particular interest is calculating the Voronoi diagram of polygonal sites, and optimal $O(n \log n)$ construction times have been achieved [18,11]. Both these algorithms adapt to computing the Voronoi diagram for a set of circular sites in two dimensions.

Naturally, one wants to develop algorithms for constructing Voronoi diagrams in three dimensions. It seems that the simplest case to consider, apart from point-sites which are fully understood, is the case of spherical sites. Here the complexity of the diagram is known to be $O(n^2)$ and can be constructed in time $O(n^2)$, from a reduction to convex polyhedra in $\mathbb{R}^5$ [3], and an optimal algorithm for their construction [9]. So the case of spherical sites is also understood.

However, these reductions for spherical sites can hardly furnish algorithms for polyhedral sites, where the complexity of the Voronoi diagram is as yet not settled. We continue, therefore, to study how more mainstream 2-dimensional methods can be adapted to three dimensions.

We base our algorithm here on Sharir's $O(n \log^2 n)$ method for circular sites in two dimensions [16], which itself is a modification of Lee and Drysdale's method [13].

For the rest of this paper we shall be considering a set S of n disjoint convex sites: $n = |S|$. The sites are assumed to have piecewise algebraic boundaries composed of a bounded number of algebraic pieces each of bounded degree. It follows that the descriptive complexity of the diagram depends only on the number n of sites and may be written as $C(n)$.

(1.1) It is to be remarked that the algebraic complexity of sites does not noticeably affect the number of features in 2-dimensional Voronoi diagrams, which have at most $17n$ features (roughly) no matter how complex the sites may be algebraically. On the other hand, in three dimensions, without the algebraic degree bound, three sites can have an arbitrary number of Voronoi edges. Figure 1 gives an example. A regular k-gon is sandwiched between two discs of equal radius, between the inner and outer radius of the k-gon. There are k unbounded Voronoi edges.

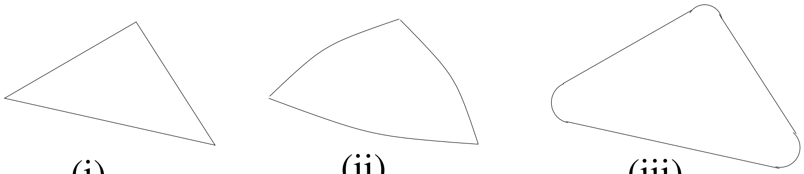

(i) (ii) (iii)

Figure 2: convex sets in 2 dimensions, (i) convex, (ii) strictly convex, and (iii) rounded.

2 Voronoi diagrams

(2.1) Three-dimensional space $\mathbb{R}^3$ is equipped with *Euclidean distance*: if $p = (x_0, y_0, z_0)$ and $q = (x_1, y_1, z_1)$, then

$$d(p,q) = \|q-p\| = \sqrt{(x_1-x_0)^2 + (y_1-y_0)^2 + (z_1-z_0)^2}.$$

If r is a nonnegative real number then $B(p,r)$ will denote the ball centred at p of radius r. It may be open or closed according to context.

$$B(p,r) = \{x \in \mathbb{R}^3 : \|x-p\| < r\} \text{ (open)}, \quad \text{or} \quad B(p,r) = \{x \in \mathbb{R}^3 : \|x-p\| \leq r\} \text{ (closed)}.$$

(2.2) Restriction on sets of sites. Throughout this paper, S will be a set of disjoint, nonempty, closed, bounded, strictly convex, rounded sets (Definition 2.3) in $\mathbb{R}^3$, which we shall call *sites*. It will be assumed that each site is a semi-algebraic set of bounded algebraic complexity.

(2.3) Definition *A set $P \subset \mathbb{R}^3$ is* strictly convex *when for any $x, y \in P$, the closed line-segment xy is contained in P and the open line-segment xy is contained in the interior of P.*

For any set P, its boundary *(those points which are interior neither to P nor its complement) is denoted ∂P.*

A convex set P is rounded *if for every point $p \in \partial P$ there exists exactly one plane through p tangent to P. Intuitively, P has no sharp corners or edges. See Figure 2.*

(2.4) Proposition *If P is bounded,[1] nonempty, closed, and convex, then for every $x \in \mathbb{R}^3$ there exists a unique point y in P closest to x. The point y depends continuously on x, $x = y$ if and only if $x \in P$, and if $x \notin P$ then $y \in \partial P$ (See, for instance,* [14, Lemma 1.1]*).* ∎

(2.5) Definition *Write $\pi_P(x)$ for the point in P closest to x. The map π_P is called* inward projection *onto P. Define $d(x,P) = \|x - \pi_P(x)\|$, the distance from x to P.*

(2.6) It follows from the above proposition that $d(x,P)$ is continuous in x.

(2.7) Definition *For any point $x \in \mathbb{R}^3$, the* clearance sphere *at x is the possibly degenerate sphere whose centre is x and whose radius is the distance from x to the closest site or sites. See Figure 3.*

The Voronoi diagram $\mathrm{Vor}(S)$ *is the set of points x in $\mathbb{R}^3$ such that the clearance sphere at x touches two or more sites in S.*

The open Voronoi cell *of a site P in $\mathrm{Vor}(S)$ consists of all points whose clearance spheres meet P and no other site. The (closed)* Voronoi cell *consists of all points whose clearance spheres meet P.*

[1] P need not be bounded, but in this paper it always is.

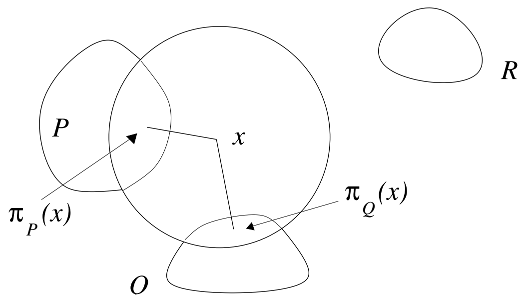

Figure 3: Clearance sphere: x is on $B(P,Q)$, common to the P- and Q-cells in $\mathrm{Vor}(P,Q,R)$, but at positive distance from the R-cell.

(2.8) Degenerate clearance spheres, that is, of zero radius, are possible: x has nondegenerate clearance sphere if and only if $x \in \mathbb{R}^3 \setminus \bigcup S$. Since the sites are disjoint, the Voronoi diagram is contained in $\mathbb{R}^3 \setminus \bigcup S$.

We remark without proof that the 'roundedness' requirement is not restrictive: suppose that for every $P \in S$, $P_\varepsilon = \{x : \ d(x,P) \leq \varepsilon\}$. Let $S_\varepsilon = \{P_\varepsilon : P \in S\}$. If ε is small enough so all sites in S_ε are disjoint, then $\mathrm{Vor}(S)$ and $\mathrm{Vor}(S_\varepsilon)$ coincide and if S satisfies all assumptions (2.2) except not all sites are rounded, then S_ε satisfies all the assumptions 2.2, including the semi-algebraic property, and all sites are rounded.

(2.9) Definition *The* bisector $B(P,Q)$ *separating P from Q consists of all points $x \in \mathbb{R}^3$ such that $d(x,P) = d(x,Q)$.*

(2.10) Lemma *$B(P,Q)$ is unbounded and homeomorphic to a plane. Explicitly, let Π be the plane bisecting the closest pair of points in P and Q: then the orthogonal projection from the bisector onto the plane is a homeomorphism.*

Partial proof. We shall show that the projection is injective.

