Chapter 8

Dirichlet–Voronoi Diagrams and Delaunay Triangulations

8.1 Dirichlet–Voronoi Diagrams

In this chapter we present very briefly the concepts of a Voronoi diagram and of a Delaunay triangulation.

These are important tools in computational geometry, and Delaunay triangulations are important in problems where it is necessary to fit 3D data using surface splines.

It is usually useful to compute a good mesh for the projection of this set of data points onto the xy-plane, and a Delaunay triangulation is a good candidate.

Our presentation will be rather sketchy. We are primarily interested in defining these concepts and stating their most important properties without proofs.

For a comprehensive exposition of Voronoi diagrams, Delaunay triangulations, and more topics in computational geometry, consult O'Rourke [**?**], Preparata and Shamos [**?**], Boissonnat and Yvinec [**?**], de Berg, Van Kreveld, Overmars, and Schwarzkopf [**?**], or Risler [**?**].

The survey by Graham and Yao [**?**] contains a very gentle and lucid introduction to computational geometry.

For concreteness, one may safely assume that we work in the affine space $\mathcal{E} = \mathbb{E}^m$, although what follows applies to any Euclidean space of finite dimension.

Given a set $P = \{p_1, \ldots, p_n\}$ of n points in $\mathcal{E}$, it is often useful to find a partition of the space $\mathcal{E}$ into regions each containing a single point of P and having some nice properties.

It is also often useful to find triangulations of the convex hull of P having some nice properties.

We shall see that this can be done and that the two problems are closely related. In order to solve the first problem, we need to introduce bisector lines and bisector planes.

For simplicity, let us first assume that $\mathcal{E}$ is a plane i.e., has dimension 2.

Given any two distinct points $a, b \in \mathcal{E}$, the line orthogonal to the line segment (a, b) and passing through the midpoint of this segment is the locus of all points having equal distance to a and b.

It is called the *bisector line of a and b*. The bisector line of two points is illustrated in Figure 8.1.

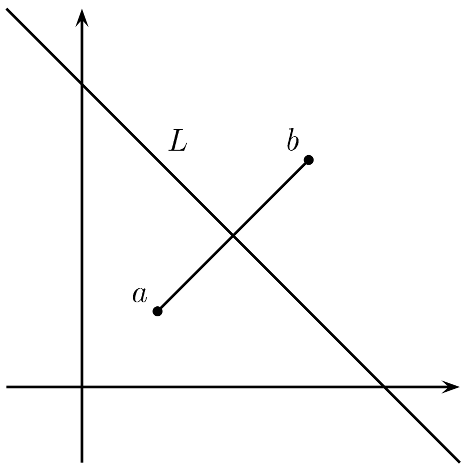

Figure 8.1: The bisector line L of a and b

If $h = \frac{1}{2}\,a + \frac{1}{2}\,b$ is the midpoint of the line segment (a, b), letting m be an arbitrary point on the bisector line, the equation of this line can be found by writing that **hm** is orthogonal to **ab**.

In any orthogonal frame, letting $m = (x, y)$, $a = (a_1, a_2)$, $b = (b_1, b_2)$, the equation of this line can be written as

$$(b_1 - a_1)x + (b_2 - a_2)y = (b_1^2 + b_2^2)/2 - (a_1^2 + a_2^2)/2.$$

The closed half-plane $H(a,b)$ containing a and with boundary the bisector line is the locus of all points such that

$$(b_1 - a_1)x + (b_2 - a_2)y \leq (b_1^2 + b_2^2)/2 - (a_1^2 + a_2^2)/2,$$

and the closed half-plane $H(b,a)$ containing b and with boundary the bisector line is the locus of all points such that

$$(b_1 - a_1)x + (b_2 - a_2)y \geq (b_1^2 + b_2^2)/2 - (a_1^2 + a_2^2)/2.$$

The closed half-plane $H(a,b)$ is the set of all points whose distance to a is less that or equal to the distance to b, and vice versa for $H(b,a)$. Thus, points in the closed half-plane $H(a,b)$ are closer to a than they are to b.

