Local differentiability and monotonicity properties of Voronoi diagrams for disjoint convex sites in three dimensions

Colm Ó Dúnlaing*
Mathematics, Trinity College, Dublin 2, Ireland

December 31, 2003

Abstract

This paper studies local properties of Voronoi diagrams of sets of disjoint compact convex sites in $\mathbb{R}^3$.

It is established that bisectors are C^1 surfaces and trisectors are C^1 curves, and that as a point moves along a trisector its clearance sphere develops monotonically (Lemma 2.4). This monotonicity property is useful in establishing the existence of Voronoi vertices bounding edges in certain situations.[1]

The paper then considers the diagram for a set of disjoint spheres. Considerations about general position are covered in detail. By letting the spheres grow from point sites till they reach their true radius, it is shown that the Voronoi cell for the smallest site has complexity $O(n)$, assuming that the sites are of at most k distinct radii. It follows that the Voronoi diagram is $O(n^2)$.

Although this is weaker than Aurenhammer's result [1] establishing $O(n^2)$ complexity with no restriction on radius, the techniques may be of value for studying more general Voronoi diagrams.

Finally, the paper shows that without the bound on the number of different radii, the cell owned by a point site can have complexity $\Omega(n^2)$.

1 Voronoi diagrams: differentiability properties

This paper considers the Voronoi diagrams of spherical sites in $\mathbb{R}^3$. For a general survey of Voronoi diagrams see [2].

The current state of knowledge about the complexity of Voronoi diagrams, in 3 dimensions, is scanty.

*e-mail: odunlain@maths.tcd.ie. Mathematics department website: http://www.maths.tcd.ie. Some of this work was presented at the eighteenth European Conference on Computational Geometry, Bonn, Germany, April 2003.

[1] Chee Yap [8] seems to have been first to exploit such monotonicity properties.

It is known to be $O(n^2)$ for spheres, whether or not they are disjoint [1]. This bound is tight even for point sites (which are discussed in this paper).

When the sites are straight lines, the complexity is known to be $o(n^{2+\epsilon})$ for all $\epsilon > 0$, granted that either the distance function is polyhedral, based on a fixed convex polyhedron [5], or the distance is Euclidean but the lines are in $O(1)$ different directions [7].

Andrew Farrell's 1994 M.Sc. dissertation [6] considers convex sites in two dimensions and includes an implementation study in the case of circular sites in two dimensions.

Let S be a finite nonempty set of pairwise disjoint closed balls in $\mathbb{R}^3$. Balls of radius 0, equivalent to points, are allowed. If x is a point outside a closed ball B of radius r and centre c, let y be the point where the line xc intersects the boundary of B. Then y is the closest point to x in B, and $|x-y| = |x-c| - r$ is the distance from x to B ($|x-y|$ is the Euclidean distance from x to y).

The following result is valid for any closed and bounded site, not just spheres.

(1.1) Lemma (i) *If B is a bounded, closed, convex, nonempty set in $\mathbb{R}^3$ and x is any point in $\mathbb{R}^3$ then there is a unique point in B closest to x.* (ii) *The map taking x to the closest point in B is continuous.* (iii) *If $x \notin B$ then the closest point is in the boundary ∂B of B.*

Proof. Let F be the family of all closed balls C centred at x such that $C \cap B \neq \emptyset$. It is nonempty, i.e., such balls exist, because B is nonempty. If $C_1, C_2, \ldots, C_k \in F$, where without loss of generality C_1 has the smallest radius, then $C_1 = C_1 \cap C_2 \cap \ldots \cap C_k$, and $C_1 \cap B = (C_1 \cap B) \cap (C_2 \cap B) \cap \ldots \cap (C_k \cap B)$. The left-hand side is nonempty, so $\{B \cap C \colon C \in F\}$ is a set of closed subsets of B with the finite intersection property, hence $B \cap \bigcap F \neq \emptyset$, since B is compact.

The intersection $\bigcap F$ is a closed ball C of radius r, say, where r may equal zero. $C \cap B$ is the nonempty intersection mentioned, and r is the minimum radius of all balls in F. If $C \cap B$ contained two points y_1 and y_2 then (since $B \cap C$ is convex) it would contain a point $z = (y_1 + y_2)/2$ strictly closer to x. In this case the closed ball centred at x of radius $|z - x|$ would intersect B and belong to F, which is false since $|z - x| < r$. Hence $C \cap B = \{y\}$ for some point y. The point y is the unique point in B closest to x, as asserted in (i).

If $x \in B$ then $r = 0$ and $y = x$.

The line-segment xy cannot intersect B except at y, because any other intersection point $z \in B \cap xy$ would have $|xz| < r$. If $x \notin B$ then y is not interior to B, because otherwise xy would intersect B in more than one point. Hence if $x \notin B$ then $y \in \partial B$, proving (iii).

For (ii), suppose that x_1 and x_2 are two points and y_1 and y_2 the points in B closest to each. We need to show that as x_2 tends to x_1, y_2 tends to y_1. This follows from the stronger result: $|y_1 y_2| \leq |x_1 x_2|$, which depends on the convexity of B.

If both points x_i are in B then $y_i = x_i$ and the result holds. If $y_1 = y_2$ then the result holds. Otherwise suppose $x_1 \notin B$ and $y_1 \neq y_2$.

The plane through y_1 perpendicular to $x_1 y_1$ separates x_1 from all of B. Therefore y_2 is in the plane or on the other side from x_1, so the angle $x_1 y_1 y_2$ is at least $90°$. Let Π_1 be the plane through y_1 perpendicular to $y_1 y_2$. Then x_1 is in the closed halfspace H_1 bounded by Π_1 and not containing y_2. The same holds even if $x_1 = y_1$.

By the same reasoning, x_2 is in the closed halfspace H_2 bounded by the plane Π_2 through y_2 perpendicular to $y_1 y_2$ and not containing y_1.

These planes are parallel and the line $y_1 y_2$ is perpendicular to them: the distance $|y_1 y_2|$ is minimal among all pairs z_1, z_2 of points, one in H_1 and one in H_2. In particular, $|y_1 y_2| \leq |x_1 x_2|$. Q.E.D.

Actually it is not necessary that B be bounded. Some closed ball K around x intersects B, $K \cap B$ is compact, and the finite intersection property still applies.

The balls in S are called (spherical) *sites*. Write $\bigcup S$ for the union of the sites in S. Given any point x in $\mathbb{R}^3 \setminus \bigcup S$, its *clearance* is the distance from x to the closest site in S, its *clearance sphere* is the sphere with centre x whose radius is the clearance of x, and the *Voronoi diagram* $\mathrm{Vor}(S)$ of S consists of all points in $\mathbb{R}^3 \setminus \bigcup S$ whose clearance spheres touch more than one site.

The Voronoi diagram is a 2-dimensional complex with faces, edges, and vertices. The faces are connected subsets of surfaces called *bisectors*.

(1.2) Definition *Let B and B' be nonempty, disjoint, closed, bounded, convex sites (point sites are allowed), Then the* (B, B')-bisector *is the Voronoi diagram of $\{B, B'\}$, that is, the set of points equidistant from B and B'. If B'' is a third site disjoint from B and B', then the* (B, B', B'')-trisector *is the set of points equidistant from the three sites.*

A face *of* $\mathrm{Vor}(S)$ *is defined as follows. Let X be the set of points in the (B, B')-bisector which are closer to B and B' than to any other site in S. Equivalently, the points in X have clearance spheres touching B and B' and no other site. An* open *(B, B')-face is a connected component of X, and a (B, B')-face is the closure of an open face.*

The definition of face is complicated because two (B, B')-faces can meet at a vertex. This is impossible for point-sites, but is otherwise possible at pinch points (1.17).

(1.3) Lemma (Tangent plane principle.) *Suppose x is a point on the (B_1, B_2)-bisector, and its clearance sphere touches these sites at p_1, p_2, respectively. Then the (p_1, p_2)-bisector is a plane tangent to the (B_1, B_2)-bisector at x.*

Sketch proof. Let Π_i be the plane through p_i perpendicular to xp_i, $i = 1, 2$. Then the (B_1, B_2)-bisector is sandwiched between the (p_1, Π_2) bisector and the (p_2, Π_1)-bisector. These two surfaces are paraboloids of revolution with a common tangent plane at x, and that tangent plane is the (p_1, p_2)-bisector. Q.E.D.

(1.4) Corollary *A bisector is continuously differentiable.*

Proof. The (B_1, B_2)-bisector has a tangent plane at any point x. Also, this tangent plane is the perpendicular bisector of points y_1 and y_2 on B_1 and B_2 respectively, closest to x. But these points depend continuously on x (Lemma 1.1), and so does the tangent plane. Q.E.D.

This result cannot be strengthened: the second derivative need not exist. For example, if $B_1 = (0, 0, 1/4)$ and B_2 is the closed line-segment connecting $(-1, 0, -1/4)$ to $(1, 0, -1/4)$, and we consider the cross-section where the (B_1, B_2)-bisector intersects the xz-plane, it is described by equations

$$z = \begin{cases} x^2 \quad \text{if } |z| \leq 1 \\ 2|x| - 1 \quad \text{if } |z| \geq 1 \end{cases}$$

The partial derivative $\partial z / \partial x$ exists and is continuous, as the above lemma predicts. Near $x = 1$, $\partial^2 z / \partial x^2$ approaches 2 from the left and 0 from the right, so it does not exist at $x = 1$. (In the case of spherical sites the bisectors are hyperboloids of revolution, infinitely differentiable.)

(1.5) Definition *Let B be one of the sites in a set S of disjoint compact convex sites. The* Voronoi cell *of B consists of all points $x \in \mathbb{R}^3 \setminus \bigcup S$ which are as close, or closer, to B than to any other site in S.*

Strictly speaking, the site owning a cell is disjoint from that cell, but by abuse of notation the cell may sometimes be considered to include the site itself.

As an example of this ambiguity, when sites all have the same radius then the cells are convex polyhedra, which would be nonsensical if one insists that sites themselves are disjoint from the cells.

The Voronoi diagram is the union of boundaries of cells of all sites in S.

(1.6) Lemma *If B and B' are spherical sites with radii r, r' respectively, where $r \geq r' \geq 0$, then the bisector of B and B' is a plane if $r = r'$ and a (single sheet of a 2-sheeted) hyperboloid of revolution, whose axis is the line joining their centres, if $r > r'$.*

In either case, the bisector partitions $\mathbb{R}^3$ into two regions, and that containing B' is convex.

Sketch proof. Let c and c' be their respective centres. The bisector is

$$\{x \in \mathbb{R}^3 \setminus (B \cup B') \colon |x - c| - r = |x - c'| - r'\}$$

or

$$\{x \in \mathbb{R}^3 \setminus (B \cup B') \colon |x - c| - |x - c'| = r - r'\}.$$

This is a plane if $r - r' = 0$ and one sheet of a 2-sheeted hyperboloid of revolution, axis as stated, if $r > r'$. The observation about the convex region follows from the general shape of such hyperboloids of revolution. See Figure 1.

Alternatively, the bisector is the same as that separating the point c' from a sphere B'' of radius $r - r'$, centre c. Let x be a point on the bisector, c'' the point in B'' closest to x: it is where the boundary of B'' intersects the line-segment xc.

The tangent plane to the bisector at x is also the bisector separating c'' from c' (Lemma 1.3). Let y be another point on the plane. It is equidistant from c'' and c', hence it is closer to B'' than to c', or to B than to B'. In other words, all points in the cell of B' in $\mathrm{Vor}(B, B')$ are on the same side of this plane. Hence the cell containing B' is convex. ∎

Remark: the argument shows something slightly more general, that is, if the closest point to x on B' is an extreme point on B', then the (B, B')-bisector is locally on one side of the tangent plane at x, i.e., x is not a saddle-point on the bisector.

(1.7) Corollary *Suppose that S is a set of disjoint spherical sites whose minimum radius is r_1. Let S' be the set of spherical sites obtained by replacing every site B in S by a site with same centre and radius $r - r_1$, where r is the radius of B.*

Then $\mathrm{Vor}(S) = \mathrm{Vor}(S')$, and the smallest sites in S' are point sites. (Proof: the clearance function is increased by r_1 uniformly.) ∎

The complexity of the Voronoi diagram for point sites is known:

(1.8) Lemma *If S is a set of n point sites, then $\mathrm{Vor}(S)$ has n cells and $O(n^2)$ faces, edges, and vertices.*

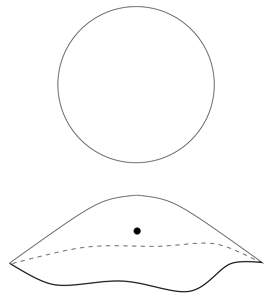

Figure 1: The bisector is a hyperboloid of revolution.

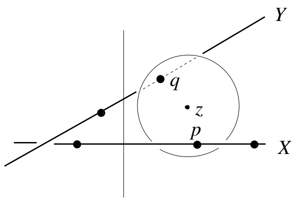

Figure 2: Sphere tangent to two lines at two given points.

Proof. Let p be any (point) site in S. For any other site p', the bisector between p and p' is a plane perpendicular to the line pp' and passing through $(p + p')/2$.