$B(P,Q)$ is a differentiable surface, and for any $x \in B(P,Q)$, if p and q are the points where its clearance sphere touches p and q respectively, then the tangent plane T at x bisects the line-segment pq, so p and q are reflections of each other in T [14, Lemma 1.3].

Without loss of generality Π is the yz-plane, with P to its left and Q to its right. If orthogonal projection from $B(P,Q)$ onto the yz-plane is not injective, then there exist two points $(\alpha_1, \beta, \gamma)$ and $(\alpha_2, \beta, \gamma)$ in $B(P,Q)$. Let Σ be the plane through these two points and parallel to the z-axis. $B(P,Q) \cap \Sigma$ contains two co-horizontal points, so it contains a local maximum or minimum in the z-direction, call it x. At this point x the tangent line L to $B(P,Q) \cap \Sigma$ is parallel to the x-axis. Let $T \supseteq L$ be the tangent plane to $B(P,Q)$ at x and p, q the points where its clearance sphere touches P and Q: p and q are mirror images in T, so $\vec{pq}$ is perpendicular to T, hence perpendicular to the x axis. This means that p and q have the same x-coordinates, which is impossible. **Q.E.D.**

(2.11) Here we introduce and define the *features* of $\mathrm{Vor}(S)$: its cells, faces, edges, and vertices. Cells have already been defined. For any sites $P, Q \in S$, let U be the set

$$\{x \in B(P,Q) : \ (\forall P' \in S)(P' \neq P \wedge P' \neq Q \Rightarrow (d(x,P) = d(x,Q) < d(x,P')))\}.$$

U is a relatively open subset of $B(P, Q)$. An *open face* (common to the boundary of the P- and Q-cells) is a (path)-connected component of U, and a *face* is the closure of an open face. The *trisector* $T(P, Q, R)$ separating three sites is the set of points equidistant from all three. Let T be the set of points

$$\{x : \ (\forall P' \in S)(P' \neq P \wedge P' \neq Q \wedge P' \neq R \Rightarrow (d(x, P) = d(x, Q) = d(x, R) < d(x, P')))\}.$$

T is a relatively open subset of the trisector $T(P, Q, R)$. An *open edge* is a connected component of T, and an edge is the closure of an open edge. A vertex is a point on the boundary of an open edge. Edges can have zero, one, or two incident vertices. Edges with no incident vertex either extend infinitely in both directions or are bounded closed loops. Edges with one incident vertex extend infinitely in one direction.

The trisector is a 1-dimensional curve, possibly with several connected components (see Figure 1). There is a bound on the number of components since sites are semi-algebraic of bounded algebraic complexity.

(2.12) Assumptions of general position. We make the following assumptions on convexity and general position for the sites in S. In addition to the restrictions stated in Paragraph 2.2,

- The leftmost endpoints of all sites have distinct x-coordinates.
- No four sites have a common tangent plane.
- If P is a site and e an edge on its cell boundary in $\mathrm{Vor}(S)$, then $\pi_P(e)$ intersects every plane orthogonal to the x-axis in a finite number of points.
- If u is a point on ∂P which is either a projected vertex or a stationary point with respect to the x-direction on a projected edge, or a seam edge (see Paragraph 3.10 below), and v is another such point, then u and v have different x-coordinates.

3 Frustums, the convex hull, and the inward projection of cell boundary

Our algorithm is classical divide-and-conquer, partitioning S into two sets, calculating the diagrams for each, and combining the results into one diagram. Later in this section (paragraph 3.15), the output of the algorithm will be specified fully.

The algorithm needs the Voronoi cells to be divided into simple pieces which will be called *frustums*. When combining two Voronoi diagrams, pairs of frustums will be processed. The combining process will be initiated from certain points called *starter points,* and it will be necessary to find the frustums containing the starter points in each diagram. Hence we need a point-location method for frustums.

(3.1) ∂C is not always homeomorphic to a plane or sphere. Let C be the cell containing a site P in $\mathrm{Vor}(S)$. The decomposition of C into frustums could be based on some kind of trapezoidal decomposition of the faces on ∂C. This is adequate if C is bounded, since ∂C is, roughly speaking,

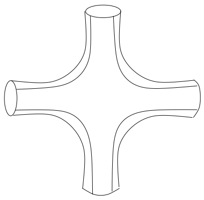

Figure 4: part of an unbounded cell C where ∂C is not homeomorphic to a plane.

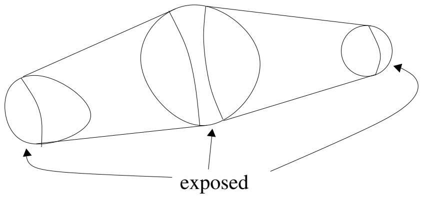

Figure 5: Convex hull, exposed and hidden regions, which are separated by seams.

a planar graph.[2] However, if C is unbounded, then ∂C need not be homeomorphic to a plane nor a sphere (Figure 4, also Figure 1).

For this reason the decomposition begins by projecting the faces, edges, and vertices from ∂C onto ∂P, using the inward projection π_P (2.5). This map is not necessarily surjective: where it isn't, the site meets the boundary of the convex hull. Convex hulls are discussed extensively in [7,17].

(3.2) Definition *The* convex hull $H(S)$ *is the smallest convex set containing* $\bigcup S$*, that is, the intersection of all closed halfspaces containing all the sites in* S*.*

Let P *be a site and* p *a point in its boundary* ∂P*. By definition,* $\partial P \subseteq P \subseteq H(S)$*. Call* p exposed *if it is in the boundary of* $H(S)$ *and* hidden *if it is interior to* $H(S)$*. See Figure 5.*

If the cell C containing a site P is bounded, then $\pi_P : \partial C \to \partial P$ is a homeomorphism. If the cell is unbounded, then $\pi_P : \partial C \to \partial P$ is not surjective, as the following lemma shows.

(3.3) Lemma *Let* P *be a site,* C *the cell containing* P *in* $\mathrm{Vor}(S)$*. Then* (i) *the restriction of* π_P *to the boundary* ∂C *of* C *is injective, and* (ii) $\pi_P(\partial C)$ *consists of the hidden points on* ∂P*.*

Proof. Given $p \in \partial P$, let T be the tangent plane to P at p, and L be the ray extending from p in the direction of the outward unit normal — that is, $L = \pi_P^{-1}(p)$.

For (i), note that if $x, x' \in L$ and $\|x' - p\| > \|x - p\|$ then the closed ball $B(x, \|x - p\|)$ intersects P only at p and is strictly interior to the open ball $B(x', \|x' - p\|)$ except at p. In particular if $x \in \partial C$, so $B(x, \|x - p\|)$ touches another site Q at a point q, then q is interior to $B(x', \|x' - p\|)$ and x', being closer to Q than to P, is not in C. Therefore L can intersect ∂C at at most one point, which proves that the restriction of π_P to ∂C is injective.

(ii) If p is exposed, then the open half-space bounded by T and containing L does not intersect $H(S)$. Therefore, for any point $x \in L$, the closed ball $B(x, \|x - p\|)$, which is interior to this halfspace except at p, intersects $H(S)$ only at p. But $H(S) \supseteq \bigcup S \supseteq P$, so x is closer to P than any other site, and cannot belong to ∂C, so $p \notin \pi_P(\partial C)$.

[2] It may have multiple edges and closed loops without incident vertices.

