Chapter 7

Voronoi diagrams and weighted complexes

This chapter presents several complexes used in subsequent chapters. *Voronoi diagrams* are among the oldest geometric complexes studied; they are unbounded polyhedral complexes that associate each point in space with the closest site in a fixed set of sites—for instance, indicating for every point in town the nearest post office, as illustrated in Figure 7.1. Voronoi diagrams arise in the study of natural phenomena such as crystal formation, meteorology, computational chemistry, condensed matter physics, and epidemiology. Historically, they have been rediscovered several times and are also known as Dirichlet tessellations, Thiessen polygons, and Wigner–Seitz cells. They are the natural combinatorial duals of Delaunay subdivisions. In later chapters, they play important roles in ensuring the topological correctness of algorithms for meshing curved surfaces.

We also study *weighted* Voronoi diagrams, in which each site is equipped with a numerical weight; a larger weight allows a site to claim more territory. Weighted Voronoi diagrams are Voronoi diagrams in which the Euclidean distance is replaced with a nonmetric measure called the *power distance*. We have already introduced weighted Delaunay triangulations, the duals of weighted Voronoi diagrams, in Section 2.8. Here we explore their properties in further detail.

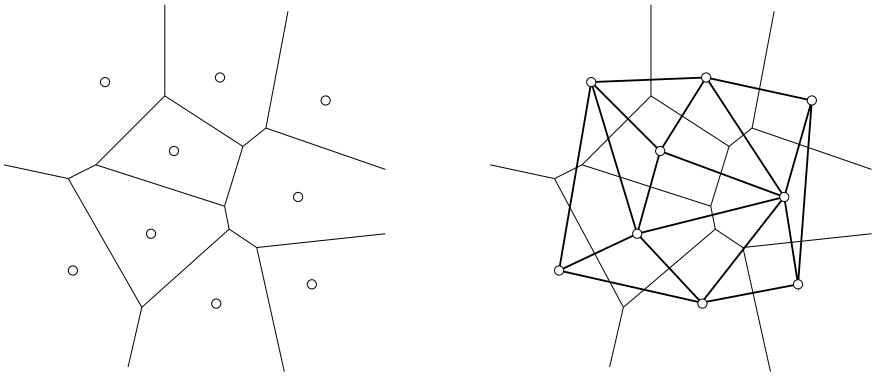

(a) (b)

Figure 7.1: (a) The Voronoi diagram of a planar point set. (b) With the Delaunay triangulation.

Lastly, we present a kind of subcomplex of a weighted Delaunay triangulation that we call a *quarantined complex*, which satisfies conditions that guarantee that the circumcenters of the simplices lie in the triangulation. Quarantined complexes generalize Proposition 6.2 to higher dimensions and to weighted Delaunay triangulations. We prove results about quarantined complexes called the Monotone Power, Orthoball Cover, and Orthocenter Containment Lemmas, and apply them in Chapters 8, 9, and 11.

7.1 Voronoi diagrams

Let S be a finite set of points, called *sites*, in $\mathbb{R}^d$. The *Voronoi cell* V_u of a site $u \in S$ is the set of all points in $\mathbb{R}^d$ that are at least as close to u as to any other site in S. Formally,

$$V_u = \{\, p \in \mathbb{R}^d : \forall\, w \in S,\ d(u,p) \leq d(w,p) \,\},$$

and u is said to be the *generating site* of V_u. Each Voronoi cell is a convex polyhedron, possibly an unbounded polyhedron. The faces of the Voronoi cells are called *Voronoi faces*. Voronoi faces of dimensions d, $d-1$, 2, 1, and 0 are called *Voronoi cells*, *Voronoi facets*, *Voronoi polygons*, *Voronoi edges*, and *Voronoi vertices*, respectively. All Voronoi faces are convex, and all except the vertices can be unbounded.

Definition 7.1 (Voronoi diagram)**.** The *Voronoi diagram* of S, denoted $\operatorname{Vor} S$, is the polyhedral complex containing the Voronoi cell of every site $u \in S$ and all the faces of the Voronoi cells.

Observe from the Voronoi diagram in Figure 7.1(a) that sites on the boundary of $\operatorname{conv} S$ have unbounded Voronoi cells. In the plane, two Voronoi edges of each unbounded Voronoi cell are rays—unless the sites are collinear, in which case all the Voronoi edges are lines.

Voronoi faces can be characterized by the sites that generate them. Define V_{uw} to be the set of points in the plane that have both sites u and w as their nearest neighbor. Formally,

$$\begin{aligned} V_{uw} &= V_u \cap V_w \\ &= \{\, p \in \mathbb{R}^d : d(u,p) = d(w,p) \text{ and } \forall\, x \in S,\ d(u,p) \leq d(x,p) \,\}. \\ V_{u_1 \ldots u_j} &= V_{u_1} \cap \ldots \cap V_{u_j} \\ &= \{\, p \in \mathbb{R}^d : d(u_1,p) = \cdots = d(u_j,p) \text{ and } \forall\, x \in S,\ d(u_1,p) \leq d(x,p) \,\}. \end{aligned}$$

If V_{uw} is nonempty, it is a face of the Voronoi diagram and of both Voronoi cells V_u and V_w, and u and w are said to be *Voronoi neighbors* and the *generating sites* or simply *generators* of V_{uw}.

If S is generic (recall Definition 4.2), a Voronoi face generated by j sites has dimension $d+1-j$. For example, in the plane, V_{uw} is an edge of two cells and V_{uwx} is a vertex of three cells. If S is not generic, the dimensionality could be lower or higher; for example, if four vertices in the plane lie in the circular order u, v, w, x on the boundary of an empty open disk, then $V_{uw} = V_{uwv} = V_{uwvx}$ is a vertex. For $j \geq 3$, the face might have any dimension less than or equal to $d-2$; for $j = 2$, any dimension less than or equal to $d-1$; for $j = 1$, it is always a d-face.

Each Voronoi facet lies on a bisector between two sites. For two distinct points $u, w \in S$, the *bisector* of u and w is the set of points at equal distance from u and w, or equivalently, the hyperplane orthogonal to the segment uw that intersects the midpoint of uw. Clearly, V_{uw} is a

subset of the bisector. The bisector cuts $\mathbb{R}^d$ into two halfspaces. Let H_{uw} be the closed halfspace containing u. Then

$$V_u = \bigcap_{w \in S \setminus \{u\}} H_{uw}.$$

Recall that Definition 1.5 defines a *convex polyhedron* to be the convex hull of a finite point set. For Voronoi diagrams, we define it more generally to be the intersection of a finite set of halfspaces, thereby permitting unbounded polyhedra.

For any finite point set S, the Voronoi diagram of S and the Delaunay subdivision of S defined in Section 4.2 are related to each other by a duality illustrated in Figure 7.1(b). A *duality* is a bijective map between the faces of one complex and the faces of another that reverses the property of inclusion; thus, if $\sigma \subset \tau$ are faces of the Delaunay subdivision of S, they map to dual faces $\sigma^* \supset \tau^*$ of the Voronoi diagram of S. Each k-face of the Delaunay subdivision is paired with a $(d-k)$-face of the Voronoi diagram. For example, each Delaunay vertex is paired with its Voronoi cell, and each Voronoi vertex is paired with a Delaunay d-face that is the union of all Delaunay d-simplices whose circumspheres are centered at that Voronoi vertex. The following theorem makes precise the nature of this duality.