The set of points x as close to p as to p' is the closed half-space bounded by this plane and containing p. In general, the Voronoi cell containing p is the intersection of $n - 1$ closed half-spaces; hence it is a convex polyhedron with $O(n)$ faces, edges, and vertices.

Since each cell has complexity $O(n)$, $\mathrm{Vor}(S)$ has complexity $O(n^2)$. Q.E.D.

(1.9) More precisely, each cell of the Voronoi diagram is a convex polyhedron with at most $2n - 7$ vertices [4].

(1.10) This bound is tight. The Voronoi diagram for n point sites in $\mathbb{R}^3$ can have complexity $\Omega(n^2)$. To see this, let X he the x-axis and Y the line $x = 0, z = 1, -\infty < y < \infty$ parallel to and above the y-axis. For any point p on X and q on Y there exists a sphere tangent to X at p and Y at q, See Figure 2. It follows that if $n/2$ point sites are arranged on X and $n/2$ on Y, then for each pair p, q of sites from these respective sets, there is a sphere, centre z, say, tangent to X at p and Y at q. Then $z \in \mathrm{Vor}(S)$ belongs to the (p, q)-bisector, and the face separating p from q is nonempty. Hence $\mathrm{Vor}(S)$ has $\Omega(n^2)$ faces.

There follows a collection of facts about Voronoi diagrams which will be needed in this paper. All of these facts apply to sets of disjoint convex compact sites, not just spherical sites.

(1.11) Lemma (Site reflection principle.) *Let B_1, B_2 be disjoint convex compact sites. Suppose that x is a point on the (B_1, B_2)-bisector, p_1 is where the clearance sphere at x touches B_1, and T is the tangent plane to the bisector at x. Then the clearance sphere touches B_2 at the perpendicular reflection of p_1 in T.*

Proof. If p_2 is the point of contact of the clearance sphere with B_2, then we know that the plane bisecting p_1p_2 is the tangent plane T (Lemma 1.3). Then p_2 is the perpendicular reflection of p_1 in this plane. Q.E.D.

(1.12) Corollary (Bisectors don't touch.) *Let B_1, B_2, B_3 be disjoint convex compact sites. It is impossible for a (B_1, B_2) and a (B_1, B_3)-bisector to touch.*

Proof. Suppose they touched at a point x, equidistant from three closest points p_i on the respective sites B_i. From the tangent plane principle, the p_1, p_2-bisector and the p_1, p_3-bisector are tangent at x, that is, they coincide, and from the site reflection principle, $p_2 = p_3$, which is impossible since the sites are disjoint. Q.E.D.

Corollary 1.4 considers differentiability in the geometric sense of possessing a well-defined tangent plane. More formally we define differentiability in terms of the *Fréchet derivative*:

(1.13) Definition *Let a function $f\colon \mathbb{R}^n \to \mathbb{R}^m$ be defined in a neighbourhood of x. If there exists a* linear *map $A\colon \mathbb{R}^n \to \mathbb{R}^m$ such that for any 'small' displacement Δx*

$$f(x + \Delta x) = f(x) + A\Delta x + o(|\Delta x|)$$

then A is unique and is called the (Fréchet) derivative *of f at x and f is said to be* differentiable *at x.*

(1.14) Lemma *Let B be a nonempty compact convex site. Then the map $d : \mathbb{R}^3 \backslash B \to \mathbb{R};\ x \mapsto |x-y|$ where y is the closest point to x in B, is a C^1 map (continuously differentiable).*

Proof. Fix x, thus also fixing y, and write d for $|x - y|$. Let $\vec{N} = (x - y)/|x - y|$. Then $\vec{N} \cdot (x - y) = d$. For small Δx, the distance from $x + \Delta x$ to the plane through y perpendicular to $x - y$ is $\vec{N} \cdot (x - y) + \vec{N} \cdot \Delta x = d + \vec{N} \cdot \Delta x$. Since this plane separates x from B, if Δx is small then the distance from $x + \Delta x$ from this plane is a lower bound on its distance from B. The distance $|x + \Delta x - y|$ is an upper bound on this distance (since $y \in B$). Note

$$|x + \Delta x - y|^2 = |x - y|^2 + 2(x - y) \cdot \Delta x + |\Delta x|^2 = d^2 + 2d\vec{N} \cdot \Delta x + |\Delta x|^2,$$

since $x - y = d\vec{N}$. Squaring the lower bound, we get

$$d^2 + 2d\vec{N} \cdot \Delta x + (\vec{N} \cdot \Delta x)^2.$$

Ignoring the term $(\vec{N} \cdot \Delta x)^2$, we get the inequalities

$$d^2 + 2d\vec{N} \cdot \Delta x \le (d(x + \Delta x))^2 \le d^2 + 2d\vec{N} \cdot \Delta x + |\Delta x|^2.$$

Take square roots: then

$$d(x + \Delta x) = d(x) + \vec{N} \cdot \Delta x + o(|\Delta x|).$$

In other words, the function $d(x)$ is differentiable with derivative $\vec{N} = (x - y)/|x - y|$. Since y, the point in B closest to x, depends continuously on x, (Lemma 1.1), the derivative is continuous: the function $x \mapsto d(x)$ is C^1. Q.E.D.

The following corollary essentially repeats Corollary 1.4.

(1.15) Corollary *Let B_1, B_2 be disjoint compact convex sites. The (B_1, B_2)-bisector is a C^1 surface except perhaps at the midpoint of the line joining a closest pair of points on B_1 and B_2.*

Proof. Let $f(x) = |x - y_1| - |x - y_2|$ where y_i are the points closest to x on B_i. This function is differentiable with derivative $N_1 - N_2$, where $N_i = (x - y_i)/|x - y_i|$ are unit normals. Suppose the derivative vanishes. Then $x - y_1 = x - y_2$, so $x = (y_1 + y_2)/2$.

For $i = 1, 2$, consider the plane through y_i perpendicular to xy_i; let H_i be the halfspace bounded by this plane, not containing x. The bounding planes are parallel and $|y_1 - y_2|$ is minimal for all pairs of points, one in H_1 and the other in H_2. Since $B_i \subseteq H_i$, y_1, y_2 are a closest pair of points from B_1, B_2 respectively and x is their midpoint.

Otherwise, the derivative is nonzero and the surfaces is a C^1 surface locally, by the Implicit Function Theorem [10]. Q.E.D.

(1.16) Corollary *Let B_1, B_2, B_3 be disjoint compact convex sites.* (i) *For any point x on the (B_1, B_2, B_3)-trisector, the trisector is tangent to the (y_1, y_2, y_3)-trisector, where each y_i is the point on B_i closest to x.* (ii) *The (B_1, B_2, B_3)-trisector is a C^1 curve.*

Proof. (i) By Lemma 1.3, the trisector is tangent to the (y_1, y_2)-, (y_2, y_3)-, and (y_1, y_3)-bisector planes. The points y_j are not collinear (since they are all equidistant from x), so these three planes intersect in a line, the (y_1, y_2, y_3)-trisector, which is tangent to the trisector.

(ii) Consider the function $f(x) = (|x - y_2| - |x - y_1|, |x - y_3| - |x - y_1|)$, where y_i are the points in B_i closest to x. The trisector is $f^{-1}(0, 0)$.

The derivative of f (multiplied by $|x - y_i|$, the same for each i) is $(y_1 - y_2, y_1 - y_3)$. This derivative is singular only when the two components are proportional, i.e., $y_1 - y_2 = \alpha(y_1 - y_3)$ for some α, and this happens only when y_1, y_2, y_3 are collinear, which is impossible. Hence x is a regular point for f and by the Implicit Function Theorem [10] the trisector is a C^1 curve near x. Q.E.D.

A Voronoi vertex is incident to four or more cells (or has four or more sites closest to it). A degenerate case is that of a pinch point:

(1.17) Definition *Let S be a set of disjoint compact convex sites. A* pinch point *in $\mathrm{Vor}(S)$ is a vertex where two edges meet with a common tangent.*

(1.18) Lemma *Let S be a set of disjoint compact convex sites, and let v be a vertex of $\mathrm{Vor}(S)$. Then v is a pinch-point if and only if among the four or more sites closest to v there are four, B_i ($1 \leq i \leq 4$, say), such that the four points p_i closest to v on B_i are coplanar.*

Consequently, either v is a pinch-point with respect to four closest sites B_i, and all four of the trisectors involving three of these sites meet tangentially, or v is not a pinch-point and all trisectors involving three of the sites closest to v meet transversally.

Proof. Let p_i be the point closest to v on B_i ($i = 1, \ldots, 4$). For any three distinct i, j, k, the (p_i, p_j, p_k)-trisector in $\mathrm{Vor}(p_1, p_2, p_3, p_4)$ is tangent to the (B_i, B_j, B_k)-trisector at v (Corollary 1.16 (i)). If the four points are not coplanar then there are four trisectors meeting transversally. If the four points are coplanar (hence concyclic) then all four of the trisectors (with three closest sites B_i) meeting at v have the same tangent.

Thus if there exist four closest sites B_i such that the closest points p_i are coplanar, v is a pinch-point, and the four trisectors indicated have a common tangent at v, and if v is a pinch-point then these sites B_i exist. If v is not a pinch-point then the trisectors determined by all possible sets of three sites closest to v meet transversally at v. Q.E.D.

(1.19) Lemma (Multi-directional principle.) *Let S be a set of disjoint compact convex sites and v a vertex in $\mathrm{Vor}(S)$, not a pinch-point. Then there exists no plane Π through v with the property that all edges incident to v meet v from the same side of Π.*

Proof. Let $\{B_i : 1 \leq i \leq k\}$ be the sites and p_i the points on B_i closest to v; $k \geq 4$ and by the previous lemma no four points p_i are coplanar. They are all on the boundary of a sphere centred at v. Let $P = \{p_1, p_2, \ldots\}$.

We begin by considering $\mathrm{Vor}(P)$. The Voronoi diagram consists of infinite Voronoi edges all meeting at the same vertex v.

Let C be the convex hull of P. Since no four points in P are coplanar, C has nonempty interior (and all faces are triangular). The plane Π can be assumed horizontal, and it is enough to show that $\mathrm{Vor}(P)$ contains at least one edge directed upwards. If p_i, p_j, p_k are corners of a face of C then $\mathrm{Vor}(P)$ contains a (p_i, p_j, p_k)-edge, which is directed away from v along the outward normal to that face. Therefore it is enough to show that there exists a face whose outward normal is directed upwards.

Let Π' be the highest horizontal plane which intersects the convex hull. If it contains more than two corners, it intersects a face; let p be interior to the face. If it contains two corners, it intersects an edge; let p be interior to the edge. Otherwise let p be the unique corner of the hull in Π'. All the faces not meeting p are at positive distance from p. Let r be the minimum of these distances and let B be the open ball, centred at p, of radius r. Then all faces of the hull which intersect B are incident to p. Let the outward normals to these faces be $\vec{N}_i$, $1 \leq i \leq \ell$. The faces satisfy equations $\vec{N}_i \cdot \vec{px} = 0$. A point x in B is in the convex hull if and only if $\vec{N}_i \cdot \vec{px} \leq 0$ for all i, $1 \leq i \leq \ell$. Choose a point x directly above p in B. It is above Π', therefore not in the convex hull, so $\vec{N}_i \cdot \vec{px} > 0$ for some i, and that outward normal $\vec{N}_i$ is directed upwards.

Since the edges of $\mathrm{Vor}(P)$ are the tangents to those edges meeting v in $\mathrm{Vor}(S)$ (Corollary 1.16 (i)), the edges incident to v cannot all be incident from the same side of Π. Q.E.D.

2 Voronoi diagrams: monotonicity properties

In this section, we include some results about the way a clearance sphere develops as its centre moves along a certain curve. These results hold for disjoint convex sites, not just spheres. The simplest is as follows

(2.1) Lemma *Suppose x is a point in the cell owned by a compact convex site B in the Voronoi diagram. Let y be the point closest to x in B. Then every point on the line between x and y is interior to the cell. Consequently the cell is star-shaped relative to B.*

Proof. Let S be the clearance sphere around x, with centre x and radius $|yx|$. Let z be a point (strictly) between x and y on the line xy.

For any point w, suppose that $|wz| \leq |yz|$. Then $|wx| \leq |wz| + |zx| \leq |yz| + |zx| = |yx|$, i.e., $|wx| \leq |yx|$ with equality only if z is on the line-segment wx — ensuring $|wx| = |wz| + |zx|$ — and $|wz| = |yz|$.

Suppose that $|wx| = |yx|$, so z is on the line-segment wx. Since $x \neq z$ we can write

$$w = x + \frac{|wx|}{|xz|}(xz) \quad \text{and} \quad y = x + \frac{|yx|}{|xz|}(xz),$$

so $w = y$.

Thus the sphere T with centre z, radius yz, is inside S, strictly so except at y. Therefore it meets no site except at y. Since $y \in B$, y is closest to z in B and the sphere T is the clearance sphere for z, so z is strictly closer to B than to any other site, i.e., interior to the cell owned by B. Q.E.D.

Lemma 2.4 is a more elaborate form of the same idea. The following lemma supports it.