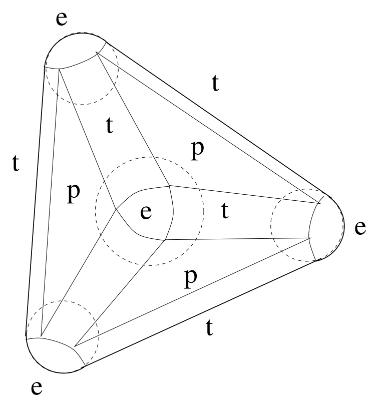

Figure 6: Convex hull of four spheres, forming a rounded tetrahedron. Exposed facets, tunnel facets, and planar facets are marked e, t, and p, respectively.

Conversely, suppose $L \cap \partial C = \emptyset$, so for all $x \in L$ the closest point in $\bigcup S$ is p. For every positive integer n, let x be the point in L at distance n from p. Then x has clearance n, so

$$V_n \supseteq \bigcup S,$$

where $V_n = \{y : \|y - x\| \geq n\}$. Therefore

$$\bigcup S \subseteq \bigcap_n V_n.$$

The set on the right is the closed halfspace bounded by T and containing P, which implies that $p \in \partial H(S)$, so p is a exposed boundary point. **Q.E.D.**

(3.4) If E comprises the exposed and H the hidden points on ∂P, then $E \cup H = \partial P$ and it can be shown that E is relatively closed and H relatively open in ∂P and that $\partial E = \partial H$ is a Jordan curve separating H from the relative interior of E. Points on the common boundary of E and H are called *seam points.*

(3.5) Definition *The* P, Q-seam on P *is the set of seam points on* ∂P *relative to* $H(P, Q)$*, the convex hull of* $\{P, Q\}$ *alone.*

(3.6) Under the assumption that P and Q are strictly convex and rounded, it can be shown that the P, Q-seam is homeomorphic to a circle. (There is also a Q, P-seam on ∂Q.)

(3.7) Structure of $\mathbf{H(S)}$. The boundary of $H(S)$ is composed of *exposed facets* on site boundaries, *tunnel facets*, and *planar facets* (*triangular facets*) meeting three sites. See [17]. Exposed facets are connected components of the exposed region, tunnel facets are surfaces generated by line-segments touching two sites, planar facets are bounded by triangles, and facets meeting four or more sites are ruled out by general position (2.12). Figure 6 illustrates these features.

If F is a planar facet, then the three corners on its boundary are the unique points of contact of F with sites in S. Its sides are shared with three tunnel facets.

(3.8) Lemma *Let* p *be a seam point on* ∂P*. Then there exists a site* Q *such that* $p = \lim_n \pi_P(x_n)$*, where* x_n *all belong to the same unbounded* P, Q*-face in* $\mathrm{Vor}(S)$*. Also,* p *is on the* P, Q*-seam.*

Proof. The point p is in the boundary of the hidden region. Therefore there exists a sequence p_n of point in the hidden region converging to p. By Lemma 3.3, there exist (unique) points $x_n \in \partial C$ such that $p_n = \pi_P(x_n)$. Since there are a finite number of different faces, all the x_n can be assumed to come from the same face of $\mathrm{Vor}(S)$.

If instead of S we turn our attention to $H(P, Q)$ and $\mathrm{Vor}(P, Q)$, the hull and diagram for $\{P, Q\}$, then p is an exposed point in $H(P, Q)$ and also the limit of hidden points relative to $\{P, Q\}$, so it is in the P, Q-seam. **Q.E.D.**

In the above lemma the points x_n do not converge, they tend to infinity. Similarly,

(3.9) Lemma *Let $p \in \partial P$ be a corner of a planar facet meeting three sites P, Q, R. Then p is common to the P, Q- and P, R-seams, and there exists a sequence of points x_n tending to infinity on the P, Q, R-trisector such that $\pi_P(x_n) \to p$. (Proof omitted.)* ∎

(3.10) Boundary of exposed facets in ∂P. By Lemma 3.8, the boundaries of the exposed facets on ∂P are contained in the union of the P, Q-seams. If p is an exposed point where the P, Q- and P, R-seams intersect, then it is one of the corners of a planar facet touching P, Q, and R. Therefore the exposed facet boundaries can be decomposed into edges and vertices, where the edges are part of seams and the vertices are corners of planar facets.

(3.11) Projecting a cell boundary on the site boundary. Let C be the cell containing the site P in $\mathrm{Vor}(S)$.

C is bounded if and only if P contains no exposed points. The boundary of C consists of faces, edges, and vertices, and if C is bounded then it is a generalised planar graph, see Paragraph 3.1. As mentioned at the start of this section, ∂C need not be homeomorphic to a plane if C is unbounded, and for this reason we project ∂C onto ∂P, producing a generalised planar graph (there can be multiple edges and closed-loop edges without incident vertices).

The subdivision of ∂P arising from $\pi_P(\partial C)$ is summarised as follows.

(3.12) Theorem *Let C be the cell containing P in $\mathrm{Vor}(S)$. The restriction of π_P to ∂C is injective, and its image consists of the hidden points in ∂P (Lemma 3.3). If C is unbounded then the remainder of ∂P consists of the exposed points, that is, the union of exposed facets. The boundary of the exposed facets consists of seam edges meeting at the corners of adjacent planar facets (3.10).* ∎

Let G be the (generalised) graph structure on ∂P, subdividing the hidden region by the projected graph structure from ∂C, and whose other faces are the exposed facets, bounded by seam edges and vertices from planar facets.

(3.13) Description of face boundaries. The boundary of a P, Q-face F in $\mathrm{Vor}(S)$ consists of some Jordan chains of edges, together with, if F is unbounded, one or more unbounded chains of edges.

If F is bounded then one of these curves bounds F from the outside, the outer boundary,and the others from the inside. If it is unbounded, it is natural to arrange the outer boundary in cyclic order around F. To do this, recall that $B(P, Q)$ can be projected homeomorphically onto a plane (Lemma 2.10); let F' be the image of F under this projection and intersect it with a large disc D, large enough to contain all bounded features and intersect all unbounded features of F'. Each unbounded component of $\partial F'$ intersects ∂D twice, and the order of these intersection points fixes an anticlockwise order between the unbounded components, and also an anticlockwise order within the unbounded components. This order is independent of D, if D is sufficiently large, and induces a corresponding order on the boundary of F.

Each intersection point on $\partial F' \cap \partial D$ is on an infinite 'branch' of a Voronoi edge bounding F. An unbounded edge with no incident vertex has two branches, an unbounded edge with an incident vertex has one, and it corresponds to a unique seam vertex (Lemma 3.9). A segment of $F' \cap \partial D$ joins two intersection points and corresponds to a P, Q-seam edge (Lemma 3.8). We say that the branches inducing two adjacent intersection points on $F' \cap \partial D$ are *adjacent*.

(3.14) Definition *An infinite branch of F is called a* vertex at infinity, *and where two branches of F are adjacent then they bound an* edge at infinity.

(3.15) Output of the Voronoi diagram computation. At this point we state explicitly what the required output should be. It also represents the input available from the two recursive calls on the two subsets of S. Since S is in general position (2.12), every edge is incident to three faces and every vertex to four cells, six faces, and four edges.

Input is a list of disjoint compact convex sites P, with suitable descriptions as semi-algebraic sets. The output is as follows.

- For each site P, the list of faces incident to its cell in $\mathrm{Vor}(S)$.
- For each face, its boundary components, in cyclic order for the Jordan-curve components, while the unbounded components are presented together in cyclic order as discussed above.
- For each unbounded branch as discussed above,
- For each face, its two incident cells.
- For each edge, its three incident faces and cells and its bounding vertices (at most two).
- For each vertex, its four incident edges, six incident faces, and four incident cells.
- A network representing the infinite branches and their incidence relationships, that is for each branch (vertex at infinity) a record describing the three adjacent edges at infinity, on the three incident faces, and for each edge at infinity, a record describing the two vertices at infinity which it connects.