Theorem 7.1. *Consider a site set $P = \{u_1, \ldots, u_j\} \subseteq S$. The polyhedron* $\operatorname{conv} P$ *is a k-face of the Delaunay subdivision of S if and only if $V_{u_1 \ldots u_j}$ is a $(d-k)$-face of* $\operatorname{Vor} S$ *and there is no site $u_{j+1} \in S \setminus P$ such that $V_{u_1 \ldots u_j} = V_{u_1 \ldots u_j u_{j+1}}$.*

P . If $\operatorname{conv} P$ is a face of the Delaunay subdivision, there exists a closed d-ball B that contains no site except $u_1, \ldots, u_j$, all of which lie on B's boundary. The Voronoi cell $V_{u_1 \ldots u_j}$ contains the center of B, so $V_{u_1 \ldots u_j}$ is nonempty and hence a face of $\operatorname{Vor} S$. Because every site in $S \setminus P$ is further from the center of B than the radius of B, there is no site $u_{j+1} \in S \setminus P$ such that $V_{u_1 \ldots u_j} = V_{u_1 \ldots u_j u_{j+1}}$.

For the reverse implication, let B be a closed d-ball whose center is an arbitrary point in the interior of $V_{u_1 \ldots u_j}$ and whose boundary passes through $u_1, \ldots, u_j$. No other site is in B, so $\operatorname{conv} P$ is a face of the Delaunay subdivision.

If $\operatorname{conv} P$ has dimension k, let Π be the $(d-k)$-flat that contains B's center and is orthogonal to $\operatorname{aff} P$. Every point on Π is equidistant from the sites in P, but no point in $\mathbb{R}^d \setminus \Pi$ is equidistant, so $V_{u_1 \ldots u_j} \subseteq \Pi$. Every point on Π whose distance from B's center is less than $d(B, S \setminus P)/2$ is closer to the sites in P than to any other site in S, and is therefore in $V_{u_1 \ldots u_j}$. It follows that $V_{u_1 \ldots u_j}$ has the same dimension as Π, namely $d-k$. □

To complete the duality, it is sometimes convenient to pair the unbounded outer face of the Delaunay triangulation with a "vertex at infinity" where every ray of the Voronoi diagram is imagined to terminate.

Thanks to duality, the bounds on the complexity of a Delaunay triangulation given in Sections 2.1 and 4.1 apply to Voronoi diagrams. An n-site Voronoi diagram in the plane has at most $2n-5$ vertices and $3n-6$ edges. An n-site Voronoi diagram in $\mathbb{R}^3$ has at most $(n^2-3n-2)/2$ vertices. Just as an n-vertex triangulation in $\mathbb{R}^d$ can have as many as $\Theta(n^{\lceil d/2 \rceil})$ d-simplices, an n-site Voronoi diagram can have up to $\Theta(n^{\lceil d/2 \rceil})$ Voronoi vertices.

To compute the Voronoi diagram of a set S of sites, first compute a Delaunay triangulation $\operatorname{Del} S$; Chapters 3 and 5 describe algorithms for doing so. Then compute the circumcenter of each d-simplex in $\operatorname{Del} S$. Where two or more adjoining d-simplices have the same circumcenter,

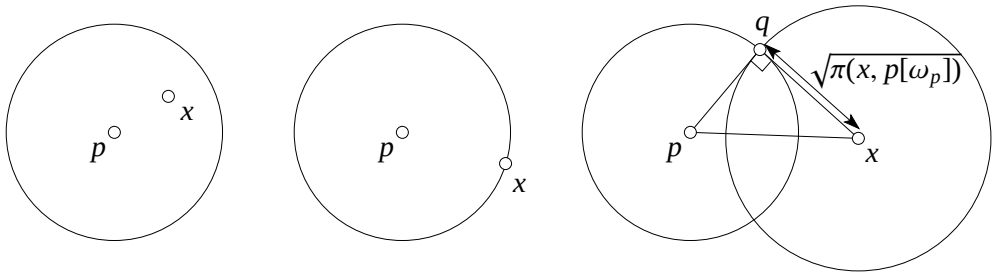

Figure 7.2: From left to right, the power distance $\pi(x, p[\omega_p])$ is negative, zero, and positive. The two balls at right are orthogonal, as $\pi(x[\omega_x], p[\omega_p]) = 0$.

fuse them into a polyhedron (having the same circumcenter), yielding the Delaunay subdivision. Finally, build the dual of the Delaunay subdivision. This last step might require no computation at all; a data structure that represents the Delaunay subdivision and the circumcenters implicitly represents the Voronoi diagram too.

7.2 Weighted Voronoi and weighted Delaunay

Just as Delaunay triangulations naturally generalize to weighted Delaunay triangulations, described in Section 2.8, Voronoi diagrams generalize to *weighted Voronoi diagrams*, the duals of weighted Delaunay subdivisions. Recall that a *weight assignment* $\omega : S \mapsto \mathbb{R}$ maps each site $u \in S$ to a weight ω_u. Negative weights are allowed. A *weighted point set* is written $S[\omega]$, and a *weighted point* is written $u[\omega_u]$.

Definition 7.2 (power distance)**.** The *power distance* between two weighted points $p[\omega_p]$ and $q[\omega_q]$ is

$$\pi(p[\omega_p], q[\omega_q]) = d(p, q)^2 - \omega_p - \omega_q.$$

We sometimes apply the notation to unweighted points, which by default have weight zero.

A weighted point $p[\omega_p]$ can be interpreted as a d-ball $B_p = B(p, \sqrt{\omega_p})$. For an unweighted point x outside B_p, the power distance $\pi(x, p[\omega_p])$ is the square of the length of the line segment that extends from x to touch B_p tangentially, illustrated in Figure 7.2. The power distance is zero if x lies on the ball's boundary, and negative if x lies inside the ball. We sometimes write the power distance $\pi(x, p[\omega_p])$ as $\pi(x, B_p)$. If the weight ω_p is negative, the ball radius $\sqrt{\omega_p}$ is imaginary, and every point lies outside the ball.

A weighted Voronoi diagram is defined like a Voronoi diagram, except that the Euclidean distance is replaced by the power distance. Observe that the larger a point's weight, the closer it is to other points, so a larger weight enables a site to claim a larger cell. The *weighted Voronoi cell* of a site $u \in S[\omega]$ is

$$W_u = \{\, p \in \mathbb{R}^d : \forall\, w \in S[\omega],\ \pi(p, u[\omega_u]) \leq \pi(p, w[\omega_w]) \,\}.$$

Counterintuitively, if u's weight is too small to compete with its neighbors, u can lie outside W_u, and W_u can even be empty (in which case u is submerged—not a vertex of the weighted Delaunay triangulation of S). Moreover, W_u is not necessarily a d-face; it can be of any dimension.