(2.2) Lemma *Let C be a circle around the origin in the yz plane, of radius R. For any point $(x,0,0)$ let B_x be the closed ball with centre $(x,0,0)$ and radius $\sqrt{x^2+R^2}$. The boundary of its intersection with the yz-plane is the circle C. Then as x increases, that part of B_x to the right of the yz-plane is monotonically increasing, and that part to the left is monotonically decreasing, with respect to set inclusion.*

Proof. Let $(x,0,0)$ and $(x',0,0)$ be nearby points on the x-axis, and consider the rightmost point in B_x, $(x+\sqrt{x^2+R^2},0,0)$. Call this point p. It is on the positive x-axis, whether or not x is positive.

Suppose $x < x'$: claim p is interior to $B_{x'}$. In other words,

$$(x + \sqrt{x^2+R^2} - x')^2 < x'^2 + R^2.$$

Expanding, we need to show

$$x^2 + x^2 + R^2 + x'^2 + 2x\sqrt{x^2+R^2} - 2xx' - 2x'\sqrt{x^2+R^2} \quad < \quad x'^2 + R^2,$$

which is equivalent to

$$2x^2 - 2xx' + 2x\sqrt{x^2+R^2} - 2x'\sqrt{x^2+R^2} < 0$$

i.e.,

$$2(x-x')(\sqrt{x^2+R^2}+x) < 0.$$

The first factor is negative and the second positive, whether or not x is positive, so the result is true. Q.E.D.

(2.3) Corollary *Let L be a line and p a point, and Π the plane through p perpendicular to L. Suppose that L is parametrised in the form $a + t\vec{N}$ where $a \in L \cap \Pi$. If t is positive (negative) say that $a + t\vec{N}$ is* in front of *(respectively, behind) Π. Let B_t be the ball with centre $a + t\vec{N}$ and touching p.*

Then as t increases, that part of B_t in front of Π is expanding, and that part behind is contracting.

Proof. Without loss of generality L is the x-axis, $\vec{N}$ is in the positive x-direction, and p is in the yz-plane. If $p \neq 0$ then the result is a direct consequence of Lemma 2.2. Things are very slightly different if $p = O$, but the argument is essentially to be found in Lemma 2.1. Q.E.D.

(2.4) Lemma (trisector monotonicity principle.) *Given three compact convex sites B_i, the clearance sphere touching these three sites evolves as follows, as its centre x moves along the (B_1, B_2, B_3)-trisector.*

Let Π be the plane passing through the three points where the sphere touches the sites. It separates the clearance sphere into a front *and* rear *part, and as x moves along the trisector, the front part expands and the rear part contracts.*

Proof. We may assume that Π is the yz-plane and the forward direction is in the direction of increasing x. Let S be the sphere around x touching the three sites. Let $x + \Delta x$ be a nearby point on the trisector, where Δx is small. Let S' be the clearance sphere around $x + \Delta x$. It intersects Π (the yz-plane) but does not contain any point p_i in its interior, so the circle $S' \cap \Pi$ is inside the circle $S \cap \Pi$ (possibly not strictly).

Since the sphere S contains no points from any B_i in its interior, but S' touches these sites, S' is partly but not entirely inside S, and $S \cap S' \neq \emptyset$. S and S' have different centres, so $S \cap S'$ is either a single point or a circle. It cannot be a single point (where S' would touch S from the inside) since S' touches three different sites, so $S \cap S'$ is a circle C.

Without loss of generality $\Delta x > 0$. By the above lemma, to the right of the plane through C, S is inside S', and to its left, S' is inside S. Therefore S' touches the three sites to the right of this plane. Let Π' be the plane containing these contact points. Since S' intersects the yz-plane inside S, C is to the right of the yz-plane, i.e., to the right of Π. In other words, C is between the two planes Π and Π', so that part of S to the right of Π' is inside S' and that part of S' to the left of Π is inside S. Q.E.D.

(2.5) Remark about Voronoi vertices. The above monotonicity principle says something about how Voronoi edges meet Voronoi vertices. As a point moves along a trisector, consider how the clearance sphere intersects a fourth site. It 'meets' the site at the front and 'leaves' the site at the rear. This might lead one to suppose that the site could induce at most two Voronoi vertices on the trisector. That supposition is, however, false, as the following example shows.

Let C be the unit circle, centre O, in the xy-plane. If there is a point-site B_0 located at the origin, then the clearance spheres centred in C and touching the origin sweep out a torus-like solid T. We can easily place spherical sites B_1 and B_2, centred on the z-axis, so they intersect T tangentially in horizontal circles. Then C is the B_0, B_1, B_2-trisector.

One can find a horizontal disc D which intersects the solid T along its boundary, and does not intersect the other three sites B_i.

Let B_3 be a horizontal plane figure bounded by a regular k-sided horizontal polygon whose sides intersect the interior of D and whose corners are outside D (and intersect the solid T). Then as a point x moves around the circle C, its clearance sphere alternately intersects and does not intersect B_3 in $2k$ phases. These phases switch at (B_0, B_1, B_2, B_3)-vertices, so there are $2k$ such vertices. See Figure 3.

As a last variant of Lemma 2.4, for simplicity concerned only with spherical sites, we have

(2.6) Lemma *Suppose p is a point site and B_2 a spherical site with centre c_2 and positive radius. Let $t \mapsto f(t)$ be a differentiable curve on the (p, B_2)-bisector which is transverse to the contours of constant clearance: equivalently, df/dt is never perpendicular to pc_2. Let B_t be the ball with centre $f(t)$ touching p and B_2, and let Π_t be the plane through p perpendicular to df/dt.*

Then Π_t contains the point where B_t touches B_2 and as t increases that part of B_t in front of Π_t is expanding and that part behind is contracting.

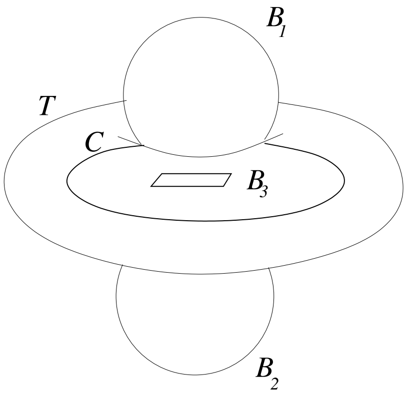

Figure 3: Four sites with $2k$ Voronoi vertices. (Site B_0 is at the origin.)

Proof. Much the same as Lemma 2.4. ∎

3 General position

Much of the paper up till now has been concerned with sets of disjoint compact convex sites. This section (but not the next) is mostly concerned with sets S of n disjoint spherical sites in $\mathbb{R}^3$.

It can be assumed that S is in general position, as specified below. First, a result about trisectors.

(3.1) Lemma *The trisector separating three spherical sites is a conic section, possibly a straight line.*

Proof. The trisector is where two bisectors intersect, or the cells owned by three sites. Without loss of generality, one of the sites is a point-site at the origin, radius zero, and the other two have centres c_1, c_2 respectively and radii r_1, r_2, respectively, possibly zero.

Points x on the trisector satisfy two equations

$$|x| = |x - c_1| - r_1 \quad \text{and} \quad |x| = |x - c_2| - r_2.$$

Therefore

$$|x|^2 + 2r_1|x| + r_1^2 = |x|^2 - 2c_1^T x + c_1^T c_1, \quad \text{so} \quad c_1^T c_1 - r_1^2 - 2c_1^T x = 2r_1|x|,$$

with a similar equation holding for the second bisector. If $r_1 = 0$ and $r_2 = 0$ then the trisector is a straight line. Otherwise, multiply the first equation by r_2 and the second by r_1 and subtract.

$$(r_2 c_1^T - r_1 c_2^T)x = r_2 c_1^T c_1/2 - r_2 r_1^2/2 - r_1 c_2^T c_2/2 + r_1 r_2^2/2 \tag{3.2}$$

Suppose $r_2 c_1^T - r_1 c_2^T$ is the zero vector but at least one radius, r_1 say, is nonzero. Then $c_2 = (r_2/r_1)c_1$, and the cone T with vertex at O and tangent to B_2 is also tangent to B_1. Without loss of generality c_1 is between O and c_2. Suppose y is a point on the (B_0, B_2)-bisector; let C be the ball centred at y and touching B_0 and B_2. Suppose it touches B_2 at z. The line-segment yz is entirely within the cone T and only meets the boundary at O, so it must intersect the interior of B_1, and the interior of the clearance sphere around y must also intersect the interior of B_1. Therefore y cannot be on the trisector: the trisector is empty.

We conclude that $r_2c_1 - r_1c_2 \neq O$ and the equation 3.2 defines a plane which contains the trisector. Since the trisector is then the intersection of a plane with either another plane or a quadric surface, it is a conic section. Q.E.D.

(3.3) Lemma *In the above lemma, if B_2 and B_3 have the same radius, then the trisector is in the bisector of their centres.*

Proof. If $r_2 = r_3 = 0$ then the (B_2, B_3)-bisector is the bisector of their centres. Otherwise equation 3.2 implies

$$(c_1 - c_2)^T x = (c_1 - c_2)^T (c_1 + c_2)/2.$$

This is again the equation of the plane bisecting the centres. Q.E.D.

(3.4) Lemma *Suppose that a (B_0, B_1, B_2) and a (B_0, B_1, B_3)-trisector have more than two points in common. Then the four sites have coplanar centres.*

Proof. Without loss of generality B_0 is a point site located at the origin. Let T_1 and T_2 be the (B_0, B_1, B_2)- and (B_0, B_1, B_3)-trisectors. If all the sites are point sites then T_1 and T_2 are straight lines which intersect in more than one point, hence coincide. Since these lines are perpendicular to the planes containing c_0, c_1, c_2 and c_0, c_1, c_3 respectively, the four point sites are coplanar and concyclic.

Suppose that one of the trisectors is a straight line. Since T_1 and T_2 intersect in more than two points, the other cannot be a non-linear conic section, so they are coincident straight lines, all four sites have the same radius, zero, and they are concyclic and coplanar point-sites.

Otherwise suppose $r_1 > 0$. T_1 and T_2 are in planes normal to the directions

$$r_1c_2 - r_2c_1 \quad \text{and} \quad r_1c_3 - r_3c_1$$

respectively (Equation 3.2), and these planes contain at least three non-collinear points in common, so they coincide and the normal directions are proportional. Since $r_1 \neq 0$ it follows that

$$c_2 = \alpha c_1 + \beta c_3 \tag{3.5}$$

for some α and β, so c_0, c_1, c_2, c_3 are coplanar (c_0 is the origin). Q.E.D.

(3.6) Definition *If an edge e is incident to four or more cells in $\mathrm{Vor}(S)$ (in its interior) it is called* degenerate. *A vertex is* degenerate *if it is a pinch-point or is incident to five or more cells in $\mathrm{Vor}(S)$.*

A set S of sites is in general position *if $\mathrm{Vor}(S)$ has no degenerate vertices and no four sites have coplanar centres.*

(By Lemma 3.4 it has no degenerate edges either.)

We shall argue that if $S = \{(c_1, r_1), \ldots, (c_n, r_n)\}$ specified by centres and radii, then the centres can be encoded in $\mathbb{R}^{3n}$, and if the centres are allowed to vary, then the degenerate placements satisfy polynomial equations and the set of nondegenerate placements is dense. Equivalently, the set of degenerate placements contains no open set. The following lemma is a direct corollary of Lemma 1 in [9], but we include a short proof:

(3.7) Theorem (polynomials are nonzero on a dense set). *Let $p(x_1, \ldots, x_N)$ be a nonzero polynomial defining a function from $\mathbb{R}^N$ to $\mathbb{R}$. Then the set $\{x \in \mathbb{R}^N : p(x) = 0\}$ has empty interior and its complement is open.*

Proof. The complement is open because the polynomial function is continuous.

If the set had nonempty interior it would contain a product of nonempty open intervals $\prod_1^N (a_j, b_j)$. We argue by induction on N. When $N = 1$, if p vanishes on a nonempty open interval, then it vanishes everywhere.

For the inductive step, suppose $p(x_1, \ldots, x_{N+1}) = \sum p_r(x_1, \ldots, x_N)x_{N+1}^r$ vanishes on a product of open intervals $\prod_1^{N+1}(a_j, b_j)$. For any $(x_1, \ldots, x_N) \in \prod_1^N (a_j, b_j)$, the polynomial $p(x_1, \ldots, x_{N+1}) = \sum p_r x_{N+1}^r$, a polynomial in one variable x_{N+1}, vanishes on (a_{N+1}, b_{N+1}), so all the p_r must vanish. So for each r, $p_r(x_1, \ldots, x_N)$ vanishes on $\prod_1^N (a_j, b_j)$. By induction all the polynomials p_r vanish identically, so p vanishes identically. Q.E.D.

(3.8) Given a set S of sites such that $\mathrm{Vor}(S)$ has degeneracies, we wish to perturb the centres but not the radii of the sites to produce a related set of sites in general position. For this reason we imagine the sets of sites being parametrised by the positions x_i of their centres, but leaving their radii fixed: a list of positions corresponds to a point in $\mathbb{R}^{3n}$ where $n = |S|$. Let $x_i = (x_{i1}, x_{i2}, x_{i3})$, so the positions correspond to $x_{11}, x_{12}, \ldots, x_{n3} \in \mathbb{R}^{3n}$.