(3.16) Lemma *For any cell C in $\mathrm{Vor}(S)$, the structure of ∂C can be extracted from the above information in linear time.*

Proof. That is, the faces, edges, and vertices can be copied in linear time, and their incidence relations generated in linear time. Faces, edges, and vertices can be copied through the various lists mentioned above.

For each edge, there are two incident faces on ∂C, and they can be extracted immediately from the three faces incident in $\mathrm{Vor}(S)$. Also, the vertices incident to each edge are available.

For each vertex, there are three incident edges on ∂C, which are easily extracted from the four incident edges in $\mathrm{Vor}(S)$. **Q.E.D.**

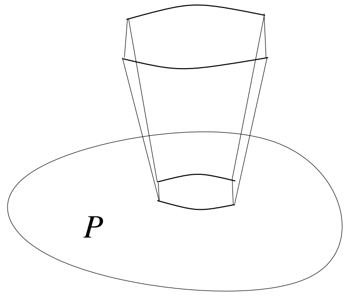

Figure 7: Quad on ∂P, and the frustum extending it.

(3.17) Lemma *For every site P, the generalised graph G can be constructed in linear time.*

Proof. For bounded edges and faces the incidences in G are inherited from ∂C (Lemma 3.16). The vertices on the exposed facets correspond to unbounded edges (Lemma 3.9). The seam edges correspond to unbounded faces (Lemma 3.8). The seam edges and vertices are inherited from the network described in Paragraph 3.15. **Q.E.D.**

(3.18) Definition *With G as introduced, its faces are subdivided into 'quads' as follows.*

Through each vertex $p \in \partial P$, let T be the plane through p parallel to the yz-plane. Introduce two new edges by following the curve $T \cap \partial P$ until it meets nearby edges of G. Do the same with points p which are local maxima or minima of edges of G, in the x-direction.

The resulting planar subdivision of ∂P is denoted $G(P, S)$, called the augmented graph*. A* quad *is a face of $G(P, S)$. Every quad is incident to at most four edges. For every quad D,*

$$C \cap \pi_P^{-1}(D)$$

is called a frustum*. See Figure 7, also Figure 12.*

The frustums provide a decomposition of $C \backslash P$ into simple regions.

It will be necessary to furnish a point-location structure whereby, given a point $x \in C \backslash P$, the correct frustum containing x can be found. Equivalently one needs to find the quad containing $\pi_P(x)$ in ∂P. This is a point-location problem in $G(P, S)$.

To construct the augmented graph $G(P, S)$ from the graph G discussed above, for each face of $G(P, S)$ one needs the vertices (and turning points) sorted with respect to the x-coordinate. This is easily done directly using sorting, but we must avoid the $\log n$ runtime factor needed to sort.

This will be considered in Section 5.

4 Divide-and-conquer algorithm

(4.1) The divide-and-conquer algorithm. The algorithm is recursive, partitioning the sites into two sets $S = L \cup R$. L consists of sites whose leftmost endpoints are further left than those in R. Recursively, $\mathrm{Vor}(L)$ and $\mathrm{Vor}(R)$ are calculated, and the *contour*, the set of points whose closest sites include at least one from L and one from R, is calculated.

(4.2) Definition *Let $\ell_1, \ldots, \ell_n$ be the leftmost endpoints of all sites in S, given in strictly (see 2.2) ascending order of x-coordinate, and let P_i be a listing of the sites in the same order. If $n > 1$ then L consists of the first $\lfloor n/2 \rfloor$ sites in the list and $R = S \backslash L$.*

Let A (respectively, B) be the set of points whose clearance spheres touch sites in L (respectively, R).

The contour *is the intersection $A \cap B$, i.e., the set of all points whose clearance spheres touch at least one site in L and at least one site in R.*

(4.3) Lemma *The contour is a surface in $\mathbb{R}^3$, possibly disconnected.*

Sketch proof. The contour is a union of faces, edges, and vertices of $\mathrm{Vor}(S)$. It is a surface in the sense that if p is a point in the contour, there exists a neighbourhood V of p such that V intersects the contour in an open disc. This is certainly true if p is interior to a face of the contour. If p is interior to an edge, then there are three incident cells (2.12) of which at least one is from A and one from B. Then p is on the union of two contour faces. If p is a vertex there are four incident cells, at least one and at most three from A. There are six faces meeting at p. If one cell is from A or one from B then V can be interior to the union of its three bounding faces. If two are from A, then V is in the union of four contour faces. **Q.E.D.**

Our regions A and B are defined the same way as regions in Sharir's algorithm, which takes time $O(n \log n)$ to combine them along the contour, leading to an $O(n \log^2 n)$ construction of 2-dimensional Voronoi diagrams [16].

(4.4) Definition *Granted a fixed upper bound on the algebraic complexity of the sites, $C(n)$ is a suitable upper bound for the maximum total number of features over all possible diagrams $\mathrm{Vor}(S)$ where S varies over all possible sets of n sites satisfying 2.12.*

(4.5) Since the diagram for point-sites has complexity at least Dn^2 for some constant D, [14], we assume that $C(n) \geq g(n)n^2$ where $g(n)$ is nondecreasing. It is for this reason that $C(n)$ is taken to be an upper bound rather than the exact estimate.

(4.6) Runtime estimate In the recursive algorithm, if we can combine A with B in time $O(C(n))$, then we get an $O(C(n))$ construction algorithm, since in that case the runtime $T(n)$ satisfies

$$T(n) \leq C(n) + 2T(n/2),$$

and we have the following

(4.7) Lemma *On the assumption that $C(n) = n^2 g(n)$ where $g(n)$ is nondecreasing, $T(n)$ is $O(C(n))$.*

Proof. Divide the recurrence by n^2, getting

$$U(n) = g(n) + \frac{1}{2}U(\frac{n}{2}),$$

where $U(n) = T(n)/n^2$. Informally put, the solution is

$$U(n) = g(n) + \frac{1}{2}g(\frac{n}{2}) + \frac{1}{4}g(\frac{n}{4}) \ldots \leq 2g(n),$$

since $g(n)$ is nondecreasing. Therefore $T(n) \leq 2n^2 g(n) = 2C(n)$. **Q.E.D.**

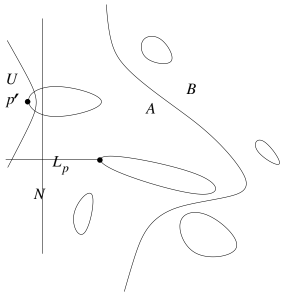

Figure 8: Illustrating Lemma 4.9 in 2-dimensional form

.

(4.8) Lemma (i) *Every Voronoi cell is star-shaped relative to the site it contains.* (ii) *Therefore it is path-connected.*

Proof. For (i), see [14, Lemma 2.1]. That is, suppose that P is the site contained in the cell, x is a point in the cell, and y the point on P closest to x (we allow $y = x$ when $x \in P$). Then the line-segment xy is entirely contained in the cell.

If w is another point in the cell, and z the point in P closest to w, the polygonal path $xyzw$ connects x to w in the cell. **Q.E.D.**

(4.9) Lemma *The region A is (path-) connected, and if G is a sufficiently large closed ball then $G \cap A$ is path connected.*

Proof. For each site P, let p be the leftmost point in P and L_p be the horizontal ray extending leftwards from p. (L_p is parallel to the x-axis.)