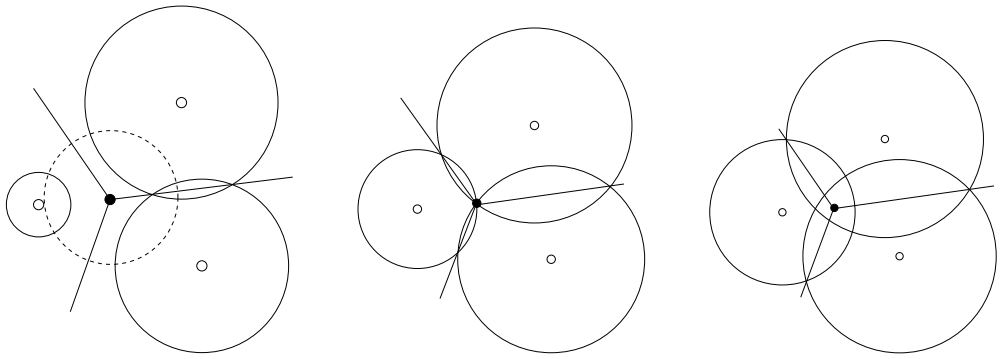

Figure 7.3: Three weighted point sets that have the same weighted Voronoi diagram, although the weights are different in each set. At left, the orthoball (dashed) centered at the weighted Voronoi vertex is orthogonal to all three weighted points. At center, the orthoball has radius zero. At right, the radius of the orthoball is imaginary.

Definition 7.3 (weighted Voronoi diagram)**.** The *weighted Voronoi diagram* of $S[\omega]$, denoted $\mathrm{Vor}\, S[\omega]$ and also known as the *power diagram*, is the polyhedral complex containing the weighted Voronoi cell of every site $u \in S$ and all the faces of the weighted Voronoi cells.

As a weighted point can be interpreted as a ball, we sometimes talk about the weighted Voronoi diagram of a set of balls. The faces of a weighted Voronoi cell are called *weighted Voronoi faces* and, like Voronoi faces, are characterized by their generating sites as

$$W_{u_1 \ldots u_j} = W_{u_1} \cap \ldots \cap W_{u_j}.$$

Unlike Voronoi faces, a weighted Voronoi face generated by two or more sites can have any dimension less than d.

Under the power distance, the bisector of $u[\omega_u]$ and $w[\omega_w]$ is still a hyperplane orthogonal to uw, but it intersects the midpoint of uw only if the points have equal weight; otherwise, it is further from the site with greater weight, and it might not intersect uw at all. As Figure 7.3 shows, if the balls centered at u and w intersect, the intersection of their boundaries lies on the bisector. The bisector cuts $\mathbb{R}^d$ into two closed halfspaces; let I_{uw} be the closed halfspace dominated by u. Then

$$W_u = \bigcap_{w \in S \setminus \{u\}} I_{uw}.$$

A fundamental property of the weighted Voronoi diagram is that there is a symmetric relationship between the sites and the vertices of the weighted Voronoi diagram, if we interpret both types of points as balls. Observe that a vertex v of the weighted Voronoi diagram is equidistant in power distance from the weighted points that generate it. Assign v a weight ϖ_v equal to the power distance from v to its generators. Then the power distance from $v[\varpi_v]$ to its generators is zero. Let B_v be the ball centered at v whose radius is $\sqrt{\varpi_v}$. We say that $v[\varpi_v]$ is *orthogonal* to its generators, and we call B_v the *orthoball* centered at v.

Definition 7.4 (orthogonal)**.** Two weighted points—equivalently, two balls—are *orthogonal* to each other if the power distance between them is zero. They are *further than orthogonal* if the power distance is positive. They are *closer than orthogonal* if the power distance is negative.

If the weights are positive, orthogonality implies that the boundaries of the balls intersect each other at right angles, as illustrated at right in Figure 7.2. To see this, consider weighted points p and x with associated balls B_p and B_x, and let q be a point on the boundary of both balls. If the balls meet at right angles, $\angle pqx$ is a right angle, and by Pythagoras' Theorem $\omega_p + \omega_x = d(p,x)^2$, so the power distance between $p[\omega_p]$ and $x[\omega_x]$ is zero.

Orthoballs generalize the notion of circumballs to simplices with weighted vertices. For example, a d-simplex with weighted vertices has a unique orthoball, namely the orthoball centered at the vertex of the weighted Voronoi diagram of the simplex's $d+1$ vertices, as illustrated at left in Figure 7.3. More generally, an orthoball of a simplex σ can be understood in two ways: as a ball that is orthogonal to σ's weighted vertices, or as the intersection of a hyperplane that includes the lifted simplex σ^+ with the paraboloid induced by the parabolic lifting map.

Definition 7.5 (orthoball)**.** Let σ be a k-simplex or a convex k-polyhedron with weighted vertices. An *orthoball* of σ is a ball (open or closed, as desired) that is orthogonal to σ's weighted vertices. In other words, the center c and radius $\sqrt{\varpi_c}$ of the orthoball satisfy $d(v,c)^2 - \omega_v - \varpi_c = 0$ for every vertex v of σ. Note that ϖ_c can be negative, in which case the radius is imaginary. An *orthosphere* of σ is the boundary of an orthoball of σ.

If σ has an orthoball (every simplex does, but not every polyhedron), there is one unique orthoball of σ whose center lies on the affine hull of σ, called the *diametric orthoball* of σ. The diametric orthoball's center is called the *orthocenter* of σ, its radius is the *orthoradius* of σ, and its intersection with $\operatorname{aff}\sigma$ is the *k-orthoball* of σ, which has the same center and radius as the diametric orthoball. An *orthodisk* is a 2-orthoball.

Analogously to circumballs, a k-simplex in $\mathbb{R}^d$ has exactly one k-orthoball, but if $k < d$ it has infinitely many d-orthoballs. The k-orthoball is a cross-section of every d-orthoball (see Lemma 7.4). Among all d-orthoballs, the diametric orthoball has minimum radius. Not every convex polyhedron has an orthoball, but every face of a weighted Delaunay subdivision does.

The alternative way to understand an orthosphere O of a simplex σ is to imagine lifting O to $\mathbb{R}^{d+1}$ with the parabolic lifting map, yielding the ellipsoid $O^+ = \{(p, \|p\|^2) : p \in O\}$, which is the intersection of a hyperplane with the paraboloid $p_{d+1} = \|p\|^2$. With this observation, every orthosphere induces a non-vertical hyperplane in $\mathbb{R}^{d+1}$ and vice versa. To see this, recall from the proof of Lemma 2.1 that if O has center o and radius $\sqrt{\varpi_o}$, the lifted ellipsoid O^+ lies on a hyperplane $h \subset \mathbb{R}^{d+1}$ whose formula is $p_{d+1} = 2o \cdot p - \|o\|^2 + \varpi_o$, where p varies over $\mathbb{R}^d$. By Definition 7.5, $\|v\|^2 - \omega_v = 2o \cdot v - \|o\|^2 + \varpi_o$ for every vertex v of σ. Recall that a weighted vertex v is lifted to a height of $\|v\|^2 - \omega_v$; it follows that $v^+ \in h$ for every vertex v of σ, so $\sigma^+ \subset h$.

This intuition fails for an orthoball of σ with imaginary radius, which induces a hyperplane that includes σ^+ but passes below the paraboloid and does not intersect it. But because imaginary radii are permitted, the one-to-one map from orthospheres to hyperplanes still holds. This map also provides an alternative way to understand the power distance; see Exercise 6.