We now consider a problem called the *post office problem* by Graham and Yao [**?**].

Given any set $P = \{p_1, \ldots, p_n\}$ of n points in the plane (considered as *post offices* or *sites*), for any arbitrary point x, find out which post office is closest to x.

Since x can be arbitrary, it seems desirable to precompute the sets $V(p_i)$ consisting of all points that are closer to p_i than to any other point $p_j \neq p_i$.

Indeed, if the sets $V(p_i)$ are known, the answer is any post office p_i such that $x \in V(p_i)$.

Thus, it remains to compute the sets $V(p_i)$. For this, if x is closer to p_i than to any other point $p_j \neq p_i$, then x is on the same side as p_i with respect to the bisector line of p_i and p_j for every $j \neq i$, and thus

$$V(p_i) = \bigcap_{j \neq i} H(p_i, p_j).$$

If $\mathcal{E}$ has dimension 3, the locus of all points having equal distance to a and b is a plane. It is called the *bisector plane of* a *and* b.

The equation of this plane is also found by writing that $\mathbf{hm}$ is orthogonal to $\mathbf{ab}$. The equation of this plane can be written as

$$(b_1 - a_1)x + (b_2 - a_2)y + (b_3 - a_3)z = \\ (b_1^2 + b_2^2 + b_3^2)/2 - (a_1^2 + a_2^2 + a_3^2)/2.$$

The closed half-space $H(a, b)$ containing a and with boundary the bisector plane is the locus of all points such that

$$(b_1 - a_1)x + (b_2 - a_2)y + (b_3 - a_3)z \leq \\ (b_1^2 + b_2^2 + b_3^2)/2 - (a_1^2 + a_2^2 + a_3^2)/2,$$

and the closed half-space $H(b, a)$ containing b and with boundary the bisector plane is the locus of all points such that

$$(b_1 - a_1)x + (b_2 - a_2)y + (b_3 - a_3)z \geq \\ (b_1^2 + b_2^2 + b_3^2)/2 - (a_1^2 + a_2^2 + a_3^2)/2.$$

The closed half-space $H(a, b)$ is the set of all points whose distance to a is less that or equal to the distance to b, and vice versa for $H(b, a)$. Again, points in the closed half-space $H(a, b)$ are closer to a than they are to b.

Given any set $P = \{p_1, \ldots, p_n\}$ of n points in $\mathcal{E}$ (of dimension $m = 2, 3$), it is often useful to find for every point p_i the region consisting of all points that are closer to p_i than to any other point $p_j \neq p_i$, that is, the set

$$V(p_i) = \{x \in \mathcal{E} \mid d(x, p_i) \leq d(x, p_j), \text{ for all } j \neq i\},$$

where $d(x, y) = (\mathbf{xy} \cdot \mathbf{xy})^{1/2}$, the Euclidean distance associated with the inner product $\cdot$ on $\mathcal{E}$.

From the definition of the bisector line (or plane), it is immediate that

$$V(p_i) = \bigcap_{j \neq i} H(p_i, p_j).$$

Families of sets of the form $V(p_i)$ were investigated by Dirichlet [**?**] (1850) and Voronoi [**?**] (1908). Voronoi diagrams also arise in crystallography (Gilbert [**?**]).

Other applications, including facility location and path planning, are discussed in O'Rourke [**?**]. For simplicity, we also denote the set $V(p_i)$ by V_i, and we introduce the following definition.

Definition 8.1.1 Let $\mathcal{E}$ be a Euclidean space of dimension $m = 2, 3$. Given any set $P = \{p_1, \ldots, p_n\}$ of n points in $\mathcal{E}$, the *Dirichlet–Voronoi diagram* $\mathcal{V}(P)$ of $P = \{p_1, \ldots, p_n\}$ is the family of subsets of $\mathcal{E}$ consisting of the sets $V_i = \bigcap_{j \neq i} H(p_i, p_j)$ and of all of their intersections.

Dirichlet–Voronoi diagrams are also called *Voronoi diagrams*, *Voronoi tessellations*, or *Thiessen polygons*. Following common usage, we will use the terminology *Voronoi diagram*.