The argument will be that degeneracies imply polynomial relationships among the $3n$ real variables $x_{i\ell}$ parametrising S. Mostly this will involve showing that a vertex closest to four given spherical sites satisfies some other polynomial equation, and invoking the following lemma.

(3.9) Lemma *Let B_1, B_2, B_3, B_4 be a set of sites whose centres x_i are allowed to vary, and let v be a vertex equidistant from these four sites; more generally suppose that we are given equations for two planes and one bisector whose intersection contains a given point v and possibly one more point.*

For any trivariate polynomial p, there is another polynomial q in the $x_{j\ell}$, whose coefficients may depend on the radii of these sites, such that $q(x_{11}, \ldots) = 0$ if and only if a point v equidistant from the sites B_i satisfies $p(v) = 0$.

Proof. The vertex v can be characterised as one of the (one or two) solutions to certain equations

$$\vec{N}_1 \cdot v = a_1, \vec{N}_2 \cdot v = a_2, b(v) = 0$$

where the first two equations define planes whose normals $\vec{N}_i$ are given by Equation 3.2, linear in the $x_{j\ell}$, and a_i are polynomials in the $x_{j\ell}$ — both involving the radii r_i also; and b is a linear or quadratic polynomial, with coefficients polynomials in the $x_{j\ell}$ and the radii, defining a bisector.

If b is linear, then we have three linearly independent equations, from which v can be expressed in terms of determinants according to Cramer's rule:

$$v = (\Delta_1/\Delta, \Delta_2/\Delta, \Delta_3/\Delta).$$

If p has degree k then $\Delta^k p(v)$ is a polynomial in the Δ_i and Δ: let $q(x_{j\ell}) = \Delta^k p(v)$, as required.

Otherwise b is a nonlinear, quadratic polynomial. The normals $\vec{N}_1$ and $\vec{N}_2$ are linearly independent, so if we set $\vec{N}_3 = \vec{N}_1 \times \vec{N}_2$ then $\vec{N}_1, \vec{N}_2, \vec{N}_3$ are linearly independent and we can calculate a unique point w such that

$$\vec{N}_i \cdot w = a_i \ \ (i = 1, 2) \quad \text{and} \quad \vec{N}_3 \cdot w = 0.$$

The coordinates of w are rational functions given by Cramer's Rule. The line L can be parametrised.

$$L = \{w + t\vec{N}_3 \colon t \in \mathbb{R}\}.$$

Find where L intersects the bisector $b(x) = 0$ by substituting $w + t\vec{N}_3$ into b, getting a quadratic equation

$$g(t) = 0.$$

The solutions have the form $a \pm b\sqrt{D}$, were a, b, and D are rational functions of the coefficients, and we get

$$v = w + (a \pm b\sqrt{D})\vec{N}_3$$

Substitute this into the polynomial p:

$$p(w + (a \pm b\sqrt{D})\vec{N}_3) = 0$$

This polynomial can be written in the form $c \pm d\sqrt{D}$,
Multiply both together

$$p(w + (a + b\sqrt{D})\vec{N}_3)p(w + (a - b\sqrt{D})\vec{N}_3) = 0$$

$$c^2 - d^2 D = 0.$$

This last is a polynomial in rational functions of determinants, so if multiplied by a suitable power of the denominators it takes the form $q(x_{j\ell}) = 0$, as required. Q.E.D.

A related result is the following:

(3.10) Lemma *Suppose that L is a line and B a bisector, specified by two linear and one quadratic polynomial equations as in the above lemma. Then there is a polynomial which vanishes if the line meets the bisector tangentially.*

Proof. According to the arguments of the above lemma, the points where the line intersects the bisector are of the form

$$w + (a \pm b\sqrt{D})\vec{N}_3.$$

If the intersection is tangential, the discriminant D is zero:

$$D = 0.$$

D is a rational function of the $x_{j\ell}$ (and the radii). If we multiply by a common denominator we obtain a polynomial equation $p(x_{11}, \ldots) = 0$. Q.E.D.

(3.11) Lemma *If S is a set of disjoint spherical sites not in general position, then there are placements arbitrarily close to S which are in general position.*

Proof. We consider every configuration as a point in $\mathbb{R}^{3n}$, specifying the centres of the sites, whose radii are left fixed. We express each degeneracy by one or more equations $p(x) = 0$. Since a finite list $p_r(x)$ of polynomials vanish at x if and only if $\sum_r (p_r(x))^2$ vanishes at x, the number of equations is irrelevant. It follows (Lemma 3.7) that every neighbourhood of a degenerate configuration contains a nondegenerate configuration.

If four centres x_i are coplanar then

$$\det(c_2 - c_1, c_3 - c_1, c_4 - c_1) = 0,$$

where det is the 3×3 determinant, a polynomial in the coordinates $x_{j\ell}$, which include the coordinates of the centres c_i.

We can assume that no four centres are coplanar.

Suppose that a vertex v is closest to five sites $B_0, \ldots, B_4$. Then v is closest to the first four sites, and also in the (B_0, B_4)-bisector, so by the above lemma a polynomial equation holds.

If v is a pinch-point, where a (B_0, B_1, B_2) and a (B_0, B_1, B_3)-trisector meet tangentially, then at most one of these trisectors is a straight line, because the site centres are not coplanar. Without loss of generality the (B_0, B_1, B_3)-trisector is not a straight line, B_0 is a point-site located at the origin, B_3 has positive radius, and the trisector is the intersection of tbe (B_0, B_3)-bisector with a unique plane Π,

Suppose L is a line tangent to the (B_0, B_1, B_3)-trisector at the point v. This trisector is the intersection of the plane Π with the (B_0, B_3)-bisector. The tangent line L is the intersection of Π with the tangent plane to this bisector at v. Hence L is contained in the tangent plane: in other words, L is tangent to the (B_0, B_3)-bisector.

If the (B_0, B_1, B_2)-trisector is a straight line then as argued above it is tangent to the (B_0, B_3)-bisector and the parameters $x_{j\ell}$ satisfy a polynomial equation (Lemma 3.10).

Otherwise the (B_0, B_1, B_2)-trisector is contained in a unique plane Π' which intersects Π in a line L (otherwise the four centres are coplanar), which is tangent to the (B_0, B_3)-bisector, and again the parameters $x_{j\ell}$ satisfy a polynomial equation. Q.E.D.

We have not discussed how selecting a configuration S' in general position near S affects the complexity of its Voronoi diagram. We should expect that $\mathrm{Vor}(S')$ should have at least as many faces as $\mathrm{Vor}(S)$. Then upper bounds on the complexity of $\mathrm{Vor}(S')$ would imply the same for $\mathrm{Vor}(S)$.

This is not necessarily so when $\mathrm{Vor}(S)$ contains pinch-points; removal of pinch-points may decrease the number of faces. The question will be reviewed in Section 9.

4 Descriptive complexity of diagrams

This section applies to sets of disjoint convex sites, not necessarily spherical. The basis of our $O(n^2)$ upper bound for the descriptive complexity of $\mathrm{Vor}(S)$ is the following lemma.

(4.1) Lemma *Let $\mathcal{C}$ be a collection of sets of disjoint spherical sites with the property that for any $S \in \mathcal{C}$, if $|S| > 1$ then any site can be removed from S to produce another member of $\mathcal{C}$.*

Suppose $M(n)$ denotes the maximum complexity of all diagrams $\mathrm{Vor}(S)$, $S \in \mathcal{C}$, $|S| = n$.

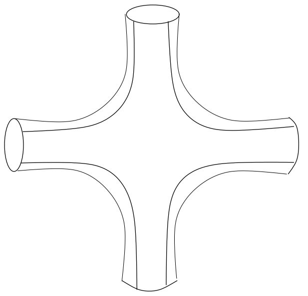

Figure 4: An unbounded cell might have high complexity.

Then $M(n)$ is $O(n^2)$ if and only if for every $S \in \mathcal{C}$, $\mathrm{Vor}(S)$ contains a cell of complexity $O(|S|)$.

Proof. Only if: suppose $M(n)$ is $O(n^2)$. Let S be a member of $\mathcal{C}$, $|S| = n$. Its complexity is at most $M(n)$. Its Voronoi diagram has n cells, so the average complexity of cells is at most $M(n)/n$ which is $O(n)$. But the average cell complexity bounds the minimum cell complexity.

If: suppose that a constant K exists such that whenever $S \in \mathcal{C}$ and $|S| = n$ then $\mathrm{Vor}(S)$ contains a cell of complexity $\leq Kn$. Given S, containing two or more sites, let $S' = S \backslash \{B\}$ where B is a site whose cell has complexity $\leq Kn$ in $\mathrm{Vor}(S)$. The complexity of $\mathrm{Vor}(S)$ is bounded by that of $\mathrm{Vor}(S')$ plus that of the cell owned by B in $\mathrm{Vor}(S)$. This yields a recurrence for $M(n)$;

$$M(n) \leq M(n-1) + Kn.$$

Also, $M(1) = 0 < K$, so $M(n) < Kn(n+1)/2$ is $O(n^2)$. Q.E.D.

(4.2) Bounded Voronoi diagrams. Choose a large solid cube K which contains in its interior all the sites in S and all the vertices and bounded edges, faces, and cells, of $\mathrm{Vor}(S)$, and let the cells, faces, and edges be intersected with this cube K. The resulting structure we call the *bounded Voronoi diagram.* Clearly its complexity is at least that of the unbounded Voronoi diagram. Bounding the Voronoi diagam simplifies Lemma 4.4 below, which would encounter difficulties with unbounded cells, whose 1-point compactifications could be of high genus.

For example, Figure 4 shows a cell generally in the shape of a cross whose four arms are supposed to go to infinity. It has been truncated, so there are four holes, but these holes and the edges and vertices bounding them should actually be at infinity. Four infinite edges (marked heavy) divide the cell boundary into two faces, and we could generalise the figure to a k-armed figure with k edges two faces, and no vertices. In any case the number of faces is always 2 and cannot bound the cell complexity.

We have not investigated whether such cells actually could arise in Voronoi diagrams of convex sites.

On the other hand, the truncated cell boundary has six faces, twelve edges, and eight vertices, satisfying Euler's formula.

The difficulty about unbounded cells does not arise with convex cells, for a rather trivial reason:

(4.3) Lemma *Having chosen a large solid cube K whose interior contains all sites, vertices, and bounded edges, faces, and cells of* $\mathrm{Vor}(S)$*, if a cell C is convex, then $C \cap K$ has at most six faces on ∂K, all of them simply connected.*

Proof. For each of the six faces F of ∂K, $F \cap C$ is convex, hence simply connected. Q.E.D.

(4.4) Lemma *If C is a cell of the bounded Voronoi diagram, whose boundary has $O(n)$ faces, then its complexity is $O(n)$.*

Proof. This should result from general planarity considerations, which say that a planar graph with m faces has $O(m)$ edges and vertices. We need to be slightly more careful because edges can be isolated, homeomorphic to circles, with no incident vertices. Let there be v vertices, e edges, and f faces ($f = n$). An edge might be incident to no vertex: in this case, since the cell is bounded, the edge is homeomorphic to a circle. There are unique faces meeting the circle from inside and outside respectively. Deleting the edge merges these faces. Thus if there are c 'circular' edges then deleting them reduces by c the number of faces.

Apart from isolated circular edges there may exist loops, that is, edges e with both ends incident to the same vertex v. Given a loop e, there is a face f meeting e from the inside. Choose two interior points u and w on e, and connect them by a new edge e' entirely within the face f (except at the endpoints). This splits the face into two, and replaces the single edge e by four edges, uv, vw, and two connecting uw. There are two new vertices. This operation increases the number of faces, edges, and vertices by $1, 3,$ and 2 respectively. Perform it as often as necessary to ensure that all vertices meet at least 3 different edges.

Let f_0, e_0, v_0 be the original number of faces, edges, and vertice; f_1, e_1, v_1 the numbers after removing isolated (circular) edges; $v_1 = v_0$ and $f_1 - f_0 = e_1 - e_0$. Let $f_2, e_2,$ and v_2 be the numbers after dealing with those vertices incident to one edge and one loop. Then

$$e_2 - e_1 = 3(f_2 - f_1), \quad \text{and} \quad v_2 - v_1 = 2(f_2 - f_1).$$

For any point x in the cell, if y is the closest point to x on the site owning the cell, then the line-segment xy is in the cell, so the cell is star-shaped relative to that site, hence contractible. It is also bounded, so its boundary is homeomorphic to a sphere. Euler's formula says

$$v_2 - e_2 + f_2 = 2$$

so $v_2 = e_2 - f_2 + 2$.

Every vertex is incident to at least three edges and every edge is incident to at exactly two vertices, following the alterations. The sum of the vertex degrees is twice the number of edges, and all vertex degrees are at least three. Hence

$$v_2 \leq \frac{2}{3} e_2.$$

Therefore $e_2 - f_2 + 2 \leq \frac{2}{3} e_2$, so $e_2 \leq 3f_2 - 6$. Since $e_1 - e_2 = 3(f_1 - f_2)$, $e_1 \leq 3f_1 - 6$. Since $e_0 - e_1 = f_0 - f_1$, $e_0 \leq f_0 + 2f_1 - 6 \leq 3f_0 - 6$ since $f_1 \leq f_0$. Therefore e_0 is $O(n)$.