Just as a simplex is Delaunay if it has an empty open circumball, a simplex is *weighted Delaunay* if it has an orthoball at a nonnegative power distance from every weighted vertex, or equivalently if it has a witness hyperplane (recall Definition 2.5). Observe that these are precisely the conditions in which its vertices $u_1, \ldots, u_j$ induce a face $W_{u_1 \ldots u_j}$ of the weighted Voronoi diagram.

Theorem 7.2. *Consider a weighted site set $P[\omega] = \{u_1, \ldots, u_j\} \subseteq S[\omega]$. The polyhedron* conv P *is a k-face of the weighted Delaunay subdivision of S if and only if $W_{u_1 \ldots u_j}$ is a $(d-k)$-face of* Vor $S[\omega]$ *and there is no site $u_{j+1} \in S \setminus P$ such that $W_{u_1 \ldots u_j} = W_{u_1 \ldots u_j u_{j+1}}$.*

P . Analogous to the proof of Theorem 7.1. □

This duality is a nearly-symmetric relationship, because every weighted vertex of a weighted Delaunay face of any dimension is orthogonal to every weighted vertex of its dual face in the weighted Voronoi diagram. Therefore, every weighted Voronoi d-cell W_u has an orthoball, namely the ball B_u induced by the weighted site $u[\omega_u]$. The only break in the symmetry arises from the unbounded cells of the weighted Voronoi diagram. If $U[\varpi]$ is the set of weighted vertices of the weighted Voronoi diagram Vor $S[\omega]$, then the weighted vertices of the weighted Voronoi diagram Vor $U[\varpi]$ include the original sites in $S[\omega]$ except those on the boundary of conv S.

A useful fact is that the orthocenter of a k-face σ of a weighted Delaunay subdivision is the point aff $\sigma \cap$ aff σ^*, where σ^* is σ's dual face in the weighted Voronoi diagram. Because every weighted Voronoi diagram is the dual of a weighted Delaunay subdivision, the complexity bounds given for Voronoi diagrams in Section 7.1 apply to weighted Voronoi diagrams as well.

7.2.1 Properties of orthoballs

This section proves some geometric properties of orthoballs, the power distance, and weighted Delaunay triangulations. The Delaunay triangulation is a special case of the weighted Delaunay triangulation, so the propositions that follow apply to ordinary Delaunay triangulations as well.

The *orthogonal projection* of a point p onto a flat or a linear cell $g \subseteq \mathbb{R}^3$ is the point $\tilde{p} \in$ aff g such that the line segment $p\tilde{p}$ is perpendicular to the affine hull aff g.

An interesting property of weighted Voronoi diagrams is that any cross-section of a weighted Voronoi diagram is a lower-dimensional weighted Voronoi diagram of a modified set of sites that lie in the cross-section. The modified sites are found by orthogonally projecting the sites onto the cross-sectional flat and adjusting their weights as described in the following proposition.

Proposition 7.3. *Let $\Pi \subset \mathbb{R}^d$ be a flat. Let $p \in \mathbb{R}^d$ be a point with weight ω_p. Let $\tilde{p}$ be the orthogonal projection of p onto Π, with weight $\omega_{\tilde{p}} = \omega_p - d(p, \tilde{p})^2$. For every point $x \in \Pi$, $\pi(x, p[\omega_p]) = \pi(x, \tilde{p}[\omega_{\tilde{p}}])$.*

P . $\pi(x, \tilde{p}[\omega_{\tilde{p}}]) = d(x, \tilde{p})^2 - \omega_{\tilde{p}} = d(x, \tilde{p})^2 + d(\tilde{p}, p)^2 - \omega_p = d(x, p)^2 - \omega_p = \pi(x, p[\omega_p])$. □

If $\omega_{\tilde{p}}$ is nonnegative, the projected site $\tilde{p}[\omega_{\tilde{p}}]$ can be interpreted as $B(\tilde{p}, \sqrt{\omega_{\tilde{p}}})$, the smallest d-ball that includes the cross-section $B(p, \sqrt{\omega_p}) \cap \Pi$. If $\omega_{\tilde{p}}$ is negative, Proposition 7.3 still applies, but the "cross-section" has imaginary radius. The same principle applies not only to weighted sites, but also to orthoballs. The following proposition shows that cross-sections of an orthoball of a simplex indicate the diametric orthoballs of its faces, as illustrated in Figure 7.4(a).

Lemma 7.4 (Orthoball Restriction Lemma)**.** *Let τ be a simplex with weighted vertices in $\mathbb{R}^d$. Let $B(c, \sqrt{\varpi_c})$ be an orthoball of τ, and let σ be a face of τ. Let $\tilde{c}$ be the orthogonal projection of c onto* aff σ*. Then the orthocenter and orthoradius of σ are $\tilde{c}$ and $\sqrt{\varpi_{\tilde{c}}}$, respectively, where $\varpi_{\tilde{c}} = \varpi_c - d(c, \tilde{c})^2$, and $B(\tilde{c}, \sqrt{\varpi_{\tilde{c}}})$ is the diametric orthoball of σ.*

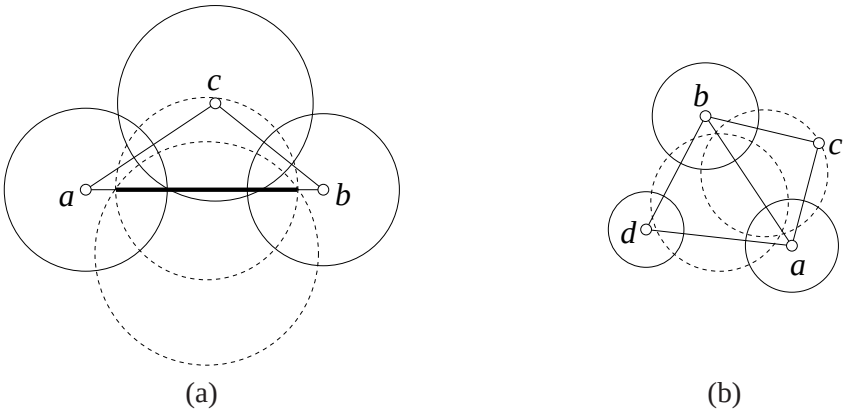

(a) (b)

Figure 7.4: (a) The intersection of edge *ab* and the orthoball of *abc* (the bigger dashed circle) is an interval (bold). The smaller dashed circle indicates the diametric ball of this interval and the diametric orthoball of *ab*. (b) The bisector of the orthoballs of *abc* and *abd* is the line that includes the edge *ab*.

P . By the definition of orthoball, $B(c, \sqrt{\varpi_c})$ is orthogonal to the weighted vertices of τ. By Proposition 7.3, $B_{\tilde{c}} = B(\tilde{c}, \sqrt{\varpi_{\tilde{c}}})$ also is orthogonal to the weighted vertices of σ, and is therefore an orthoball of σ. Because $\tilde{c} \in \operatorname{aff} \sigma$, $B_{\tilde{c}}$ is the diametric orthoball. □

Just as every weighted Voronoi facet lies on a bisector of two weighted points, every weighted Delaunay facet not on the boundary of the triangulation lies on a bisector of two orthoballs of d-simplices, as illustrated in Figure 7.4(b). See the proof of Lemma 7.5 for confirmation.