Suppose $P \in L$. Consider a point q moving leftwards along L_p. The spheres with centre q and radius $\|q - p\|$ have rightmost point p, which is left of all sites in R, so q is always closer to P than to any site in R: it may be closer to other sites in L, but in any case $L_p \subseteq A$.

Let P' be the unique site in L with leftmost left endpoint, and let that endpoint be p'. There is a vertical plane N, normal to the x-axis, which separates p' from the left endpoints of all other sites in S. All sites in S except P' are to the right of N. The hyperboloid of revolution with focus p' and directrix N separates into two open regions of which one, U say, is convex. Every point in U is closer to p' than to any site except P', so U is contained in the cell owned by P' in $\mathrm{Vor}(S)$. Therefore $U \subseteq A$. Each one of the rays L_p (except $L_{p'}$) intersects the hyperboloid and so intersects U, and also $L_{p'} \subseteq U$. See Figure 8.

Let x and y be two points in A. Let L_p and L_q be the horizontal rays connected with the cells containing x and y respectively. By Lemma 4.8, there are paths in A connecting x to p and y to q. Since L_p and L_q are contained in A and both intersect the convex set $U \subseteq A$, there is a path joining p to q in A, so there is a path joining x to y in A, as required.

If G is a sufficiently large closed ball, so it contains all the sites, and all the rays L_p intersect U within G, then the intersections of the cells with G are star-shaped, and by the same arguments $G \cap A$ is path connected. **Q.E.D.**

The recursive routine needs to construct the contour by developing it locally. It is necessary to furnish starting points from which to develop the contour, one for each contour component. Corollary 4.10 below could say ‘exactly one component.’ That is implicit in Lemma 4.11.

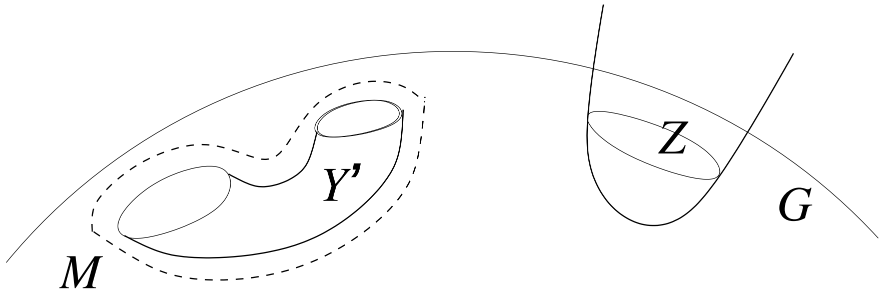

Figure 9: Corollary 4.10: where Y' and Z are separated by M.

(4.10) Corollary *Every component of B is incident to at most one contour component.*

Proof (sketch). We use a theorem of Alexander's [12, Chapter VII, Theorem 1], that if M is a smooth compact 2-manifold in $\mathbb{R}^3$ then $\mathbb{R}^3 \backslash M$ has two components of which M is their common boundary.[3]

Let X be a component of B.

Suppose that ∂X contains two contour components Y and Z. If Y is bounded then one can 'push' this surface component slightly into B to obtain a smooth manifold M with the property that M encloses all cells from A incident to Y, but cells incident to Z are outside M. Since $M \subseteq B$ this contradicts the connectedness of A.

If Y and Z are unbounded then let G be a large closed ball centred at the origin which contains all sites and bounded features of $\mathrm{Vor}(S)$ and intersects all unbounded features, and such that $G \cap A$ is connected (Lemma 4.9). It can be sufficiently large to include all points on the contour which are stationary points with respect to distance from the origin. Then ∂G intersects Y in a finite union of disjoint (topological) circles, where a neighbourhood of each circle is entirely within A on one side and within B on the other.

Choose a point ∞ on $B \cap Z \cap \partial G$, so that for any (topological) circle K on ∂G not meeting ∞, the component of $\partial G \backslash K$ containing ∞ is the outside, and the other is regarded as the inside, of K.

$Y \cap G$ is a manifold with boundary $D = Y \cap \partial G$: D is a disjoint union of (topological) circles. Classify the faces of $\partial G \backslash D$ as follows. F_0 is the face of $\partial G \backslash D$ outside all the circles in D. It includes the point ∞. This face has level zero. Inductively, the face at level $i+1$ are those faces incident to the inner boundary of faces at level i.

Complete the manifold with boundary $Y \cap G$ to a manifold Y' by adding all odd-level faces. Push Y' slightly into $W = (B \cap G) \cup \partial G \cup (\mathbb{R}^3 \backslash G)$ to form a smooth manifold M which is entirely contained in W and disjoint from Z, and whose interior contains F_0, hence contains a point p from A. The exterior of M contains $\infty \in A$. Also since $G \cap A$ is compact, it can be assumed that M has positive distance from $G \cap A$.

By Alexander's theorem every path from p to ∞ intersects M. But $A \cap G$ is path-connected (Lemma 4.9), so it contains a path from p to ∞ in $A \cap G$. This path intersects M at a point at positive distance from $G \cap A$, a contradiction. **Q.E.D.**

(4.11) Lemma *In $O(n^2)$ time overall, at least one point on each contour component can be computed.*

[3] This is the Jordan Curve Theorem in 3 dimensions.

Proof. The ideas in [13,16] work well. We take all the sites Q in R, and connect them to sites in L, so the line-segments connecting them necessarily cross the contour.

Given a site Q in R, choose any reference point q in Q, and calculate the site P in L closest to q, and the point p on P closest to q. The line-segment pq is entirely contained in the cell owned by P in $\mathrm{Vor}(L)$. Since one end is closer to L and the other to R, pq intersects the contour at a point x. The line-segment px is contained in the cell owned by P in $\mathrm{Vor}(S)$. If any point y interior to xq were in that cell, then it would not be star-shaped, contradicting Lemma 4.8. Hence all of xq is in the region B, and therefore in the connected component of B containing q.

For each site Q_i in R, let C_i denote the cell owned by P in $\mathrm{Vor}(\{P, Q_i\})$. By Lemma 4.8 the line-segment pq intersects its boundary, the P, Q_i-bisector, in at most one point x_i. If it exists, x_i is the point on pq equidistant from p and Q_i, and its existence can be checked, and x_i computed, in time $O(1)$.

The point x is where pq meets the boundary of $\bigcap C_i$, and $px = \bigcap px_i$, so x is the x_i closest to p on pq. Thus x can be computed in time $O(n)$, and x is a contour starting point.

Overall, for all sites Q, the cost is $O(n^2)$. Every contour component is incident to at least one cell from R, and this procedure returns a point on the component from that cell (Corollary 4.10). In other words, there is at least one starting point on every contour component. **Q.E.D.**

Recall that $C(n)$ bounds the maximum number of features on n-site Voronoi diagrams. We now return to consideration of frustums (paragraph 3.18).

(4.12) Lemma *The subdivision of cells of* $\mathrm{Vor}(S)$ *into frustums yields* $O(C(n))$ *frustums.*

Proof. The boundary of each frustum has bounded algebraic complexity, and each frustum contains at least one of $O(C(n))$ features — the projection of a Voronoi vertex or a vertex bounding a seam edge or a point where a projected Voronoi edge or a seam edge is orthogonal to the x-axis, and each such point is incident to a bounded number of frustums. **Q.E.D.**

Up to now the semi-algebraic character of sites has not been used. Now,

(4.13) Lemma *If F and F' are frustums, then $F \cap F'$ is a semi-algebraic set of bounded complexity. (This will be addressed in Section 6.)* ∎

The question of finding the frustums containing contour starting points is deferred to Section 5.