The sum of all vertex degrees is at most twice the number of edges, so $v_0 \leq 2e_0$ is $O(n)$ also. Q.E.D.

(4.5) Corollary *Let C be a cell of* $\mathrm{Vor}(S)$ *with $O(n)$ faces. If C is convex, then it has complexity $O(n)$.*

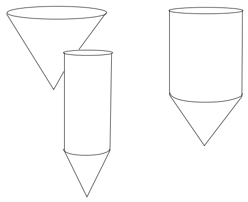

Figure 5: Lower-dimensional version of $S(r)$.

Proof. In the bounded Voronoi diagram enclosed by a solid cube K, $C \cap K$ has at most six more faces than C has (Lemma 4.3). Therefore $C \cap K$ has $O(n)$ faces and complexity $O(n)$ by the above lemma. Since the complexity of $C \cap K$ is at least that of C, the result follows. Q.E.D.

5 Inflating sites and fully general position

We are given a set S of n disjoint spherical sites. We begin with a set of point sites located at the centres of the sites in S, and call this set $S(0)$. We imagine the sites being 'inflated' to their correct size: an increasing parameter r is given, and $S(r)$ is the set of sites with their radius bounded by r. As r increases, the sites inflate until all have reached their correct radius. We study how the Voronoi diagram evolves.

(5.1) Definition *Let S be a set of n disjoint spherical sites. For any $r \geq 0$, if B is a site in S of radius R, then the* truncated site $B(r)$ *is the closed ball with the same centre as B and with radius* $\min(r, R)$.

The r-bounded version $S(r)$ *of S is the set of n sites*

$$\{B(r)\colon B \in S\}.$$

A site B is stable *in $S(r)$ if its radius is $\leq r$, otherwise it is* expanding.

(5.2) The parametrised set $S(r)$ can be viewed as a subset of $\mathbb{R}^4$, adding the parameter r in an extra dimension. If we consider the analogous situation with circular sites in two dimensions, a site of radius r_i sweeps out the union of a solid cone and a solid cylinder, meeting at height r_i. See Figure 5.

(5.3) Fully general position. The sets $S(r)$ of sites is not always in general position, even when S is. However, S is said to be in *fully general position* if

(i) S is in general position (3.6).

(ii) For each radius r_i, $S(r_i)$ is in general position.

(iii) If a vertex in $\mathrm{Vor}(S(r))$ has 5 closest sites, then not all of them are stable and not all of them are expanding.

(iv) For no r does $\mathrm{Vor}(S(r))$ contain a vertex closest to more than 5 sites.

(v) For no r does $S(r)$ contain a point which is simultaneously a pinch-point and a 5-site vertex.

(vi) For no five sites B_i, $1 \le i \le 5$, does the trisector separating the first three sites lie in the plane bisecting c_4c_5, the centres of the last two sites, or meet this plane tangentially.

(vii) For no five sites B_i, $1 \le i \le 5$, does the line L equidistant from c_3, c_4, c_5 meet the (B_1, B_2)-bisector at the point closest to c_1 on L.

(viii) If x is a pinch-point with respect to a point-site p and sites $B_2, B_3(r), B_4(r)$ in $\mathrm{Vor}(S(r))$ then the tangent T common to the edges is not parallel to the axis pc_2 of the (p, B_2)-bisector.

(ix) For no three sites B_1, B_2, B_3 is the centre c_1 equidistant from the centres c_2 and c_3.

As in the section on general position, a set of n spherical sites, whose radii are fixed, can be represented by a point in $\mathbb{R}^{3n}$ encoding the position of their centres.

(5.4) Definition *Let f be a feature (face,edge,vertex) of* $\mathrm{Vor}(S(r))$. *Its* type *is* $\mathrm{S}^a\mathrm{E}^b$ *where a and b are the numbers of stable and expanding sites closest to f.*

(5.5) Lemma *Given a set S of sites not in fully general position, every neighbourhood of S in $\mathbb{R}^{3n}$ contains configurations in fully general position.*

Proof. As in (Lemma 3.11), it is enough to show that nontrivial polynomial equations cover all configurations which violate any of the requirements for fully general position.

Let r_i be the set of radii of sites in S; without loss of generality the smallest radius is zero.

The requirement that S and all the sets $S(r_i)$ of sites are in general position is met by excluding configurations satisfying a polynomial equation (Lemma 3.11). This covers requirements (i) and (ii).

Let v be a vertex in $S(r)$ with 5 or more closest sites. If 5 sites are stable (case SSSSS) then v is closest to these sites in $\mathrm{Vor}(S)$, which is false. If 5 sites are expanding (case EEEEE), for any radius $r_i \le r$ among the sites in S, the 5 sites would be closest to v in $\mathrm{Vor}(S(r_i))$, which is also false. Thus requirement (iii) is met whenever the first two are met.

(iv) We consider the possibility that $S(r)$ contain vertices with 6 closest sites. As before, it is sufficient to show that when such vertices exist, the site centres satisfy polynomial equations independent of r. Cases SEEEEE and SSSSSE are excluded by (iii).

If four sites are expanding (case SSEEEE) then v is equidistant from their centres and

$$v = (\Delta_1/\Delta, \Delta_1/\Delta, \Delta_3/\Delta)$$

where Δ and Δ_i are polynomials in the centres x_i, independent of r, and v is on the bisector of two stable sites; hence we can deduce

$$\Delta \times b'(\Delta_1, \Delta_2, \Delta_3) = 0 \tag{5.6}$$

where b' is a polynomial with coefficients independent of r.

Suppose that four sites are stable and two are expanding (case SSSSEE). The one or two vertices closest to the four sites can be characterised by a bisector B and two planes P_1 and P_2, so $B \cap P_1 \cap P_2$ contains one or two points including v. The bisector of the two expanding sites is another plane P_3. B, P_1, P_2, and P_3 are all definable by equations independent of r.

Either $P_1 \cap P_2 \cap P_3$ is a line, and the normals are linearly dependent, or $P_1 \cap P_2 \cap P_3 = \{v\}$, and v can be substituted into the equation for B, yielding another polynomial equation like (5.6).

The final possibility is that three sites are stable and three are expanding (case SSSEEE). The trisector of the stable sites is an intersection $B \cap P_1$ where B is a bisector and P_1 a plane independent of r. The trisector of the expanding sites is an intersection $P_2 \cap P_3$ of planes definable by equations independent of r. Again either the three planes intersect in a line or they intersect at v alone, which also belongs to the bisector B, and again we derive a polynomial equation like (5.6).

Thus whenever there exists an r such that $\mathrm{Vor}(S(r))$ has vertices with six or more closest sites, the site centres satisfy polynomial equations whose coefficients are independent of r. Thus every configuration violating requirement (iv) satisfies a nontrivial polynomial equation.

(v) We need to show that a polynomial equation covers the situation where five sites $B_i(r)$ $(0 \leq i \leq 4)$ are closest to the same vertex v, which is also a pinch-point with respect to four of them.

When there is a pinch-point, there is a point v on a trisector whose tangent at v is perpendicular to all vectors $\vec{p_i p_j}$ where p_j are the points of contact with the clearance sphere at v (Lemma 1.18). Pinch points do not occur when r is one of the radii r_i, nor do four stable sites admit a pinch-point, since S is in general position, nor do four expanding sites, since their centres would be coplanar. Therefore we can choose p_i on a stable site and p_j on an expanding site. There are several cases depending on the type of the pinch-point (number of expanding sites involved).

Let B_0, B_1, B_2, B_3, B_4 be the sites, arranged in ascending order of their (true) radii.

In each case, we shall express v as one of the one or two points of intersection of a bisector with a line or a trisector with a plane, as required for Lemma 3.9, and furnish a displacement vector $\vec{N}$ parallel to the trisector tangent at v, and two contact points p_i, p_j where p_j is on an expanding site. All data except p_j will be independent of r.

$$\vec{N} \cdot \vec{p_i p_j} = 0, \quad \therefore \quad \vec{N} \cdot \vec{v p_j} = \vec{N} \cdot \vec{v p_i}.$$

Write $\vec{vp_i}$ as $|\vec{vp_i}|\vec{vc_i}/|\vec{vc_i}|$, similarly for $\vec{vp_j}$. But $|\vec{vp_i}| = |\vec{vp_j}|$, so if we clear the denominators we get

$$|\vec{vc_i}|\vec{N} \cdot \vec{vc_i} = |\vec{vc_j}|\vec{N} \cdot \vec{vc_j}.$$

This equation is independent of r, and if we square both sides we have a nonzero polynomial equation satisfied by v, and can apply Lemma 3.9.

In each case, therefore, we need only furnish equations for v and a vector $\vec{N}$ parallel to a trisector tangent at v. If the trisector is a straight line then it is the trisector of three site centres c_i, c_j, c_k and we can take $\vec{N} = \vec{c_i c_j} \times \vec{c_i c_k}$. Otherwise it is a conic section, a transverse intersection of a plane with a quadric bisector, and we can write $\vec{N} = \vec{N'} \times \nabla b(v)$, where $\vec{N'}$ is normal to the plane and b is a quadratic polynomial whose zeroes form the bisector; $\nabla b(v)$ is its gradient, which is normal to the tangent plane at v.

We distinguish the type of the vertex from the type of the pinch-point. For example, a pinch-point coinciding with a type SSSSE vertex must be of type SSSE since there are no SSSS pinch-points.

Case (v.a): SSSSE, with B_4 expanding, SSSE pinch-point. In this case v is closest to four stable sites (B_0, B_1, B_2, B_3), which furnishes equations for v independent of r. Suppose that B_1, B_2, B_3, B_4 define the pinch-point. All four of the trisectors for these sites have the same tangent at v (Lemma 1.18), so the (B_1, B_2, B_3)-trisector defines $\vec{N}$. Otherwise the pinch-point is defined by B_0, B_2, B_3, B_4, but the argument is the same.

Case (v.b): SSSEE $(B_3, B_4$ expanding), SSSE or SSEE pinch-point. In this case v is in the intersection of the (B_0, B_1, B_2)-edge with the plane bisecting c_3c_4, furnishing equations for v independent of r. If the pinch-point is SSSE then the first three sites furnish a trisector whose tangent defines $\vec{N}$. Otherwise the pinch-point is of type SSEE. The two edges meeting the point are of type SEE: (B_0, B_1, B_3)- and (B_0, B_1, B_4). By Lemma 3.3 these trisectors are contained in the c_3c_4 bisector plane, call it E, so their tangent is in E. They are also in the (B_0, B_1)-bisector, call it B, so their tangent is in the tangent plane at v to B. Thus the tangent is tangent to $E \cap B$. Both E and B are independent of r, and $\vec{N}$ can be defined using them.

Case (v.c): three expanding sites, B_2, B_3, B_4 (case SSEEE). In this case v is in the intersection of the (c_2, c_3, c_4)-trisector with the (B_0, B_1)-bisector. If the pinch-point involves the three expanding sites, the same trisector furnishes the vector $\vec{N}$. Otherwise it involves two expanding sites and two stable, and $\vec{N}$ can be derived from the intersection of the bisector plane separating the expanding sites with the bisector separating the stable sites, as in (v.b).

Case (v.d): four expanding sites, with non-coplanar centres (SEEEE). In this case v is the unique point equidistant from their centres, the pinch-point involves B_0, and $\vec{N}$ can be derived from the trisector separating three of the centres c_i, c_j, c_k $(i, j, k \geq 1)$.

(vi): Given five sites B_i, let T be the (B_0, B_1, B_2)-trisector and E the plane bisecting c_3c_4. If T is a straight line, it is perpendicular to the plane $c_0c_1c_2$, and if it is tangent to E, i.e., lies within E, then $(\vec{c_0c_1} \times \vec{c_0c_2}) \cdot \vec{c_3c_4} = 0$, a polynomial equation in the $x_{j\ell}$.

Otherwise T is the transverse intersection of a plane E' with the (B_0, B_1)-Bisector. If it lies within E then the planes E' and E coincide so

$$|\vec{N} \times \vec{N}'|^2 = 0,$$

a polynomial equation. Otherwise if it is tangent to E then $E' \cap E$ intersect in a unique line L and L meets the (B_0, B_1)-bisector tangentially. In this case there is a polynomial p such that $p(x_{j\ell}) = 0$ (Lemma 3.10).

(vii): let L be the line equidistant from c_2, c_3, c_4 and v the point on L closest to c_0. It is the orthogonal projection of c_0 onto L. Let $\vec{N}_1 = \vec{c_2c_3}$, $\vec{N}_2 = \vec{c_2c_4}$, and $\vec{N}_3 = \vec{N}_1 \times \vec{N}_2$. These vectors are linearly independent, and by solving equations

$$\vec{N}_1 \cdot w = a_1, \vec{N}_2 \cdot w = a_2, \vec{N}_3 \cdot w = 0$$

we get a point w on the line L, and $\vec{N}_3$ is parallel to the line. From these it is easy to calculate the projection v of c_0 onto L. The coordinates of v are rational functions of the $x_{j\ell}$. Then $b(v) = 0$ where b is a linear or quadratic equation defining the (B_0, B_1)-bisector, and by multiplying by a suitable polynomial to clear the denominators, we get a polynomial $p(x_{j\ell}) = 0$.