As intersections of convex sets (closed half-planes or closed half-spaces), the *Voronoi regions* $V(p_i)$ are convex sets. In dimension two, the boundaries of these regions are convex polygons, and in dimension three, the boundaries are convex polyhedra.

Whether a region $V(p_i)$ is bounded or not depends on the location of p_i.

If p_i belongs to the boundary of the convex hull of the set P, then $V(p_i)$ is unbounded, and otherwise bounded.

In dimension two, the convex hull is a convex polygon, and in dimension three, the convex hull is a convex polyhedron.

As we will see later, there is an intimate relationship between convex hulls and Voronoi diagrams.

Generally, if $\mathcal{E}$ is a Euclidean space of dimension m, given any two distinct points $a, b \in \mathcal{E}$, the locus of all points having equal distance to a and b is a hyperplane.

It is called the *bisector hyperplane of a and b*. The equation of this hyperplane is still found by writing that $\mathbf{hm}$ is orthogonal to $\mathbf{ab}$. The equation of this hyperplane can be written as

$$(b_1 - a_1)x_1 + \cdots + (b_m - a_m)x_m = (b_1^2 + \cdots + b_m^2)/2 - (a_1^2 + \cdots + a_m^2)/2.$$

The closed half-space $H(a, b)$ containing a and with boundary the bisector hyperplane is the locus of all points such that

$$(b_1 - a_1)x_1 + \cdots + (b_m - a_m)x_m \leq (b_1^2 + \cdots + b_m^2)/2 - (a_1^2 + \cdots + a_m^2)/2,$$

and the closed half-space $H(b, a)$ containing b and with boundary the bisector hyperplane is the locus of all points such that

$$(b_1 - a_1)x_1 + \cdots + (b_m - a_m)x_m \geq (b_1^2 + \cdots + b_m^2)/2 - (a_1^2 + \cdots + a_m^2)/2.$$

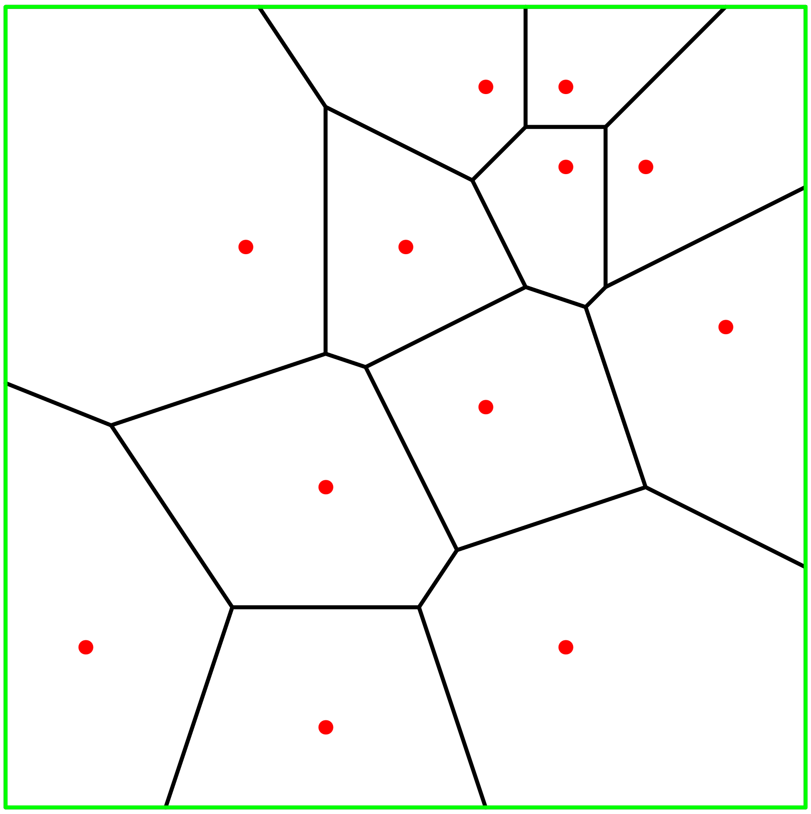

Figure 8.2: A Voronoi diagram

The closed half-space $H(a, b)$ is the set of all points whose distance to a is less than or equal to the distance to b, and vice versa for $H(b, a)$.