(4.14) To construct a contour component can be reduced to 'meshing' frustums. A starting-point is taken, which is contained in a pair V, W of frustums, V from $\mathrm{Vor}(L)$ and W from $\mathrm{Vor}(R)$: therefore the contour intersects both these frustums. Graphs representing the incidence relations between frustums (in $\mathrm{Vor}(L)$ and $\mathrm{Vor}(R)$) are assumed available.

If P and Q are the sites whose cells contain these frustums in $\mathrm{Vor}(L)$ and $\mathrm{Vor}(R)$ respectively, the contour piece is part of the (P, Q)-bisector, call it B. The boundary of $B \cap V \cap W$ can be calculated. This boundary is incident to other frustums, which can be marked for later processing. Generally during the processing, the boundary of that part of the contour component already developed is retained, with incidence relation connecting it to the frustums in $\mathrm{Vor}(L)$ and $\mathrm{Vor}(R)$, and a queue of frustum pairs incident to this boundary is retained. As a frustum pair is removed from the queue, the boundary of their intersection with the contour is computed, and used to adjust the developing boundary.

Next we estimate the overall cost of computing the contour.

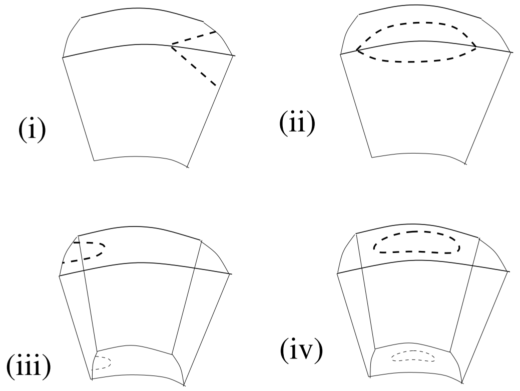

Figure 10: Illustrating the four cases in Lemma 4.15.

(4.15) Lemma *The overall cost of computing the contour by meshing frustum pairs is* $O(C(n))$.

Proof. We argue that whenever we process a frustum pair V, W, where V is a frustum from a cell C in $\mathrm{Vor}(L)$, and W from a cell in $\mathrm{Vor}(R)$, one can link the operation to one of $O(C(n))$ features.

The frustum V is a region of the form

$$V = C \cap \pi_P^{-1}(D),$$

where D is a 'quad' on ∂P and C is the cell containing P in $\mathrm{Vor}(L)$. The quad is bounded by at most two projected edges and two pseudo-edges, orthogonal to the x-axis. The frustum has at most four 'inner corners,' (pseudo-)vertices bounding the quad D, and four 'outer corners,' (pseudo-)vertices on $F \cap \partial C$, which may be considered points at infinity if V is unbounded. The base of V on ∂P remains closer to P than to any other site in S, so no frustum can be entirely deleted in forming $\mathrm{Vor}(S)$. See Figure 10. Analysis of the features linked to the meshing operation raises four cases. Let Q be the site in R whose cell contains W.

Case (i): when a corner, necessarily an outer corner, of V is closer to Q than P then a corner is 'absorbed.' (It is on the other side of the contour.) The local computation can be charged to a corner that is absorbed. This charge is also legitimate for corners at infinity. Similarly if a corner of W becomes closer to P.

Case (ii): when points on one of the (outer) sides of V are closer to Q, but the corners are not, one or more sides of V intersects the P, Q-bisector. Case (ii) is where the bisector intersects a true edge of $\mathrm{Vor}(L)$. If one of these intersecting sides is part of a true Voronoi edge, then it intersects the contour at two vertices of $\mathrm{Vor}(S)$, which can be charged with the local computation: there are at most $C(n)$ such vertices.

Case (iii): where an (outer) pseudo-edge contains points closer to Q but its endpoints don't, the P, Q-bisector intersects the outer face in a curve meeting this pseudo-edge at least twice. The intersection is in an edge of $\mathrm{Vor}(S)$. By a form of Rolle's Theorem, that edge includes a local maximum with respect to the x-direction, and there are $O(C(n))$ such points in $\mathrm{Vor}(S)$.

Case (iv): the bisector intersects the top face of the frustum but intersects no frustum edges. We have a closed-loop Voronoi edge in $\mathrm{Vor}(S)$ and there are at most $C(n)$ such edges. **Q.E.D.**

(4.16) Assembling $\mathrm{Vor}(S)$**.** At this point we have generated a list of pairs (F_1, F_2) of frustums, F_1 from $\mathrm{Vor}(L)$ and F_2 from $\mathrm{Vor}(R)$ respectively, meeting along the contour. There are $O(C(n))$ such pairs. We need a description of $\mathrm{Vor}(S)$ as specified in paragraph 3.15.

For each pair (F_1, F_2), there are a bounded number of intersections of $F_1 \cap F_2$ with contour features, and we can generate a set of triples

$$(F_1, F_2, f)$$

where f varies over these feature intersections. Since every frustum has a face on a site boundary, a frustum may be reduced in volume because it intersects the contour, but it cannot be removed completely.

For every frustum F from L (respectively, R) not intersecting the contour, so it is preserved intact in $\mathrm{Vor}(S)$, generate triples $(F, \emptyset, f)$ or $(\emptyset, F, f)$.

Form a graph whose nodes are these triples, with edges joining (F_1, F_2, f) to (G_1, G_2, g) if f and g are part of the same feature meeting at adjacent frustums.

The connected components of this graph correspond to the features of $\mathrm{Vor}(S)$. The features can be labelled and specified in time $O(C(n))$.

A face is incident to a cell if and only if it intersects F_1 or F_2 for at least one pair (F_1, F_2) where F_1 or F_2 is nonempty and contained in that cell.

An edge is incident to a face if and only if both edge and incident face similarly intersect some F_1 or F_2.

A vertex is incident to an edge if and only if both vertex and incident edge similarly intersect some F_1 or F_2.

Two unbounded branches of an unbounded edge are adjacent along a face in the sense given paragraph 3.15 if and only if they are connected by a chain of frustums intersecting the face, the chain not meeting another branch between them.

Thus the features of $\mathrm{Vor}(S)$, and their incidence relations, can be computed from the frustums and starting points in time $O(C(n))$ overall.

5 Cavities and the quad subdivision

In this section we consider how to construct the quad decomposition on which the frustums are based.

(5.1) Mapping ∂P **into the plane.** Let P be a site in S, C its cell in $\mathrm{Vor}(S)$. For purposes of point-location, it is convenient to separate ∂P into its 'front' and 'back' parts. These two parts are separated by the set X of all points on ∂P where the tangent plane is parallel to the y-axis.

(5.2) Lemma *The set X is a Jordan curve on ∂P.*

Sketch of proof. Since P is compact, strictly convex and rounded, the map from ∂P to S^2, where

$$S^2 = \{(x, y, z) : x^2 + y^2 + z^2\},$$

taking each boundary point p to the outer unit normal at that point, is a homeomorphism. X is the inverse image of

$$Y = \{(x, 0, z) : x^2 + z^2 = 1\}.$$

Since Y is a circle, X is a Jordan curve. ∎

Recall that G is the graph structure on ∂P arising from the inward projection of ∂C, together with the exposed faces on ∂P. The same curve X separates G into front and back parts, which can each be projected injectively onto the $\mathbb{R}^2$ by orthogonal projection. So we can simplify the picture by assuming that the graph G is a planar graph (bounded by the projection of X).