7.3 Quarantined complexes

In our presentation of Ruppert's algorithm, Proposition 6.2 shows that when no subsegment is encroached, every triangle's circumcenter lies in the domain, so the algorithm can safely insert the circumcenter of any skinny triangle. In this section, we generalize the proposition to higher-dimensional triangulations and weighted Delaunay triangulations to support the Delaunay refinement algorithms in this book. We study triangulations called *quarantined complexes* that have the following desirable properties.

Definition 7.6 (quarantined complex)**.** Let $\mathcal{P}$ be a PLC in $\mathbb{R}^d$. let $S \subset |\mathcal{P}|$ be a finite point set that includes the vertices of $\mathcal{P}$, and let ω be a weight assignment such that a subcomplex $\mathcal{Q}$ of $\operatorname{Del} S[\omega]$ is a Steiner triangulation of $\mathcal{P}$. The *dimension* k of both $\mathcal{P}$ and $\mathcal{Q}$ is the dimension of their highest-dimensional cell, which is not necessarily d. A j-simplex in $\mathcal{Q}$ is called a *boundary simplex* of $\mathcal{Q}$ if $j < k$ and it is included in a linear j-cell in $\mathcal{P}$. We call $\mathcal{Q}$ a quarantined complex if it satisfies the following conditions.

(i) The dimension of aff $|\mathcal{Q}|$ is equal to the dimension of $\mathcal{Q}$.

(ii) Every vertex in $\mathcal{Q}$ has nonnegative weight.

(iii) The power distance between every pair of vertices in $\mathcal{Q}$ is nonnegative.

(a) (b) (c)

Figure 7.5: (a) A quarantined complex. The weighted vertices are represented by balls. (b) This one-dimensional complex is not a quarantined complex because it has dimension one and its affine hull has dimension two. (c) This two-dimensional complex is not a quarantined complex because two weighted vertices, represented by the balls at upper left, are closer than orthogonal, and the diametric orthoball of the boundary simplex at right is closer than orthogonal to a vertex.

(iv) For every boundary simplex σ of $\mathcal{Q}$ and every vertex v in $\mathcal{Q}$, the power distance between $v[\omega_v]$ and the diametric ball B_σ of σ is nonnegative, i.e. $\pi(v[\omega_v], B_\sigma) \geq 0$.

Figure 7.5 illustrates a quarantined complex and some complexes that are not quarantined. Note that $\mathcal{P}$ is sometimes a subcomplex of a larger PLC; perhaps $\mathcal{P}$ contains a linear cell and its faces. Observe that if every vertex has weight zero, condition (iii) is trivially satisfied, as $\pi(u,v)^2 = d(u,v)^2 > 0$, and condition (iv) requires each boundary simplex to have an open diametric ball that contains no vertex. In Ruppert's algorithm, condition (iv) holds when no subsegment is encroached.

Definition 7.7 (balls in a quarantined complex)**.** For any quarantined complex $\mathcal{Q}$ of dimension k in $\mathbb{R}^d$, let Ball($\mathcal{Q}$) denote the set of closed balls that contains $B(v, \sqrt{\omega_v})$ for every vertex v in $\mathcal{Q}$, the diametric orthoball of every k-simplex in $\mathcal{Q}$, and the diametric orthoball of every boundary simplex of $\mathcal{Q}$ that is not a vertex.

Figure 7.6 illustrates the balls in Ball($\mathcal{Q}$) for one quarantined complex. No vertex in $\mathcal{Q}$ is closer than orthogonal to any ball in Ball($\mathcal{Q}$) except that vertex's own ball: this is true for vertex balls by condition (iii), for diametric orthoballs of boundary simplices by condition (iv), and for diametric orthoballs of k-simplices because $\mathcal{Q}$ is a Steiner weighted Delaunay triangulation.

7.3.1 The Monotone Power Lemma

The Monotone Power Lemma states that in a k-dimensional quarantined complex, the diametric orthoballs of the k-simplices and boundary simplices that intersect an arbitrary ray have power distances from the origin of the ray that increase monotonically along the ray, as illustrated in Figure 7.7. The result also holds for the power distances from a point whose orthogonal projection onto the affine hull of the quarantined complex is the origin of the ray. We use the lemma to prove the correctness of mesh generation algorithms in Chapters 8, 9, and 11. To apply the lemma, let $\mathcal{Q}$ be a quarantined complex in $\mathbb{R}^d$, let Π be the affine hull of $|\mathcal{Q}|$, and let σ and τ be simplices in $\mathcal{Q}$.

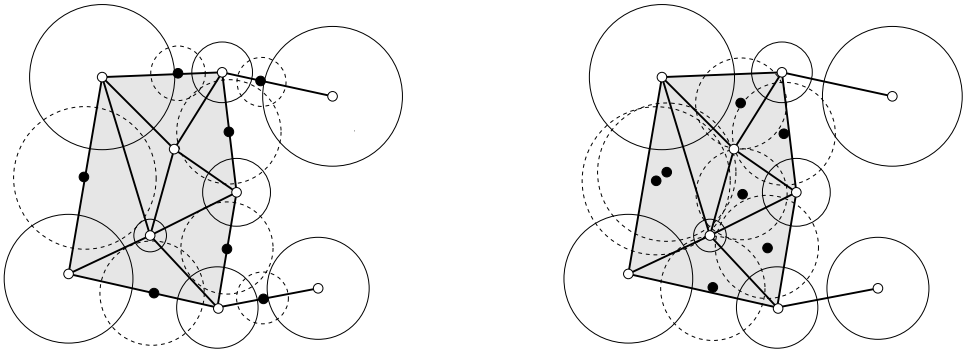

(a) (b)

Figure 7.6: Balls in a quarantined complex. (a) The weighted vertices (solid circles) and the diametric orthoballs of the boundary edges (dashed). (b) Orthoballs of triangles (dashed).

Lemma 7.5 (Monotone Power Lemma)**.** *Let Π be a flat in $\mathbb{R}^d$. Let $\vec{\gamma} \subset \Pi$ be a ray that intersects two simplices $\sigma, \tau \subset \Pi$ in that order—meaning that the ray strikes a point in σ then a distinct point in τ. (Either point may be in $\sigma \cap \tau$, and one simplex may be a face of the other.) Let B_σ and B_τ be orthoballs of σ and τ whose centers lie on Π. Suppose that no vertex of σ is closer than orthogonal to B_τ and no vertex of τ is closer than orthogonal to B_σ. Let $p \in \mathbb{R}^d$ be an unweighted point whose orthogonal projection onto Π, written $\tilde{p}$, is the source of $\vec{\gamma}$. Then $\pi(p, B_\sigma) \leq \pi(p, B_\tau)$.*

P . If $B_\sigma = B_\tau$, then $\pi(p, B_\sigma) = \pi(p, B_\tau)$ and the result follows, so assume that $B_\sigma \neq B_\tau$. Because $p\tilde{p}$ is orthogonal to Π, $\pi(p, B_i) = d(p, \tilde{p})^2 + \pi(\tilde{p}, B_i)$ for $i \in \{\sigma, \tau\}$, and it suffices to prove that $\pi(\tilde{p}, B_\sigma) \leq \pi(\tilde{p}, B_\tau)$.