Figure 8.2 shows the Voronoi diagram of a set of twelve points.

In the general case where $\mathcal{E}$ has dimension m, the definition of the Voronoi diagram $\mathcal{V}(P)$ of P is the same as Definition 8.1.1, except that $H(p_i, p_j)$ is the closed half-space containing p_i and having the bisector hyperplane of a and b as boundary.

Also, observe that the convex hull of P is a convex polytope.

We will now state a lemma listing the main properties of Voronoi diagrams.

It turns out that certain degenerate situations can be avoided if we assume that if P is a set of points in an affine space of dimension m, then no $m+2$ points from P belong to the same $(m-1)$-sphere.

We will say that the points of P are in *general position*.

Thus when $m = 2$, no 4 points in P are cocyclic, and when $m = 3$, no 5 points in P are on the same sphere.

Lemma 8.1.2 *Given a set $P = \{p_1, \ldots, p_n\}$ of n points in some Euclidean space $\mathcal{E}$ of dimension m (say $\mathbb{E}^m$), if the points in P are in general position, then the Voronoi diagram of P satisfies the following conditions:*

(1) Each region V_i is convex and contains p_i in its interior.

(2) Each vertex of V_i belongs to $m+1$ regions V_j and to $m+1$ edges.

(3) The region V_i is unbounded iff p_i belongs to the boundary of the convex hull of P.

(4) If p is a vertex that belongs to the regions $V_1, \ldots, V_{m+1}$, then p is the center of the $(m-1)$-sphere $S(p)$ determined by $p_1, \ldots, p_{m+1}$. Furthermore, no point in P is inside the sphere $S(p)$ (i.e., in the open ball associated with the sphere $S(p)$).

(5) If p_j is a nearest neighbor of p_i, then one of the faces of V_i is contained in the bisector hyperplane of (p_i, p_j).

(6)

$$\bigcup_{i=1}^{n} V_i = \mathcal{E}, \quad \text{and} \quad \overset{\circ}{V}_i \cap \overset{\circ}{V}_j = \emptyset, \quad \text{for all } i, j,\ i \neq j,$$

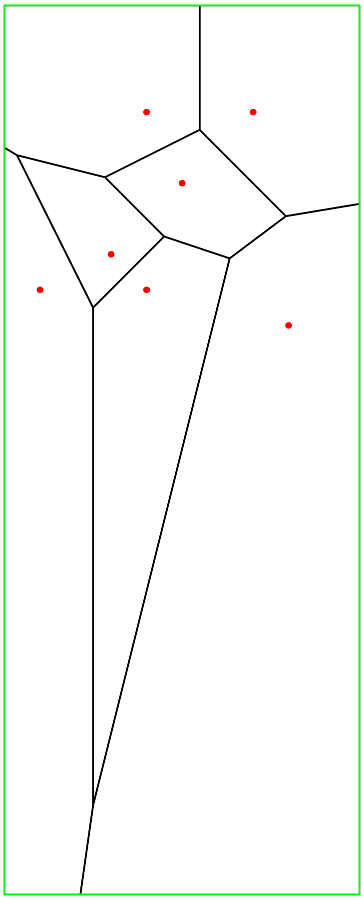

Figure 8.3: Another Voronoi diagram

where $\overset{\circ}{V}_i$ *denotes the interior of* V_i.

For simplicity, let us again consider the case where $\mathcal{E}$ is a plane. It should be noted that certain Voronoi regions, although closed, may extend very far.

Figure 8.3 shows such an example.

It is also possible for certain unbounded regions to have parallel edges.

There are a number of methods for computing Voronoi diagrams. A fairly simple (although not very efficient) method is to compute each Voronoi region $V(p_i)$ by intersecting the half-planes $H(p_i, p_j)$.