The graph G can be split along this curve in linear time. The front and back halves each project injectively onto the xz plane under orthogonal projection. Furthermore, trapezoidal pseudo-edges project onto vertical line-segments. Vertices project onto vertices and points on edges where the tangent plane is normal to the x-axis project onto points where the tangent is vertical. Piecewise algebraic edges project onto piecewise algebraic edges. Thus we can for the rest of this section confine our attention to calculating a subdivision of a plane (generalised) graph into quads by vertical edges.

(5.3) Proposition *Let F be a face on G, with its boundary presented as $k+1$ Jordan chains. Then the face can be subdivided into quads in time $O(v + k\log^2 k)$ where v is the total number of vertices bounding F.*

Proof. If $k = 0$ then Chazelle's linear-time algorithm [8] for triangulating a simple polygon provides a quad decomposition for simple polygons, and it can be adapted to our problem here where the polygons are not piecewise linear but piecewise algebraic of bounded degree.

More generally, the bounding polygons can be combined into one which can be processed by Chazelle's algorithm, within the stated time-bound. See [5].

See also [2]. ∎

We shall prove that the sum of all the k mentioned in the above proposition is $O(n\log^* n)$. Just as faces of the original graph G on ∂P either come from ∂C by inward projection, or are exposed facets, we can maintain the distinction in the plane graph G. Let us call them *exposed* and *hidden* faces of G.

(5.4) Cavities. Let F be a P,Q-face in $\mathrm{Vor}(S)$, $\pi_P(F)$ its projection. If F is unbounded then its unbounded edges projected onto edges which are connected together by seam edges, so all the unbounded edges on F together form part of the outer boundary. If F is bounded then, since the P,Q-bisector is unbounded, F has a well-defined outer boundary which projects into $\partial\pi_P(F)$.

If $\pi_P(F)$ is not simply connected then ∂F contains closed loops different from its outer boundary, each of which can be surrounded by a loop L in the relative interior of F.

The loop L can be extended to a topological sphere T contained in the cells of P and Q and not meeting any other cell.

To define T, connect points in L to corresponding points in $\pi_P(L)$ and $\pi_Q(L)$ along lines of projection. This produces a topological cylinder bounded by $\pi_P(L)$ and $\pi_Q(L)$. The Jordan curve $\pi_P(L)$ separates ∂P into two discs; add to the cylinder the disc containing $\pi_P(s)$. Similarly add a disc in ∂Q, completing the topological sphere T. T is contained in $C \cup C'$, the cells containing P and Q in $\mathrm{Vor}(S)$, and intersects no other cell. According to Alexander's Theorem [12, Chapter VII, Theorem I], T separates $\mathbb{R}^3$ into two connected open regions, one bounded, of which it is the common boundary. The bounded region contains part of ∂F. Therefore it intersects other cells besides C and C', and cells are connected, so it contains a union of cells of $\mathrm{Vor}(S)$. These cells are 'sandwiched' between C and C'. See Figure 11.

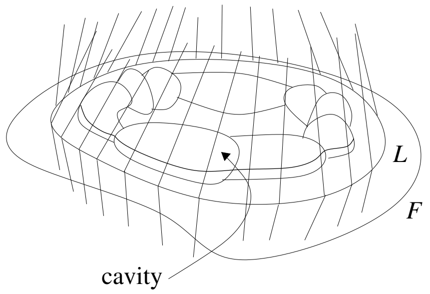

Figure 11: cavity between P- and Q-cells. The hatched lines indicate part of the enclosing topological sphere T.

(5.5) Definition *A* cavity *is a connected union of cells of* $\mathrm{Vor}(S)$*, incident to the inner boundary of a face F and sandwiched between the two cells meeting at F as discussed above.*

(5.6) Theorem *The total number of vertices common to a cavity and a surrounding face, counted over all faces, is* $o(n \log^* n)$.

Proof. We count the Voronoi edges as they meet the boundary in clockwise order. Claim that for any sites X, Y meeting the cavity boundary, their longest alternating subsequence

$$\ldots X \ldots Y \ldots X \ldots Y \ldots$$

is bounded, because if we consider the Voronoi diagram of P, Q, X, Y, ignoring the other sites, then each alternation $X \ldots Y$ is marked by a Voronoi vertex. The number of such Voronoi vertices is bounded since the sites have bounded algebraic complexity.

In other words, the cavity boundary edges form a Davenport-Schinzel sequence of bounded type and there are $o(r \log^* r)$ of them [1]. Since the sites in different cavities are distinct, $\sum r \leq n$, so the overall number of such vertices is $o(n \log^* n)$. **Q.E.D.**

Now we can finish the subdivision of the projected graph G in ∂P, as described before Lemma 3.17.

(5.7) Lemma *All the faces can be subdivided into quads in overall time* $O(C(n))$.

Proof. Let F be a face on G. By Proposition 5.3, F can be subdivided into quads in time $O(v + k \log^2 k)$ where v is the total number of vertices bounding F and $k+1$ the number of boundary components. By Theorem 5.6, $\sum k$ is $o(n \log^* n)$, where the sum is over all hidden faces (projected from cell boundaries). The total number of vertices and edges on the exposed facets is bounded by the complexity of the convex hull $H(S)$, which is also bounded by $C(n)$), so the overall time is $O(C(n) + n \log^2 n \log^* n)$, which is $O(C(n))$ since $C(n)$ is $\Omega(n^2)$. **Q.E.D.**

Having taken the front and back of ∂P separately, split the graph by the Jordan curve separately, and projecting into the xz-plane, we produced the graph G in the plane (there are two such graphs), which we now have subdivided into quads. Let $G(P, S)$ denote the quad subdivision of G — it is the front or back of the true graph $G(P, S)$ on ∂P (Definition 3.18).

It remains to construct a monotone subdivision of $G(P, S)$.

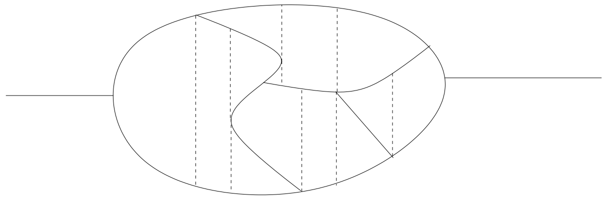

Figure 12: the monotone subdivision $G(P, S)$.

(5.8) Definition *A planar region is* monotone *if its intersection with any vertical line is connected [10, Chapter 11]. A* monotone subdivision *of the plane is a subdivision into monotone regions.*

In [10, Chapter 11] a planar-point location structure is built for monotone subdivisions. The construction is given for straight-edge subdivisions without vertical edges, but it is remarked there that the method applies generally, in particular to our graph $G(P, S)$, if infinite horizontal rays are added extending from the leftmost and rightmost points of $G(P, S)$, to produce a monotone subdivision. See figure 12. Therefore

(5.9) Proposition *Let $G(P, S)$ have k features. In $O(k)$ time a data-structure can be constructed to locate the quad containing a given point in time $O(\log^2 k)$ [10, Lemmas 11.8, 11.9].* ∎

Taking all sites P together we get quad search structures for each site boundary in time $O(C(n))$.

(5.10) Corollary *Search structures can be constructed in time $O(C(n))$ which solve the following queries in time $O(\log^2 n)$: Given a point x in the cell owned by P in $\mathrm{Vor}(S)$, but not in P itself, return the frustum (or few neighbouring frustums) containing x.*

Proof. In time $O(\log^2 n)$, calculate the unique quad or quads on ∂P containing $\pi_P(x)$. These quads are the bases of the frustums containing x. **Q.E.D.**

6 The features of $\mathrm{Vor}(S)$ are semi-algebraic

A set $S \subseteq \mathbb{R}^n$ is a **algebraic half-space** if it satisfies a polynomial inequality. If the inequality is strict, then the halfspace is open, otherwise it is closed. It is **basic set** if it is the intersection of finitely many algebraic half-spaces. It is **semialgebraic** if it is the Boolean combination of algebraic half-spaces.