(viii): The tangent T is in the bisector of c_3c_4, so if (viii) is not true then $\vec{pc_2} \cdot \vec{c_3c_4} = 0$.

(ix): Otherwise $(c_1 - (c_2 + c_3)/2) \cdot (c_3 - c_2) = 0$, a polynomial equation. Q.E.D.

6 5-site vertices

In this section S is assumed to be a set of disjoint spherical sites in fully general position. As noted in Section 3, by proceeding to general position we may reduce the number of faces, but the assumption of general position will continue until this question is considered in Section 9.

(6.1) Definition *Let B_1 and B_2 be sites in S, r a parameter. A point x is* equidistant from B_1 and B_2 in $S(r)$ *(or in* $\mathrm{Vor}(S(r))$*) if x is equidistant from the truncated sites $B_i(r)$; similarly one speaks of x being closer to B_1 in $S(r)$, etcetera.*

For the rest of this paper, we consider a site of minimum radius in S, without loss of generality a point-site p.

(6.2) Definition *$C(r)$ is the cell containing p in* $\mathrm{Vor}(S(r))$.

(6.3) Lemma *$C(r)$ is always convex. (Immediate from Lemma 1.6.)* ∎

(6.4) In this section we consider how the face structure of $C(r)$ changes when $C(r)$ is incident to a 5-site vertex. The vertex will be labelled v, and the sites involved will be labelled p, B_2, B_3, B_4, and B_5: v is equidistant from the truncated sites $B_i(r)$.

It is assumed that these sites are listed so the stable ones are listed before the expanding ones. Since the sites are in general position, p is stable and B_5 is expanding.

Since S is in fully general position, $r \neq 0$. Parameters s and t will always be 'close to r' and less than r, and greater than r, respectively.

(6.5) Lemma *At an* SSSSE *vertex on $C(r)$ an* SSSS *vertex is removed and a new face is introduced to $C(r)$.*

Proof. Since S is in fully general position, r is not the radius of any site in S.

In $S(s)$, v is further from B_5 than from the other sites, so B_5's cell in $\mathrm{Vor}(S(s))$ is a positive distance from v, and there is no (p, B_5)-face near v. Note that v is an SSSS-vertex of $C(s)$.

In $S(t)$, v is closer to B_5, so the line-segment pv intersects the $(p, B_5(t))$-bisector at a point x. Since x is on the line-segment pv and v is equidistant from p, B_2, B_3, and B_4, in $S(t)$, x is closer to p than to B_2, B_3, and B_4, so it is on the interior of a (p, B_5)-face in $C(t)$. This is the face introduced to $C(r)$. The point v is a positive distance from $C(t)$: it is the SSSS-vertex removed. Q.E.D.

(6.6) Lemma *At an* SSSEE *vertex on $C(r)$ no new face can be introduced.*

Proof. Let e be the (p, B_2, B_3)-trisector, so it contains v, and let E be the plane equidistant from $B_4(r)$ and $B_5(r)$ (and from their centres). These last two sites are expanding and the others stable.

The only faces of $C(r)$ incident to v are closest to one of the four sites B_i, $2 \leq i \leq 5$. To prove that no new faces are introduced at v, it is enough to prove the following claim: if $s < r$ is close to r, then $C(s)$ has four faces close to v, one for each site B_i.

Since v is closer to p, B_2, and B_3 than to B_4 and B_5 in $S(s)$, there is a (p, B_2, B_3) edge containing v in $C(s)$.

The clearance sphere of v in $S(r)$ touches five sites. Consider the plane Π through the points of contact with p, B_2, and B_3. Since S is in fully general position and v is a 5-site vertex in $\mathrm{Vor}(S(r))$, it is not a pinch-point, and neither of the other two contact points of the clearance sphere are in the plane Π (Lemma 1.18).

If the clearance sphere touches the other two sites on the same side of Π, say that there is a unilateral contact, otherwise say there is a bilateral contact. Let e be the (p, B_2, B_3)-trisector.

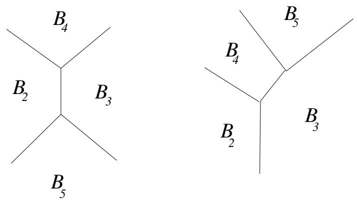

Figure 6: Evolution of diagram ($t > r$) from an SSSEE verex, bilateral and unilateral contacts respectively.

As a point v' moves away from v along e in $\mathrm{Vor}(S(r))$, on a side of Π containing the one or two contacts, the contact points become strictly closer to v' than to the sites p, B_2, B_3 (Lemma 2.4). This means that B_3 or B_4, or both if the contact point is unilateral, intersect the sphere around v' touching p, B_2, B_3. It follows that the edge containing v in $\mathrm{Vor}(S(s))$ (recall $s < r$ is close to r) is bounded by a vertex also closest to B_4 or B_5, both vertices existing in the case of a bilateral contact.

The case of a bilateral contact is simplest: the edge containing v and contained in e in $\mathrm{Vor}(S(s))$ meets a (p, B_1)-face and a (p, B_2)-face at v, and a (p, B_4) and a (p, B_5)-face at its bounding vertices. This proves the claim for a bilateral contact. See Figure 6.

We consider a unilateral contact. The trisector e meets the plane E separating $B_4(r)$ and $B_5(r)$ transversally at v, by requirement (vi) for fully general position (paragraph 5.3).

Writing $e = B \cap E'$ where B is a bisector and E' is a plane, consider a point x moving from v along e in $C(s)$ in the direction of these contact points. The edge e meets the plane E transversally, so x moves into one of the halfspaces bounded by E and away from the other: without loss of generality, into the halfspace containing B_4. The edge is bounded by a (p, B_2, B_3, B_4)-vertex, call it w, and w is incident to a (p, B_4)-face in $C(s)$. The edge meets the B_4-cell transversally at w by the trisector monotonicity principle (Lemma 2.4.) Hence w is nondegenerate, not a pinch-point (Lemma 1.18). Since w is incident to a (p, B_i)-face for $2 \leq i \leq 4$, it is sufficient to show that there is a (p, B_5)-face nearby.

There are three other edges incident to w, and one of them, at least, must be directed towards the plane E (Lemma 1.19). Suppose it is the (p, B_3, B_4)-edge. Following this edge towards E we reach a point on E, which is a (p, B_3, B_4, B_5)-vertex, a nondegenerate vertex incident to a (p, B_5)-face in $C(s)$.

In other words, if $s < r$ is sufficiently close to r then v is close to four faces of $C(s)$. This proves the claim in the case of a unilateral contact.

Since each of the four sites B_i contributes a face to $C(s)$ near v, no new faces can be introduced near v in $C(r)$. Q.E.D.

In the above argument it is unnecessary to speculate about when faces can be lost at an SSSEE vertex.

(6.7) Lemma *At an* SSEEE *vertex on* $C(r)$ *a face is removed from* $C(r)$ *and an* SEEE *vertex introduced.*

Proof. Again, let v be a vertex in $C(r)$ with five closest sites p, B_2, B_3, B_4, B_5, where the last three sites are expanding.

If $s < r$ is close to r then the clearance sphere of v in $\mathrm{Vor}(S(s))$ touches p and B_2 but not the other three sites. Therefore v is interior to a (p, B_2)-face.

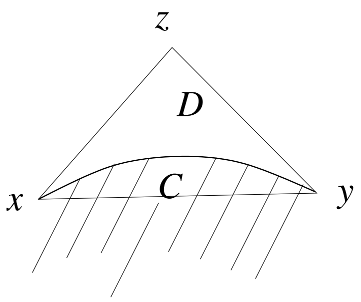

Figure 7: argument to show x and y belong to the same face of $C'(r)$.

If $t > r$ is close to r then the clearance sphere of v in $\mathrm{Vor}(S(t))$ touches B_3, B_4, B_5 but not p nor B_2. Thus v is on a (B_3, B_4, B_5)-trisector, which is a straight line L. By requirement (vii) for fully general position, v is not the point closest to p on L.

As a point x moves from v closer to p along L it enters the interior of $C(r)$ in $\mathrm{Vor}(S(r))$. Then it is interior to $C(t)$ for $t > r$ sufficiently close to r. Therefore there is a point x where L meets the boundary of $C(t)$, close to v, and x is at a positive distance from the B_2-cell. The vertex x is a (p, B_3, B_4, B_5)-vertex, where three faces meet, as claimed. Q.E.D.

(6.8) Lemma *With fixed point-site p and parameter r, let*

$$S'(r) = \{p\} \cup \{B(r)\colon B \in S, \quad \textit{radius} \quad \geq r\}.$$

Then the cell containing p in $\mathrm{Vor}(S'(r))$ contains at most one $(p, B(r))$-face for each $B \neq p \in S'(r)$. Hence $C(r)$ contains $O(n)$ SEEE- and SEEEE-vertices.

Proof. Let $C'(r)$ be the cell containing p in $\mathrm{Vor}(S'(r))$. Every SEEE vertex (and SEEEE-vertex) on $C(r)$ is a vertex of $C'(r)$, so it is enough to show that $C'(r)$ has $O(n)$ vertices. Since $C'(r)$ is convex (Lemma 6.3), by Corollary 4.5 it is enough to show that $C'(r)$ has at most $n-1$ faces.

Given a site B with $B(r) \in S'(r)\backslash\{p\}$, let $B' = B(r)$. Let x, y be points in the (p, B')-bisector, and strictly further from all other sites in $S'(r)$. We need to show that x and y belong to the same open face (Definition 1.2) of $C'(r)$.

Let D be the cell containing B' in $\mathrm{Vor}(S'(r)\backslash\{p\})$, a set of sites all of radius r, so D is a convex polytope, and x and y are interior to D. So is z, the centre of B'. Since D is convex, it contains the triangle xyz and its interior. This triangular region intersects the cell $C'(r)$ in a convex set containing the line xy but not containing z. See Figure 7

It intersects the boundary of $C'(r)$ in a curve connecting x to y in $C'(r)$. This boundary curve is interior to D, so it is entirely within the same open (p, B')-face. Hence there is at most one (p, B')-face. Q.E.D.

(6.9) Lemma *At an* SEEEE *vertex v on $C(r)$ either one* SEEE *vertex is removed and a new* SE *face is introduced, or one* SEE *edge is exchanged for another, but the number of faces is unchanged.*

Proof. The SEEEE vertex is a vertex of the convex cell $C'(r)$ discussed in the above Lemma 6.8, so the argument can be simplified by confining our attention to $C'(s)$ and $C'(t)$. Let Q be the clearance sphere around v in $\mathrm{Vor}(S'(r))$, touching p and touching the other sites B_i at points p_i. Since the centres c_i are not coplanar, neither are the points p_i ($p_i = v + \alpha \vec{vc_i}$, where $\alpha = |pv|/(|pv| + r)$, a relation which preserves coplanarity).

The plane $p_2p_3p_4$ divides the sphere Q into two regions, one containing p_5. In terms of the trisector monotonicity principle (2.4), one side is towards p_5 and one away from p_5. Similarly for $p_2p_3p_5$ and two other planes.

These four planes divide Q into ten open regions, one for each face and one for each edge of the tetrahedron $p_2p_3p_4p_5$. Corresponding to the face $p_2p_3p_4$ we have a region with a 'bias,' 'away from p_5' and 'towards' the other three points. This region is bounded by three arcs given by the $p_2p_3p_5$-, the $p_2p_4p_5$-, and the $p_3p_4p_5$-planes. We call this a +,+,+,- region to indicate the bias away from p_5.

Corresponding to the edge p_2p_3 we have a region bounded by two arcs where Q intersects the $p_2p_3p_4$- and the $p_2p_3p_5$-planes. Points in this region are biased towards p_2 and p_3 and away from p_4 and p_5. We call this a +,+,-,- region.

The two possibilities in this lemma depend on whether p is in a +,+,+,- or a +,+,-,- region. By requirement (ix) of fully general position, p is interior to one of these regions.

First note that if $s < r$ is close to r then the point on the line-segment vc_i equidistant from p and $B_i(s)$ is interior to a (p, B_i)-face. Hence $C'(s)$ contains (p, B_i)-faces close to v for $2 \leq v \leq 5$. We need only consider the situation for $t > r$ close to r.

Case (i): p is in a +,+,+,- region of Q. Without loss of generality, p is biased away from p_5 and towards the other three sites. The point v is outside $C'(t)$, being more remote from p than the other sites $B_i(t)$.

Because of the bias of p, as a point w moves away from v along the $(B_3(r), B_4(r), B_5(r))$-trisector, the clearance sphere (the sphere touching the four expanding sites, not necessarily p) recedes from p (trisector monotonicity principle 2.4). Therefore this trisector meets $C'(r)$ only at v. Similarly for the other two trisectors not involving $B_5(r)$. Therefore the polyhedral cell D containing $B_5(r)$ in $\mathrm{Vor}(c_2, c_3, c_4, c_5)$ intersects $C'(r)$ in exactly one point v. If $t > r$ then $C'(t) \subseteq C'(r)$, and $v \notin C'(t)$, so D does not intersect $C'(t)$. There is no $(p, B_5(t))$ face.