One way to do this is to construct successive convex polygons that converge to the boundary of the region.

At every step we intersect the current convex polygon with the bisector line of p_i and p_j. There are at most two intersection points. We also need a starting polygon, and for this we can pick a square containing all the points.

A naive implementation will run in $O(n^3)$.

However, the intersection of half-planes can be done in $O(n \log n)$, using the fact that the vertices of a convex polygon can be sorted.

Thus, the above method runs in $O(n^2 \log n)$. Actually, there are faster methods (see Preparata and Shamos [**?**] or O'Rourke [**?**]), and it is possible to design algorithms running in $O(n \log n)$.

The most direct method to obtain fast algorithms is to use the "lifting method" discussed in Section 8.4, whereby the original set of points is lifted onto a paraboloid, and to use fast algorithms for finding a convex hull.

A very interesting (undirected) graph can be obtained from the Voronoi diagram as follows: The vertices of this graph are the points p_i (each corresponding to a unique region of $\mathcal{V}(P)$), and there is an edge between p_i and p_j iff the regions V_i and V_j share an edge.

The resulting graph is called a *Delaunay triangulation* of the convex hull of P, after Delaunay, who invented this concept in 1934. Such triangulations have remarkable properties.

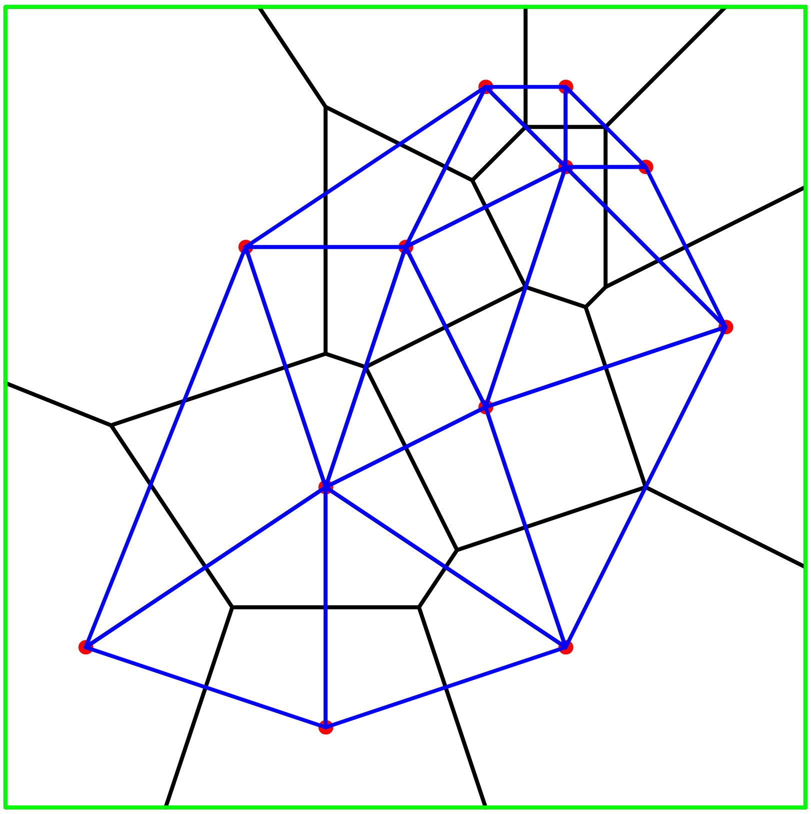

Figure 8.4: Delaunay triangulation associated with a Voronoi diagram

Figure 8.4 shows the Delaunay triangulation associated with the earlier Voronoi diagram of a set of twelve points.

One has to be careful to make sure that all the Voronoi vertices have been computed before computing a Delaunay triangulation, since otherwise, some edges could be missed.