Let $H = \{x \in \mathbb{R}^d : \pi(x, B_\sigma) = \pi(x, B_\tau)\}$ be the bisector of B_σ and B_τ with respect to the power distance. Every vertex of σ is orthogonal to B_σ and orthogonal or further than orthogonal to B_τ, and therefore lies in the closed halfspace $\{x \in \mathbb{R}^d : \pi(x, B_\sigma) \leq \pi(x, B_\tau)\}$ bounded by H. It follows that every point in σ lies in the same halfspace. Symmetrically, every point in τ lies in the closed halfspace $\{x \in \mathbb{R}^d : \pi(x, B_\sigma) \geq \pi(x, B_\tau)\}$. Because $\vec{\gamma}$ intersects a point in σ then a point in τ, its origin $\tilde{p}$ is in the former halfspace—that is, $\pi(\tilde{p}, B_\sigma) \leq \pi(\tilde{p}, B_\tau)$. □

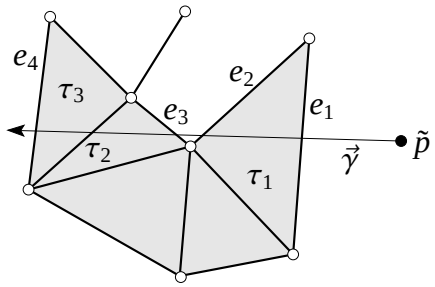

Figure 7.7: A demonstration of the Monotone Power Lemma. The ray $\vec{\gamma}$ crosses the simplices e_1, τ_1, e_2, e_3, τ_2, τ_3, e_4 in that order. The power distances from $\tilde{p}$ to their diametric orthoballs increase monotonically in that order.

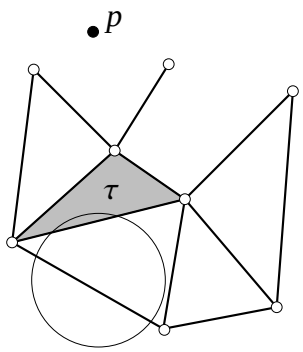

Figure 7.8: All vertices are unweighted. A weighted point $x[\varpi_x]$ is drawn as a circle, at a positive power distance from every vertex. The power distance from p to the circumball of the triangle τ is no greater than the power distance from p to $x[\varpi_x]$.

7.3.2 The Orthoball Cover Lemma

The balls of a quarantined complex $\mathcal{Q}$ cover its underlying space—every point in $|\mathcal{Q}|$ is in some ball in $\text{Ball}(\mathcal{Q})$. More generally, suppose that a weighted point $x[\varpi_x]$ in $|\mathcal{Q}|$ has nonnegative power distances from the weighted vertices of a quarantined complex $\mathcal{Q}$. Then for any point p in space, there is a ball in $\text{Ball}(\mathcal{Q})$ at least as close to p in power distance as $x[\varpi_x]$ is, as illustrated in Figure 7.8. (See Exercise 10 for the relationship between these two claims.)

The Orthoball Cover Lemma below generalizes this claim further. It is a crucial step to proving the Orthocenter Containment Lemma, and it is also the key to guaranteeing boundary conformity for a postprocessing step in Chapter 11 that modifies vertex weights to eliminate sliver tetrahedra from a mesh.

Lemma 7.6 (Orthoball Cover Lemma)**.** *Let $\mathcal{Q}$ be a quarantined complex in $\mathbb{R}^d$, with a weight assignment ω to the vertices of $\mathcal{Q}$. Let $x[\varpi_x]$ be a weighted point in* aff $|\mathcal{Q}|$ *such that $\pi(v[\omega_v], x[\varpi_x]) \geq 0$ for every vertex v in $\mathcal{Q}$. Let p be a point in $\mathbb{R}^d$, and let $\tilde{p}$ be the orthogonal projection of p onto* aff $|\mathcal{Q}|$*. If $x\tilde{p}$ intersects $|\mathcal{Q}|$, then there is a ball $B \in \text{Ball}(\mathcal{Q})$ such that $\pi(p, B) \leq \pi(p, x[\varpi_x])$.*

P . If p coincides with a vertex $v \in \mathcal{Q}$, the ball associated with $v[\omega_v]$ satisfies the claim because $\pi(p, v[\omega_v]) = -\omega_v$, which is nonpositive because $\mathcal{Q}$ is a quarantined complex, and $\pi(p, x[\varpi_x]) = \pi(v[\omega_v], x[\varpi_x]) + \omega_v \geq \omega_v$, which is nonnegative.

Consider a point $p \in |\mathcal{Q}|$ that is not a vertex. Let σ be the highest-dimensional simplex in $\mathcal{Q}$ that contains p. Either σ has the same dimensionality as $\mathcal{Q}$ or σ is a boundary simplex. In either case, the diametric orthoball B_σ of σ is in $\text{Ball}(\mathcal{Q})$. We will see that B_σ satisfies the claim.

Following the proof of Lemma 2.1, associate the weighted point $x[\varpi_x]$ with the hyperplane $h \subset \mathbb{R}^{d+1}$ whose formula is $q_{d+1} = 2x \cdot q - \|x\|^2 + \varpi_x$, where q varies over $\mathbb{R}^d$. The power distance $\pi(v[\omega_v], x[\varpi_x])$ is the signed vertical distance of v^+ above h (see Exercise 6), and the requirement that $\pi(v[\omega_v], x[\varpi_x]) \geq 0$ for every vertex v in $\mathcal{Q}$ is equivalent to requiring that every lifted vertex $v^+ \in \mathbb{R}^{d+1}$ lie on or above h. Therefore, the entire lifted complex $\mathcal{Q}^+ = \{\sigma^+ : \sigma \in \mathcal{Q}\}$ lies on or above h.

Likewise, $\pi(p, B_\sigma)$ is the signed vertical distance of p^+ above σ^+. As σ^+ lies on or above h, $\pi(p, B_\sigma) \leq \pi(p, x[\varpi_x])$ and B_σ satisfies the claim.

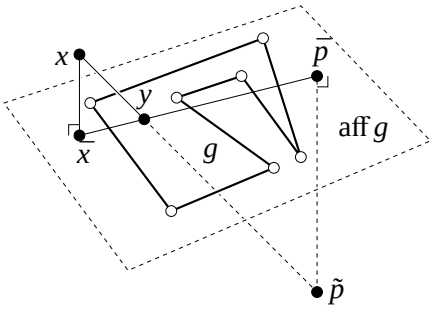

Figure 7.9: The linear cell g of least dimension through which $x\tilde{p}$ exits $|\mathcal{Q}|$.

Consider a point $p \notin |\mathcal{Q}|$. If $\tilde{p} \in |\mathcal{Q}|$, the reasoning above establishes that there is a ball $B \in \text{Ball}(\mathcal{Q})$ such that $\pi(\tilde{p}, B) \leq \pi(\tilde{p}, x[\varpi_x])$. Because $p\tilde{p}$ is orthogonal to aff $|\mathcal{Q}|$ and B is centered on aff $|\mathcal{Q}|$, $\pi(p, B) = d(p, \tilde{p})^2 + \pi(\tilde{p}, B) \leq \pi(p, x[\varpi_x])$ and B satisfies the claim.