As a point w moves away from v along the $(B_2(r), B_3(r), B_4(r))$-trisector then its clearance sphere recedes from $B_5(r)$ and absorbs p in its interior (trisector monotonicity). One can argue as previously that if $t > r$ is sufficiently close to r then w meets a $(p, B_2(t), B_3(t), B_4(t))$-vertex. This covers case (i).

Case (ii): p is in a +,+,-,- region. Without loss of generality it is biased towards p_2 and p_3 and away from p_4 and p_5. By considering a moving point w as in the previous case, one concludes that the $(B_i(r), B_j(r), B_k(r))$-edge from v is locally interior to $C'(r)$ if it is 'opposite' B_4 or B_5 and meets the edge only at v if it is 'opposite' B_2 or B_3. Let E be the $(B_2(r), B_3(r))$ face incident to v in $\mathrm{Vor}(c_2, c_3, c_4, c_5)$. Thus E intersects $C'(r)$ in a single point and does not intersect $C'(t)$ if $t > r$, so there is no $(p, B_2(t), B_3(t))$-edge if $t > r$. But the $(B_2(t), B_3(t), B_4(t))$-trisector meets $C'(t)$ at a $(p, B_2(t), B_3(t), B_4(t))$-vertex; similarly there is a $(p, B_2(t), B_3(t), B_5(t))$-vertex. These bound a $(p, B_2(t), B_3(t))$-edge on $C'(t)$. For $2 \leq i \leq 5$ there is a (p, B_i)-face incident to one of these vertices.

Similarly if $s < r$ is sufficiently close to r there is a $(p, B_4(s), B_5(s))$-edge incident to all four faces at its interior and its endpoints. There is no $(p, B_2(s), B_3(s))$-edge. Q.E.D.

(6.10) Corollary *There are $O(n)$ SEEEE-vertices overall, for all r.*

Proof. At an SEEEE-vertex, either a new face is introduced, and that face is a face of $\mathrm{Vor}(S'(r))$ where $S'(r)$ is as in the above lemma, or the number of faces does not change.

Let $C'(r)$ be the cell containing p in $\mathrm{Vor}(S'(r))$, where $S'(r)$ is as in the above lemma. At an SEEEE vertex on $C(r)$, at most one new face is introduced to $C(r)$, and also to $C'(r)$, so it is the only face owned by some truncated site $B(t)$ in $C'(t)$, for $t > r$ close to r.

Next consider when a (p, B)-face is removed from $C'(r)$. As in the previous lemma, let D be the cell containing B in $\mathrm{Vor}(S'(r)\backslash\{p\})$. It intersects $C'(r)$ in a single point. Since $C'(r)$ continues to shrink, there can never again be a (p, B)-face. Thus a (p, B)-face is introduced at most once and there are $O(n)$ SEEEE-vertices incident to $C'(r)$. Q.E.D.

7 Face split-points

Recall that every trisector is a conic section.

(7.1) Lemma *Let x be a point on a (p, B_2, B_3)-edge of $\mathrm{Vor}(S(r))$, where B_2 has positive radius, smaller than the radius of B_3. There is a unique plane Π containing the (p, B_2, B_3)-trisector. The tangent to the trisector at x divides Π into two halfplanes of which one contains the trisector.*

Then near x, the tangent remains in that cell of $\mathrm{Vor}(S(r))$ containing B_3, the site with larger radius.

Proof. Let the lines connecting x to the centres of B_2 and B_3 meet their boundaries at p_2 and p_3, contact points of the clearance sphere. The tangent T is the (p, p_2, p_3)-trisector. Let B_2' be a 'shadow site' located inside B_3, so B_2' has the same radius as B_2 and it and B_3 share a common tangent-plane at p_3. Note that x is on the (p, B_2, B_2')-trisector, which has the same tangent T at x. The cell owned by p in $\mathrm{Vor}(p, B_2, B_3)$ is convex (Lemma 6.3), and T is tangent to it, so T only meets this cell at x. Therefore there exists an open ball V around x such that $V \cap T$ is contained in the union of the B_2- and B_3-cells, and also in the (B_2, B_2')-bisector.

Any point y close to x in $V \cap T$ is equidistant from B_2 and B_2'. The distance from y to B_2' is $|y - q_2|$, where q_2 is closest to y on B_2'. The line-segment yq_2 intersects the boundary of B_3 strictly between y and q_2, since B_2' is inside B_3. Therefore y is closer to B_3 than to B_2', hence closer to B_3 than to B_2 or to p. Q.E.D.

(7.2) Merge-points and split-points. Let x be a pinch-point relative to four sites (truncated to radius r) p, B_2, B_3, B_4, in $S(r)$ On the boundary of the p-cell $C(r)$ there are four faces incident to x.

Suppose the sites are labelled so the leftmost face incident to x is a (p, B_2)-face, the rightmost a (p, B_4)-face, and the two faces in between are (p, B_3)-faces. There are two incident edges, a (p, B_2, B_3)-edge on the left and a (p, B_3, B_4)-edge on the right.

It is possible that the two (p, B_3)-faces actually be the same face, being connected by paths not passing through x, but that is unimportant.

Since S is in fully general position, at least one site is expanding. It is possible that two sites are expanding, but B_2 and B_4 cannot both be expanding by Lemma 7.1, since the (p, B_2, B_3)-edge would meet its tangent at x from the right, and the (p, B_3, B_4)-edge would meet its tangent at x, the same line, from the left, whereas the first edge should be to the left of the second near x.

(7.3) Definition *The pinch point is classified as a* merge point *or a* split point *as follows:*

- *If B_3 is expanding but B_2 and B_4 are stable, it is a merge point.*

- *If two of the sites, B_3 and B_4, say, are expanding, it is a merge point.*
- *If B_3 is not expanding, so B_2 or B_4 is, it is a split point.*

(7.4) Lemma *At a pinch point x in $C(r)$, two (p, B_3)-faces merge or split according as x is a merge point or a split point.*

Proof. (See Figure 8.)

Case (a): only B_3 is expanding. For $s < r$ near r, x is equidistant from p, B_2 and B_4, but further from $B_3(s)$, so it is interior to an edge separating the (p, B_2)- and (p, B_4)-faces. Let e be the (p, B_2, B_4)-trisector. Near x, e is in the $(p, B_3(r))$-bisector, strictly interior except at x. Let u and v be two points in e, on either side of x, and close to x. The sphere centred at u, and touching B_2 and B_4, intersects the interior of $B_3(r)$, hence intersects the interior of $B_3(s)$ for all $s < r$ sufficiently close to r. Similarly for v. Therefore if $s < r$ is sufficiently close to r then x is closer to B_2 and B_4 than to $B_3(s)$, that is, x is interior to a (p, B_2, B_4)-edge, while u and v are on the same trisector but not on the edge; therefore, the edge has vertices in $\mathrm{Vor}(S(s))$ between u and v. These must be (p, B_2, B_3, B_4)-vertices, incident to two B_3-faces.

If $t \geq r$ is close to r then the same trisector is as close to $B_3(t)$ as to p, B_2 and B_4 near x, strictly closer when $t > r$. Then choosing u and v on the same trisector, for any w between u and v let $f_t(w)$ be the unique point equidistant from p and $B_3(t)$ on the line-segment wp. Notice that this line-segment is in the cell owned by p in the related Voronoi diagram $\mathrm{Vor}(p, B_2, B_4)$, so z is closer to p than to B_2 and B_4. Therefore f_t is a path connecting $f_t(u)$ to $f_t(v)$ in the $(p, B_3(t))$-bisector. If $t > r$ then the path is strictly interior to a $(p, B_3(t))$-face. Also, $f_r(u)$ and $f_r(v)$ are interior to the two B_3-faces in $C(r)$. By making t sufficiently close to r we can make f_t arbitrarily close to f_r, and in particular we can make $f_t(u)$ arbitrarily close to $f_r(u)$, similarly for $f_t(v)$. Thus the two (p, B_3)-faces in $C(r)$ merge together to one in $C(t)$.

Case (b): B_3 and B_4 are expanding, B_2 stable. Let e be the $(p, B_2, B_3(r))$-trisector. If $t > r$ then points on e are closer to $B_3(t)$ than to p and B_2. Taking intersections with line-segments as in Case (a), we get a path f_t of points in the $(p, B_3(t))$-bisector, closer to p than to B_2, and also closer to p than $B_4(t)$ because the line-segments do not cross the $(B_3(r), B_4(r))$-bisector, which coincides with the $(B_3(t), B_4(t))$-bisector. Thus we get a path strictly interior to a $(p, B_3(t))$-face, and this path connects points close to points in the two (p, B_3)-faces in $C(r)$.

We shall invoke Lemma 2.6. Let P be the plane through x normal to the edges' tangents at x in $C(r)$. By requirement (viii) for fully general position, P is not perpendicular to the axis pc_2, and its intersection with the (p, B_2)-bisector is a curve whose tangent is never perpendicular to that axis. Therefore this curve can be parametrised as a path $f(\tau)$ satisfying the requirement of Lemma 2.6. The pinch-point $x = f(\tau)$ for some τ; without loss of generality $x = f(0)$, and at $\tau = 0$ $df/d\tau$ is directed away from the plane through the two contact points and perpendicular to $df/d\tau$. Lemma 2.6 discusses the growth of balls B_τ centred at $f(\tau)$ and touching p (and $B_2(t)$).

As τ increases the two contact points of B_τ with $B_3(r)$ and $B_4(r)$ become interior to B_τ, so they enclose small open neighbourhoods of these contact points. For $s < r$ sufficiently close to r, these neighbourhoods intersect $B_3(s)$ and $B_4(s)$. Since the ball B_0 touching p with centre $f(0) = x$ does not, there exists a least $\tau' > 0$ such that $B_{\tau'}$ touches $B_3(s)$ or $B_4(s)$. Let $y = f(\tau')$.

At $t = 0$, df/dt is not parallel to the $(B_3(s), B_4(s))$-bisector (a plane), and it is directed towards $B_4(s)$ and away from $B_3(s)$. This holds for t sufficiently close to 0, so we can assume that y is

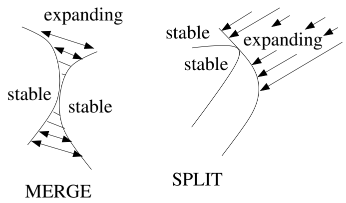

Figure 8: merge point and split point

closer to $B_4(s)$ than to $B_3(s)$, so it is interior to the $(p, B_2, B_4(s))$-edge in $C(s)$, at positive distance from $B_3(s)$. In other words, this edge exists in $C(s)$. As previously we can argue that $C(s)$ contains $(p, B_3(s))$ faces close to the $(p, B_3(r))$-faces. Thus we have shown that two (p, B_3)-faces merge at x.

Case (c): B_2 and B_3 are not expanding but B_4 is. The common tangent to the two edges bounds them on the right. Let e be the (p, B_2, B_3)-trisector. For $s < r$ sufficiently close to r, all points in e are closer to these three sites than to $B_4(s)$, and therefore the edge e bounds the same (p, B_3)-face near x in $C(s)$.

If $t > r$ is sufficiently close to r then x is closer to $B_4(t)$ than to the other three sites. Let e be the (p, B_2, B_3)-trisector. Using the trisector monotonicity principle (Lemma 2.4) we can find points u and v close to x and equidistant from p, B_2, B_3, and $B_4(t)$. They are vertices bounding a $(p, B_2, B_4(t))$-edge. Points on this edge are at positive distance from the two (p, B_3)-faces meeting u and v. Thus near x one (p, B_3)-face splits into two. Q.E.D.

(7.5) Transverse pinch-points. When a pinch-point v is on $C(r)$, but the face being pinched is not on $C(r)$, then there is an edge e on $C(r)$ containing v and another edge e' tangent to e at v. We call this a *transverse pinch-point*.

In comparing with pinch-points as previously discussed, the site p will play the rôle of B_2, and we shall suppose that the face being pinched is on a $(B_2(r), B_3(r))$-bisector. It is possible that B_2 be expanding. The fourth site $B_4(r)$ owns a cell $D(r)$ in $\mathrm{Vor}(S(r))$ touching $C(r)$

We take p, B_2, B_3, B_4 in that order and classify the pinch-point according to which sites are expanding: SSSE, SSES, SSEE, SESE, SEES, SEEE. The conclusions are as follows.

SSSE: Viewed from outside $C(r)$, the $(p, B_2, B_4(r))$- and $(p, B_3, B_4(r))$-faces are on opposite sides of their common tangent, so this is in fact a merge-point on $C(r)$. Two $(p, B_4(t))$-faces merge into one.

SSES: This is a merge-point in our previous classification. The edges e, e' get separated, $C(t)$ gets separated from $D(t)$, whereas for $s < r$ there is a $(B_2, B_4(s))$-face near v. A face is deleted.

SSEE: This is a merge-point in our previous classification. A face is deleted.

SESE: Same as SSEE.

SEES: Edges e and e' get separated. A face is deleted.

SEEE: Edges e and e' get separated. A face is deleted.