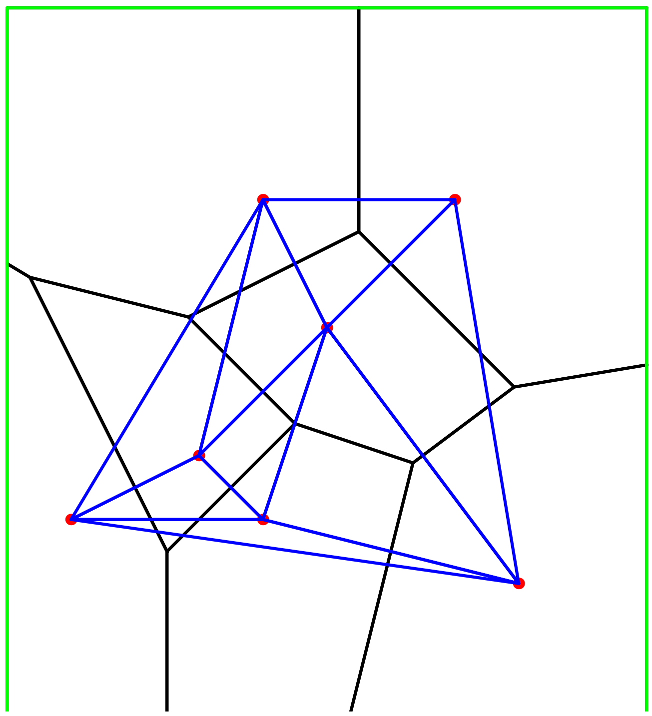

Figure 8.5: Another Delaunay triangulation associated with a Voronoi diagram

In Figure 8.5 illustrating such a situation, if the lowest Voronoi vertex had not been computed (not shown on the diagram!), the lowest edge of the Delaunay triangulation would be missing.

The concept of a triangulation can be generalized to dimension 3, or even to any dimension m.

But first, we need to define a triangulation precisely, and for this, we need to review what is a simplicial complex.

8.2 Simplicial Complexes and Triangulations

A simplex is just the convex hull of a finite number of affinely independent points, but we also need to define faces, the boundary, and the interior of a simplex.

Definition 8.2.1 Let $\mathcal{E}$ be any normed affine space, say $\mathcal{E} = \mathbb{E}^m$ with its usual Euclidean norm. Given any $n+1$ affinely independent points $a_0, \ldots, a_n$ in $\mathcal{E}$, the *n-simplex (or simplex) σ defined by $a_0, \ldots, a_n$* is the convex hull of the points $a_0, \ldots, a_n$, that is, the set of all convex combinations $\lambda_0 a_0 + \cdots + \lambda_n a_n$, where $\lambda_0 + \cdots + \lambda_n = 1$ and $\lambda_i \geq 0$ for all i, $0 \leq i \leq n$. We call n the *dimension* of the n-simplex σ, and the points $a_0, \ldots, a_n$ are the *vertices* of σ.

Given any subset $\{a_{i_0}, \ldots, a_{i_k}\}$ of $\{a_0, \ldots, a_n\}$ (where $0 \leq k \leq n$), the k-simplex generated by $a_{i_0}, \ldots, a_{i_k}$ is called a *face* of σ. A face s of σ is a *proper face* if $s \neq \sigma$ (we agree that the empty set is a face of any simplex).

For any vertex a_i, the face generated by $a_0, \ldots, a_{i-1}, a_{i+1}$, $\ldots, a_n$ (i.e., omitting a_i) is called the *face opposite* a_i. Every face that is an $(n-1)$-simplex is called a *boundary face*.

The union of the boundary faces is the *boundary of* σ, denoted by $\partial\sigma$, and the complement of $\partial\sigma$ in σ is the *interior* $\mathrm{Int}\,\sigma = \sigma - \partial\sigma$ of σ. The interior $\mathrm{Int}\,\sigma$ of σ is sometimes called an *open simplex*.

It should be noted that for a 0-simplex consisting of a single point $\{a_0\}$, $\partial\{a_0\} = \emptyset$, and $\mathrm{Int}\,\{a_0\} = \{a_0\}$.

Of course, a 0-simplex is a single point, a 1-simplex is the line segment (a_0, a_1), a 2-simplex is a triangle (a_0, a_1, a_2) (with its interior), and a 3-simplex is a tetrahedron (a_0, a_1, a_2, a_3) (with its interior); see Figure 8.6.