A predicate $P(x)$ is **semialgebraic** iff the set $\{x \in \mathbb{R}^n : P(x)\}$ is semialgebraic. A function $f(p)$ is **semialgebraic** if its graph $\{(p, f(p)) : p \in \mathbb{R}^n\}$ is a semialgebraic set.

A set $S \subseteq \mathbb{R}^n$ is a i**-cell** (or simply, cell) if it is homeomorphic to an open i-ball ($i = 0, \ldots, n$). The following result may be called the **fundamental theorem of semialgebraic sets**:

(6.1) Proposition *(a) Semialgebraic sets are closed under projection.*
(b) A semialgebraic set is a finite union of cells where these cells are semialgebraic sets and form a regular cell complex. ∎

These results follow from the theory of cylindrical cell decomposition [15,6].

(6.2) Theorem *Let A, B be semialgebraic sets. Then the following are semialgebraic:*
(a) The boundary ∂A.
(b) The distance function $d(p, A)$.
(c) The convex hull $H(A)$.
(d) Each connected component of A.

Proof. (a) We have $p \in \partial A$ iff for all $r > 0$, there is some $q_1 \in A$ such that $|p-q_1|^2 < r^2$ and another $q_2 \notin A$ such that $|p - q_2|^2 < r^2$
(b) We have $x = d(p, A)$ iff there is some $q \in A$ such that $|p - q|^2 = x^2$ and for all $r \in A$, $|p - r|^2 \geq x^2$.
(c) Suppose $A \subseteq \mathbb{R}^n$. Then $p \in H(A)$ iff there exists $p_0, \ldots, p_n \in A$ and $\alpha_0, \ldots, \alpha_n \in [0, 1]$ such that $p = \sum_{i=0}^{n} \alpha_i p_i$ and $1 = \sum_{i=0}^{n} \alpha_i$.
(d) By the fundamental theorem of semialgebraic sets, we can partition A into finitely many cells. Clearly each connected component of A is the finite union of these cells. **Q.E.D.**

Let K be a collection of pairwise disjoint semialgebraic bodies in $\mathbb{R}^n$. For $L \subseteq K$, we define

$$V(L) = V_K(L) := \{p \in \mathbb{R}^n : (\forall A, B \in L)(\forall C \in K \setminus L)[d(p, A) = d(p, B) \leq d(C)]\}.$$

Each connected component of $V(L)$ is a **Voronoi cell**.

(6.3) Corollary *Voronoi cells are semialgebraic.*

Proof. $V(L, K)$ is semialgebraic. By the previous theorem (d), its connected components (i.e., Voronoi cells) are semialgebraic. **Q.E.D.**

(6.4) Lemma *If S is a set of disjoint strictly convex compact semialgebraic sites and S_ε consists of the expanded version of these sites, then the expanded (rounded) sites are semialgebraic. (See Paragraph 2.8).*

Proof. If P is a strictly convex compact semialgebraic site, its expansion is $P_\varepsilon = \{p : (\exists q \in P)(|p - q|^2 \leq \varepsilon^2)\}$. This is a semialgebraic set. **Q.E.D.**

(6.5) Lemma *The convex hull of semialgebraic sites is semialgebraic. (Theorem 6.2* (c).*)* $\blacksquare$

(6.6) Lemma *The projection of Voronoi edges and faces from the boundary of the P-cell in* $\mathrm{Vor}(S)$ *onto ∂P is semialgebraic.*

Proof. These are projections $\pi_P(A)$ of semialgebraic sets A onto the semialgebraic set ∂P.

$$x \in \pi_P(A) \iff (\exists y \in A)(\forall z \in P)(|x - y|^2 \leq |z - y|^2)$$

Q.E.D.

(6.7) Lemma *If P and Q are disjoint sites then the P, Q-seam is semialgebraic.* $\blacksquare$

(6.8) Lemma *Frustums are semialgebraic.* $\blacksquare$

7 Acknowledgements

The authors are grateful to Bernard Chazelle and Mike Goodrich for information about polygon triangulation.

8 References

1. Pankaj Agarwal and Micha Sharir (1995). *Davenport-Schinzel sequences and their geometric applications.* Cambridge University Press.

2. Nancy M. Amato, Michael T. Goodrich, and Edgar A. Ramos (2000). Linear time triangulation of a simple polygon made easier via randomization. *Proc 16th Annual ACM Symposium on Computational Geometry,* 12 pages.

3. Franz Aurenhammer (1987). Power diagrams: properties, algorithms, and applications. *SIAM Journal of Computing* **16:1**, 78–96.

4. Franz Aurenhammer (1990). Voronoi diagrams— A survey of a fundamental geometric data structure. *ACM Computing Surveys* **23:3**, 345–405.

5. Reuven Bar-Yehuda and Bernard Chazelle (1994). Triangulating disjoint Jordan chains. *International Journal of Computational Geometry and Applications* **4:4,** 475–481.

6. Saugata Basu, Richard Pollack, and Marie-Françoise Roy (2003). *Algorithms in Real Algebraic Geometry*, Springer Series on Algorithms and Computation in Mathematics.

7. Jean-Daniel Boissonnat, André Cérézo, Olivier Devillers, Jacqueline Duquesne, and Mariette Yvinec (1996). An algorithm for constructing the convex hull of a set of spheres in dimension d. *Computational Geometry: Theory and Applications* **6:2**, 123–130.

8. Bernard Chazelle (1991). Triangulating a simple polygon in linear time, *Discrete and Computational Geometry* **6,** 485–524.

9. Bernard Chazelle (1993). An optimal convex hull algorithm in any fixed dimension. *Discrete and Computational Geometry* **10**, 377-409.

10. Herbert Edelsbrunner (1987). *Algorithms in Combinatorial Geometry.* EATCS monographs on Theoretical Computer Science, Springer-Verlag.

11. Steven Fortune (1987). A sweepline algorithm for Voronoi diagrams. *Algorithmica* **2:3**, 153–174.

12. André Gramain (1971). *Topologie des surfaces.* Presses Universitaires de France.

13. D.T. Lee and S. Drysdale (1981). Generalizations of Voronoi diagrams in the plane. *SIAM Journal of Computing* **10:1**, 73–87.

14. Colm Ó Dúnlaing (2003). Local differentiability and monotonicity properties of Voronoi diagrams for disjoint convex sites in three dimensions. Report TCDMATH 03–09.

15. Jacob T. Schwartz and Micha Sharir (1983). On the piano movers' problem: II. General techniques for computing topological properties of real algebraic manifolds. *Advances in Appl. Math.* **4**, 298–351.

16. Micha Sharir (1987). Intersection and closest-pair problems for a set of planar discs, in *Planning, Geometry, and Complexity of Robot Motion,* ed. Schwartz, Sharir, and Hopcroft, Ablex Publishing Corporation, New Jersey 1987, pp 214–239.

17. Nicola Wolpert (2002). An exact and efficient approach for computing a cell in an arrangement of quadrics. Doctoral dissertation, University of the Saarland, Saarbrücken.

18. Chee K. Yap (1987). An $O(n \log n)$ algorithm for the Voronoi diagram of a set of simple curve segments. *Discrete and Computational Geometry* **2**, 365–393.