If $p \notin |\mathcal{Q}|$ and $\tilde{p} \notin |\mathcal{Q}|$, we prove the lemma by induction on the dimension k of $\mathcal{Q}$. In the base case, $k = 0$ and $\tilde{p}$ must be in $|\mathcal{Q}|$, so the lemma holds. If $k > 0$, assume for the inductive step that the lemma holds for quarantined complexes of dimension less than k. Because $x\tilde{p}$ intersects $|\mathcal{Q}|$, a walk from x to $\tilde{p}$ leaves $|\mathcal{Q}|$ at a point y on the boundary of $|\mathcal{Q}|$. Let $\mathcal{P}$ be the PLC that $\mathcal{Q}$ triangulates, let $g \in \mathcal{P}$ be the linear cell of least dimension that contains y, as illustrated in Figure 7.9, and observe that $\tilde{p} \notin \text{aff}\, g$. Let $\mathcal{Q}|_g = \{\sigma \in \mathcal{Q} : \sigma \subseteq g\}$ be the subset of simplices that triangulate g, and observe that $\mathcal{Q}|_g$ is a quarantined complex of dimension less than k.

Let $\bar{x}$ be the projection of x onto aff g, and let $\varpi_{\bar{x}} = \varpi_x - d(x, \bar{x})^2$. By Proposition 7.3, the weighted point $\bar{x}[\varpi_{\bar{x}}]$ is at nonnegative power distances from the weighted vertices of g because $x[\varpi_x]$ is. Observe that the projection $\bar{p}$ of p onto aff g is also the projection of $\tilde{p}$ onto aff g, and that $\overline{x\bar{p}}$ intersects g because g, $x\tilde{p}$, and $\overline{x\bar{p}}$ all contain y, as illustrated. By the inductive assumption, there is a ball $B \in \text{Ball}(\mathcal{Q}|_g)$ such that $\pi(p, B) \leq \pi(p, \bar{x}[\varpi_{\bar{x}}])$.

The ball B satisfies the claim by the following reasoning. Because $x\tilde{p}$ intersects g but $\tilde{p} \notin \text{aff}\, g$, either x lies on g or $x\tilde{p}$ crosses aff g. In the former case, $\bar{x} = x$; in the latter case, $\angle x\bar{x}\tilde{p} > 90°$ and therefore $\angle x\bar{x}p > 90°$. In either case, $d(p, \bar{x})^2 + d(\bar{x}, x)^2 \leq d(p, x)^2$. It follows that $\pi(p, B) \leq d(p, \bar{x})^2 - \varpi_{\bar{x}} = d(p, \bar{x})^2 + d(\bar{x}, x)^2 - \varpi_x \leq d(p, x)^2 - \varpi_x = \pi(p, x[\varpi_x])$. □

The proof of the Orthoball Cover Lemma uses only the first two properties of a quarantined complex, and it does not use the weighted Delaunay property. Hence, it holds for a wide class of PLC triangulations besides quarantined complexes.

7.3.3 The Orthocenter Containment Lemma

The Orthocenter Containment Lemma states that all the orthocenters of the k-simplices and boundary simplices in a quarantined complex lie in its underlying space. Delaunay refinement algorithms insert orthocenters of simplices to obtain domain conformity and high element quality, so it is important to ensure that the orthocenters lie in the domain.

Lemma 7.7 (Orthocenter Containment Lemma)**.** *Let $\mathcal{Q}$ be a quarantined complex of dimension k in $\mathbb{R}^d$. The orthocenter of every k-simplex lies in $|\mathcal{Q}|$, as does the orthocenter of every boundary*

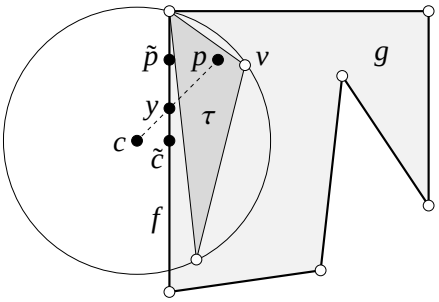

Figure 7.10: A face f of g separates the simplex $\tau \subseteq g$ from its orthocenter c.

simplex of $\mathcal{Q}$. More specifically, let $\mathcal{P}$ be the PLC that $\mathcal{Q}$ triangulates. If an i-simplex $\tau \in \mathcal{Q}$ is included in a linear i-cell $g \in \mathcal{P}$, then the orthocenter of τ lies in g.

P . Suppose for the sake of contradiction that some i-simplex $\tau \in \mathcal{Q}$ is included in a linear i-cell $g \in \mathcal{P}$ but has an orthocenter $c \notin g$. Let p be a point in the interior of τ. The line segment from p to c leaves g at a point y on g's boundary, illustrated in Figure 7.10. Let $f \in \mathcal{P}$ be the face of g of least dimension that contains y, and observe that $c \notin \text{aff } f$ and $p \notin \text{aff } f$. Let $\tilde{c}$ be the orthogonal projection of c onto aff f. Because pc crosses f, $\angle p\tilde{c}c > 90°$.

Let $\sqrt{\varpi_c}$ be the orthoradius of τ, and let $\varpi_{\tilde{c}} = \varpi_c - d(c,\tilde{c})^2$. By Proposition 7.3, the weighted point $\tilde{c}[\varpi_{\tilde{c}}]$ has nonnegative power distances from the weighted vertices of f because $c[\varpi_c]$ does. Let $\mathcal{Q}|_f = \{\sigma \in \mathcal{Q} : \sigma \subseteq f\}$ be the subset of simplices that triangulate f. Observe that $\mathcal{Q}|_f$ is a quarantined complex. Let $\tilde{p}$ be the orthogonal projection of p onto aff f, and observe that $\tilde{p}\tilde{c}$ contains y and therefore intersects f, as illustrated. By the Orthoball Cover Lemma (Lemma 7.6), there is a ball $B \in \text{Ball}(\mathcal{Q}|_f)$ such that $\pi(p,B) \leq \pi(p,\tilde{c}[\varpi_{\tilde{c}}])$. From this and the fact that $\angle p\tilde{c}c > 90°$, we have $\pi(p,B) \leq d(p,\tilde{c})^2 - \varpi_{\tilde{c}} = d(p,\tilde{c})^2 + d(\tilde{c},c)^2 - \varpi_c < d(p,c)^2 - \varpi_c = \pi(p,c[\varpi_c])$. Because $p \in \tau$ and $\pi(p,B) - \pi(p,c[\varpi_c])$ is a linear function of p, there is a vertex v of τ for which $\pi(v,B) < \pi(v,c[\varpi_c])$. Because $c[\varpi_c]$ is orthogonal to $v[\omega_v]$, we conclude that $\pi(v[\omega_v],B) < \pi(v[\omega_v],c[\varpi_c]) = 0$. But this implies that v is closer than orthogonal to a ball in Ball($\mathcal{Q}$), which is not permitted in a quarantined complex. □

The proof can be interpreted as saying that if the orthocenter of τ lies outside g, with a face f of g separating τ from its orthocenter, then some vertex of τ encroaches upon a boundary simplex included in f or is too close to a vertex of f for $\mathcal{Q}$ to be a quarantined complex. The encroaching vertex is the vertex of τ that lies furthest toward B's side of the bisector between B and the orthocenter $c[\varpi_c]$.