We now state a number of properties of simplices, whose proofs are left as an exercise.

Clearly, a point x belongs to the boundary $\partial\sigma$ of σ iff at least one of its barycentric coordinates $(\lambda_0, \ldots, \lambda_n)$ is zero, and a point x belongs to the interior $\mathrm{Int}\,\sigma$ of σ iff all of its barycentric coordinates $(\lambda_0, \ldots, \lambda_n)$ are positive,

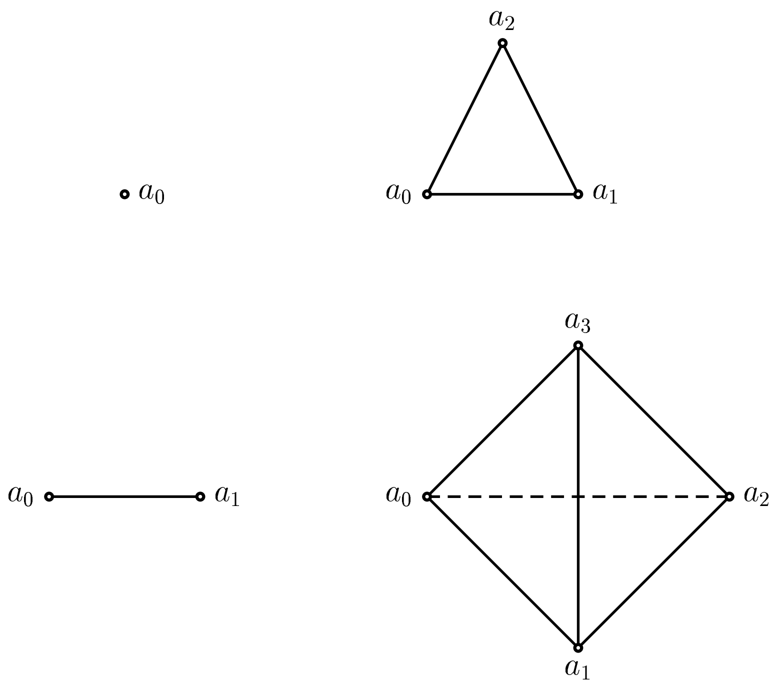

Figure 8.6: Examples of simplices

i.e., $\lambda_i > 0$ for all i, $0 \leq i \leq n$.

Then, for every $x \in \sigma$, there is a unique face s such that $x \in \mathrm{Int}\, s$, the face generated by those points a_i for which $\lambda_i > 0$, where $(\lambda_0, \ldots, \lambda_n)$ are the barycentric coordinates of x.

A simplex σ is convex, arcwise connected, compact, and closed. The interior $\mathrm{Int}\, \sigma$ of a simplex is convex, arcwise connected, open, and σ is the closure of $\mathrm{Int}\, \sigma$.

We now need to put simplices together to form more complex shapes, following Munkres [**?**].

Definition 8.2.2 A *simplicial complex in* $\mathbb{E}^m$ (for short, a *complex* in $\mathbb{E}^m$) is a set K consisting of a (finite or infinite) set of simplices in $\mathbb{E}^m$ satisfying the following conditions:

(1) Every face of a simplex in K also belongs to K.

(2) For any two simplices σ_1 and σ_2 in K, if $\sigma_1 \cap \sigma_2 \neq \emptyset$, then $\sigma_1 \cap \sigma_2$ is a common face of both σ_1 and σ_2.

If $\sigma \in K$ is a simplex of $n+1$ elements, then its dimension is n, and it is called an *n-simplex*.

A 0-simplex $\{x\}$ is called a *vertex*.

The *dimension* of the simplicial complex K is the maximum of the dimensions of simplices in K.

Condition (2) guarantees that the various simplices forming a complex are glued nicely. It can be shown that the following condition is equivalent to condition (2):

(2′) For any two distinct simplices σ_1, σ_2, $\operatorname{Int}\sigma_1 \cap \operatorname{Int}\sigma_2 = \emptyset$.