8 The upper bound

We consider the evolution of $C(r)$ as r increases (up to the maximum radius occurring in S).

$C(0)$ is a convex polyhedron with at most $n-1$ faces. As r increases, some of these faces become curved, and new faces appear and disappear.

(8.1) Lemma *Let S be parametrised as usual by the coordinates of its centres. If S is in general position, and S' is sufficiently close to it, then S' is in general position and $\mathrm{Vor}(S')$ is homeomorphic to S.*

Sketch proof. S' is in general position because the set of configurations in general position is open, since polynomials are continous. If S' is sufficiently close to S then one can set up correspondences between all features (sites, vertices, edges, faces) in each Voronoi diagram, and homeomorphisms connecting them. ▮

(8.2) Lemma *The structure of $C(r)$ changes only at 5-site vertices or at pinch-points or transverse pinch-points.*

Sketch proof. If $S(r)$ is in general position, then one can argue somewhat as in the previous lemma that $\mathrm{Vor}(S(t))$ is homeomorphic to $\mathrm{Vor}(S(r))$ for t close to r. Also, if $S(r)$ is degenerate but the degeneracies do not affect $C(r)$, then $C(r)$ is homeomorphic to $C(t)$.

Otherwise, $C(r)$ must include 5-site vertices, or pinch-points, or transverse pinch-points. Q.E.D.

(Recall that two faces cannot touch in their interiors (Corollary 1.12)). The face-structure of $C(r)$ can change as follows.

- SSSSE vertex: a face is introduced.
- SSSEE vertex: a face possibly deleted, none introduced.
- SSEEE vertex: a face is deleted.
- SEEEE vertex: a face possibly introduced, none deleted.
- Merge point: two faces merge.
- Split point: a face splits.
- Transverse pinch-point of any other kind: two faces merge.

We are concerned mostly with when the number of faces increases. This is at SSSSE vertices, at SEEEE vertices, and at split points. Overall, $O(n)$ faces can be introduced at SEEEE vertices (Corollary 6.10). When a face is introduced at an SSSSE-vertex, an SSSS vertex is deleted.

Let the radii of sites in S be listed

$$r_0 = 0 < r_1 < r_2 < \ldots < r_{k-1}.$$

We consider the evolution of $C(r)$ in $k-1$ phases. In the j-th phase $r_{j-1} < r < r_j$.

Suppose that an SSSS vertex is deleted in the j-th phase; by induction we can assume that $C(r_{j-1})$ has complexity $O(n)$. At an SSSSE vertex, an SSSS vertex is deleted and an SSSE face introduced; this can happen $O(n)$ times in the j-th phase.

It remains to bound the number of faces added at split points during the j-th phase. So long as such a face is incident to some vertex which can be counted, we have an upper bound. For this reason we count the evolution of vertices at various events affecting $C(r)$, as tabulated below. The $\pm$ notation indicates the number of vertices of various kinds gained and lost at the event. For example, at an SSSSE event, an SSSS vertex is lost and three SSSE vertices are gained. Events with the same label can have different possible outcomes.

Event type	Vertices gained/lost
SSSSE vertex	$-$ SSSS $+3$ SSSE
SSSEE vertex bilateral	-2 SSEE $+2$ SSEE
SSSEE vertex bilateral	-2 SSSE $-$ SSEE $+$ SSEE
SSSEE vertex unilateral	-2 SSSE $-$ SSEE $+$ SSEE
SSEEE vertex	-3 SSEE $+$ SEEE
SEEEE vertex	-3 SEEE $+$ SEEE
SEEEE vertex	-2 SEEE $+$ SEEE
SSSE merge point	ignore
SSEE merge point	ignore
transverse pinch	ignore
split point	$+2$ SSSE

The events which are ignored are events in which the numbers of faces, edges, and vertices all decrease. The table shows that the total number of vertices increases only for two events: SSSSE and split point. Each SSSSE event trades one SSSS vertex for three SSSE, so there are $O(n)$ in the j-th phase. The new vertices, created at SSSSE events and split points, are all of type SSSE. Let us call those SSSE vertices created at split points *split vertices*. These are the only vertices on $C(r)$ which are created without removing some other vertex, and the only faces which are not accounted for by considering incident vertices are those faces created by split events and bounded only by SSSE split vertices.

Consider when an SS face is split by an SE face, as in Figure 8. The face is bounded by the convex side of an SSS edge, and this edge touches an SSE edge. If the SSS edge bounds a (p, B_2) face on its concave side and a (p, B_3) face on the other, and the expanding face is a (p, B_4)-face, then

$$\text{radius}(B_2) < \text{radius}(B_3).$$

If there are more split faces (p, B_5), (p, B_6) ... along the boundary of F then

$$\text{radius}(B_2) < \text{radius}(B_3) < \text{radius}(B_5) < \text{radius}(B_6) \ldots$$

Therefore there cannot be more than j split faces occurring consecutively along the boundary of F. Thus the number of split faces adjacent to F is proportional (the constant of proportionality involves j, which is bounded) to the number of otherwise accounted vertices adjacent to F; hence there are $O(n)$ faces on $C(r)$. Summarising:

(8.3) Theorem *$C(r)$ has complexity $O(n)$, hence $\text{Vor}(S)$ has complexity $O(n^2)$.*

Proof. The complexity bound on $C(r)$ has been established as discussed with the aid of Corollary 4.5. The second part is immediate from 4.1. Q.E.D.

9 General position and complexity

We began with a set of n disjoint spherical sites, and selected a nearby set S' in general position, and then another nearby set in fully general position. If S' has as many faces as S then its complexity is at least that as S.

This is not necessarily the case. When removing 5-site vertices, or coplanar sets of centres, it can be shown that the number of faces cannot decrease (because the faces involved in $\mathrm{Vor}(S)$ are close to similar faces in $\mathrm{Vor}(S')$). However, this is not the case with pinch-points. When pinch-points are removed the nearby faces sometimes stay separate and sometimes merge. It must be ensured that faces do not merge.

Suppose x is a pinch-point involving a point-site p and sites $B_2, B_3, \ldots B_k$, $k \geq 4$, so the (p, B_i, B_{i+1})-edges all touch at x, leave the sites p, B_2, B_k fixed and move the other sites slightly away from p. Then x is interior to the (p, B_2, B_k)-edge and the number of faces $(2k-2)$ is preserved.

However, there may be other pinch-points which are displaced at the same time, and the total number of faces may be reduced. The difficulty can be met. It is not necessary that the total number of faces be preserved or increased, but the total number of faces incident to the p-cell in $\mathrm{Vor}(S')$, where p is a point site.

Among all the pinch-points on the p-cell in $\mathrm{Vor}(S)$, let P be the set of sites which contribute an outermost face to at least one pinch-point on this cell, and choose a site B of smallest radius in P. If we move that site slightly closer to p then all pinch-points involving B are removed in the way that we desired, i.e., the number of incident faces is preserved. Repeating the process we remove all the pinch-points from the p-cell.

Having removed all pinch-points involving faces on the p-cell, we transform a configuration S to a new one S' without any such pinch-points and with at least as many faces on the p-cell. Apply the usual arguments to obtain a configuration S'' in fully general position and 'close' to S' when considered as a point in $\mathbb{R}^3$. Since S'' is in fully general position, the line-segment $S'S''$ in $\mathbb{R}^3$ is not entirely in the zero-set of any of the polynomials whose vanishing is necessary to violate fully general position. In other words, this line-segment contains only finitely many sets of sites (encoded as points) not in fully general position. By moving S'' sufficiently close to S' on the same line-segment, we can ensure that all configurations on the line-segment $S'S''$, except perhaps S', are in fully general position. In particular, there are no pinch-points on the cell owned by p in any of the configurations between S and S''. It follows that in $\mathrm{Vor}(S'')$ the cell owned by p has at least as many faces as in $\mathrm{Vor}(S)$.

10 The bound on number of radii is essential

The folllowing construction shows that the if the number of distinct radii is unbounded, then the cell owned by a point site p can have quadratic complexity.

It is already known that cells in the Voronoi diagram of spheres can have quadaratic complexity, and the question is studied in all dimensions in [3]. However, we believe that this particular result, where the cell is owned by a *point site*, was hitherto unknown.

The construction begins in two dimensions. Let C be the circle with centre $(0, 1)$ and radius 1; p will be located at the origin, and C will be the clearance circle for the point $(0, 1)$, which will be a vertex of high degree.

(10.1) Lemma *The discs D_k with centres $(1/2^k, 0)$ and tangent to C are all disjoint.*

Proof. Note that for any $a \neq 0$, $(1 + a/2)^2 > 1 + a$, so $\sqrt{1 + a} < 1 + a/2$.

The radius of D_k is

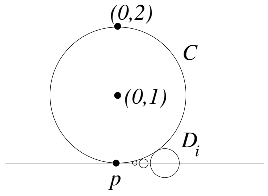

Figure 9: circle C touching p and n sites D_i.

$$\sqrt{1+\left(\frac{1}{2^k}\right)^2}-1<\frac{1}{2^{2k+1}}.$$

Discs D_k and D_ℓ with $k<\ell$ have centres separated by the distance $1/2^k - 1/2^\ell \geq 1/2^{k+1}$. The sum of their radii is less than $(1/4^k + 1/4^\ell)/2 < 1/4^k$. This is $\leq 1/2^{k+1}$ for all $k \geq 1$, so the circles are disjoint. Q.E.D.

Choose the discs $D_1, \ldots, D_n$ as sites in a 2-dimensional Voronoi diagram. Then the centre of C is a Voronoi vertex where the cell owned by p at the origin meets the n cells owned by the sites D_i. See Figure 9.

Now rotate the circle C and discs D_i around the x-axis. The discs D_i sweep out spherical sites B_i in $\mathbb{R}^3$. The centre of C sweeps out a circle E which is in fact the (degenerate) Voronoi edge closest to p and the n sites B_i. The highest point $(0,2)$ sweeps out a circle P.

The rotated circle C produces a (pinched) torus T of which P is the largest bounding circle. The torus T is also the region swept out by the clearance spheres centred on the circle E. If we place n point-sites $p_1, \ldots, p_n$ on the bounding circle P, as a point x sweeps around the circle E, then its clearance sphere touches each one of these sites in turn.

Now displace the points p_i slightly inside the torus, so that the point sites p_i now contribute n faces to the cell of p in the Voronoi diagram. The circle E is broken up into n Voronoi edges, still degenerate.

Finally, expand the sites B_i, $i = 1, \ldots, n-1$, so that the single circle E, a degenerate edge, splits into $n-1$ circular edges close to E. The faces contributed by the p_i split each of these circles into n edges, so the resulting Voronoi cell has more then $n(n-1)$ edges with $2n+1$ sites. In other words, it has quadratic complexity.

The general idea is illustrated in Figure 10.

11 Acknowledgement

The author is grateful to Pankaj Agarwal for useful pointers to the literature in this area.

12 References

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

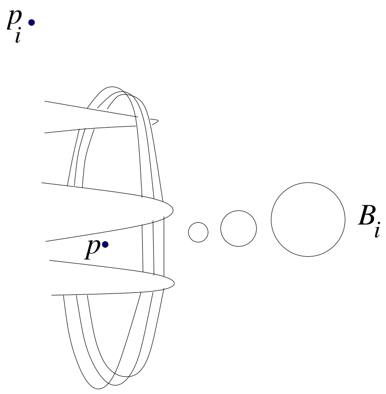

Figure 10: The final construction.

2. Franz Aurenhammer (1990). Voronoi diagrams— A survey of a fundamental geometric data structure. *ACM computing surveys* **23.3**, 345–405.

3. Jean-Daniel Boissonnat and Menelaos I. Karavelas (2003). On the combinatorial complexity of Euclidean Voronoi cells and convex hulls of d-dimensional spheres. *Proc. 14th Annual ACM-SIAM Symposium on Discrete Algorithms, Baltimore, Maryland, 2003,* 305–312.

4. Arne Brøndstred (1983). *An introduction to convex polytopes.* Springer Graduate Texts in Mathematics 90.

5. L. Paul Chew, Klara Kedem, Micha Sharir, Boris Tagansky, and Emo Welzl (1998). Voronoi diagrams of lines in 3-space under polyhedral convex distance functions. *Journal of Algorithms 29*, 238–255.

6. Andrew Farrell (1994). Fortune's Voronoi diagram algorithm extended to convex sites. Report TCDMATH 97–03. (M.Sc. dissertation, Mathematics Department, Trinity College, Dublin.)

7. Vladlen Koltun and Micha Sharir (2002). Three dimensional Euclidean Voronoi diagrams of lines with a fixed number of orientations. *Proc. 18th European Workshop on Computational Geometry*, 1–3.

8. Colm Ó Dúnlaing, Michar Sharir, and Chee K. Yap (1986). Generalized Voronoi diagrams for moving a ladder. I: Topological analaysis. *Communications on Pure and Applied Mathematics* **XXXIX**, 423–483.

9. J.T. Schwartz (1980). Fast probabilistic algorithms for verification of polynomial identities. *J. ACM* **27:4**, 701–717.

10. Michael Spivak (1998). *Calculus on Manifolds.* Addison-Wesley (First published Benjamin, 1965).