7.4 Notes and exercises

Gustav Dirichlet [81] first used two-dimensional and three-dimensional Voronoi diagrams in the study of quadratic forms in 1850. The general Voronoi diagram in $\mathbb{R}^d$ is named after Georgy Feodosevich Voronoi [220] (sometimes written Voronoy) who defined it in 1908. Voronoi was one of Boris Delaunay's two doctoral advisors.

The informal use of the Voronoi diagram dates back to Descartes's *Principia Philosophiae* [70]

of 1644. In an illustration, he decomposes space into convex, Voronoi-like regions that contain one star each. Each star seems to exert more influence in its region than the other stars. The concept of Voronoi diagrams has been conceived in several fields and given different names, such as *domains of action* in crystallography [160], *Wigner–Seitz zones* in metallurgy [224], and *Thiessen polygons* in geography and meteorology [114, 214].

The power distance as a generalized distance function was known to Voronoi. The weighted Voronoi diagram, also known as the *power diagram*, has been used in problems in packings and coverings of spheres [177] and number theory [135]. In computational geometry, Aurenhammer [9] pioneered the study of the weighted Voronoi diagram. For in-depth surveys of Voronoi diagrams, see Aurenhammer [10] and Aurenhammer and Klein [11].

The Monotone Power Lemma generalizes a result of Edelsbrunner [85], who uses it to show that if the vertices of a weighted Delaunay triangulation are generic, then from the perspective of any fixed point in $\mathbb{R}^d$, it is impossible to have a cycle of d-simplices in which each simplex overlaps the next simplex in the cycle. The insight is that any ray shot from the viewpoint intersects d-simplices in increasing order according to the power distances of their orthoballs from the viewpoint. The result does not hold if the vertices are not generic, because it is possible to have a cycle of overlapping simplices that have the same orthoball.

The Orthocenter Containment Lemma began with Ruppert's proof of Proposition 6.2 for unweighted points in the plane [180], was partly extended to three dimensions by Shewchuk [198], and was extended to weighted Delaunay triangulations by Cheng and Dey [48].

Exercises

1. Draw the Voronoi diagram and Delaunay triangulation of the following point set.

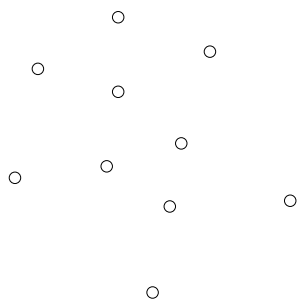

2. Let Y and Z be two point sets in the plane, together having n points. Consider overlaying their Voronoi diagrams $\mathrm{Vor}\,Y$ and $\mathrm{Vor}\,Z$: that is, draw the edges of both Voronoi diagrams and add a new vertex wherever an edge of $\mathrm{Vor}\,Y$ intersects an edge of $\mathrm{Vor}\,Z$, thereby creating a polyhedral complex called the *overlay* of $\mathrm{Vor}\,Y$ and $\mathrm{Vor}\,Z$. Formally, for each face $f \in \mathrm{Vor}\,Y$ and each face $g \in \mathrm{Vor}\,Z$, the overlay contains the face $f \cap g$ if it is nonempty.

   (a) Give an example showing how to choose sets Y and Z so the overlay has $\Theta(n^2)$ faces. Make sure it is clear how to generalize your example as $n \to \infty$.

   (b) Prove that for any two point sets Y and Z, the overlay of their Voronoi diagrams has only $O(n)$ *unbounded* faces.

3. The worst-case complexity of a Voronoi diagram of n sites is $\Theta(n^2)$ in three dimensions and $\Theta(n^{\lceil d/2 \rceil})$ in $\mathbb{R}^d$.

   (a) Prove that a planar cross-section of a d-dimensional Voronoi diagram—in other words, the polygonal complex formed by intersecting a 2-flat with a Voronoi diagram—has complexity no greater than $O(n)$.

   (b) Does a planar cross-section of a d-dimensional Delaunay triangulation also always have a complexity of $O(n)$? Explain.

4. Show that the weighted Voronoi bisector of two weighted points is a hyperplane orthogonal to the line through the two points. Then show that if the points have nonnegative weights and are further than orthogonal to each other, the bisector lies between them.

5. Show that at most two triangular faces of a tetrahedron can lie between the tetrahedron's interior and its circumcenter, but if its vertices have nonzero weights, three triangular faces can lie between its interior and its orthocenter.

6. Let $p[\omega_p]$ and $q[\omega_q]$ be two weighted points in $\mathbb{R}^d$, each representing either a weighted vertex or an orthoball. Lemma 2.1 associates each ball B_p in $\mathbb{R}^d$ with a hyperplane h_p in $\mathbb{R}^{d+1}$, and thereby associates each weighted point $p[\omega_p]$ with a hyperplane h_p. Show that the power distance $\pi(p[\omega_p], q[\omega_q])$ is equal to the vertical distance of p^+ above h_q, or symmetrically, the vertical distance of q^+ above h_p.

7. Let $\mathcal{T}$ be a triangulation of a point set in the plane, and let x be an arbitrary point in the plane. Say that a triangle $t_2 \in \mathcal{T}$ is behind a triangle $t_1 \in \mathcal{T}$ from x's perspective if one of the following conditions is satisfied.

   - Some ray shot from x intersects first t_1 and then t_2.
   - There is a triangle $t \in \mathcal{T}$ such that t is behind t_1 and t_2 is behind t from x's perspective. (Note that there need not be a single ray originating at x that intersects both t_1 and t_2, and that this definition is recursive.)

   Draw a triangulation $\mathcal{T}$ in the plane and a viewpoint x such that some triangle in $\mathcal{T}$ is behind itself. In other words, there is a cyclical list of triangles such that each triangle is behind the preceding triangle in the cycle.

8. Consider a triangulation comprising just one triangle abc and its faces. Even for this simple triangulation, the consequence of the Monotone Power Lemma can fail to hold if $c[\omega_c]$ is closer than orthogonal to the diametric orthoball of ab. Design a counterexample that demonstrates this failure.

9. Consider a triangulation comprising just one edge ab and its vertices. The consequence of the Orthoball Cover Lemma can fail to hold if the condition that $\pi(a[\omega_a], x[\varpi_x]) \geq 0$ is not satisfied. Design a counterexample that demonstrates this failure.

10. Let $\mathcal{Q}$ be a quarantined complex with a weight assignment ω. Let $x[\varpi_x]$ be a weighted point such that $x \in |\mathcal{Q}|$ and $\pi(x[\varpi_x], v[\omega_v]) \geq 0$ for every vertex v in $\mathcal{Q}$. Show that $B(x, \sqrt{\varpi_x})$ is included in the union of the balls in Ball($\mathcal{Q}$).

11. Consider a triangulation comprising just one triangle τ and its faces. Suppose that every vertex weight is zero and τ has an obtuse angle. Show that the consequence of the Orthocenter Containment Lemma fails to hold, and specify which property of a quarantined complex the triangle violates.