The union K_g of all the simplices in K is a subset of $\mathbb{E}^m$.

We can define a topology on K_g by defining a subset F of K_g to be closed iff $F \cap \sigma$ is closed in σ for every simplex $\sigma \in K$.

It is immediately verified that the axioms of a topological space are indeed satisfied.

The resulting topological space K_g is called the *geometric realization of K*.

A *polytope* is the geometric realization of some simplicial complex.

A polytope of dimension 1 is usually called a *polygon*, and a polytope of dimension 2 is usually called a *polyhedron*.

It can be checked that each region V_i of a Voronoi diagram is a (convex) polytope.

In the sequel, we will consider only finite simplicial complexes, that is, complexes K consisting of a finite number of simplices.

In this case, the topology of K_g defined above is identical to the topology induced from $\mathbb{E}^m$.

In this case, for any simplex σ in K, $\mathrm{Int}\,\sigma$ coincides with the interior $\overset{\circ}{\sigma}$ of σ in the topological sense, and $\partial\sigma$ coincides with the boundary of σ in the topological sense. We can now define triangulations.

First, assume that $\mathcal{E} = \mathbb{E}^2$. Given a subset S of $\mathbb{E}^2$, a *triangulation of* S is a finite complex K of dimension 2 such that S is the union of the 2-simplices in K.

Equivalently, S is the union of the (closed) triangles in K. Thus, a triangulation of S specifies a way of cutting up S into a collection of (closed) triangles that intersect nicely.

Next, if $\mathcal{E} = \mathbb{E}^3$, given a subset S of $\mathbb{E}^3$, a triangulation of S is a finite complex K of dimension 3 such that S is the union of the 3-simplices in K.

Equivalently, S is the union of the (closed) tetrahedra in K. Thus, a triangulation of S specifies a way of cutting up S into a collection of (closed) tetrahedra that intersect nicely. In general, we have the following definition.

Definition 8.2.3 Given a subset S of $\mathbb{E}^m$ (where $m \geq 2$), a *d-triangulation of S* (where $d \leq m$) is a finite complex K such that

$$S = \bigcup_{\substack{\sigma \in K \\ \dim(\sigma) = d}} \sigma,$$

i.e., such that S is the union of all d-simplices in K

Given a finite set P of n points in the plane, and given a triangulation of the convex hull of P having P as its set of vertices, observe that the boundary of P is a convex polygon.

Similarly, given a finite set P of points in 3-space, and given a triangulation of the convex hull of P having P as its set of vertices, observe that the boundary of P is a convex polyhedron.

It is interesting to know how many triangulations exist for a set of n points (in the plane or in 3-space), and it is also interesting to know the number of edges and faces in terms of the number of vertices in P.

These questions can be settled using the Euler–Poincaré characteristic.

We say that a polygon in the plane is a *simple polygon* iff it is a connected closed polygon such that no two edges intersect (except at a common vertex).

Lemma 8.2.4

(1) For any triangulation of a region of the plane whose boundary is a simple polygon, letting v be the number of vertices, e the number of edges, and f the number of triangles, we have the "Euler formula"

$$v - e + f = 1.$$

(2) For any polytope S homeomorphic to a closed ball in $\mathbb{E}^3$ and any triangulation of S, letting v be the number of vertices, e the number of edges, f the number of triangles, and t the number of tetrahedra, we have the "Euler formula"

$$v - e + f - t = 1.$$

(3) Furthermore, for any triangulation of the polyhedron $B(S)$ that is the boundary of S, letting v' be the number of vertices, e' the number of edges, and f' the number of triangles, we have the "Euler formula"

$$v' - e' + f' = 2.$$

Proof. We only sketch the proof. More details can be found in O'Rourke [**?**], Risler [**?**], or books on algebraic topology, such as Massey [**?**] or Munkres [**?**].

The proof of (1) is by induction on the number f of triangles.

The proof of (2) is by induction on the number t of tetrahedra.

The proof of (3) consists in first flattening the polyhedron into a planar graph in the plane. This can be done by removing some face and then by deformation. The boundary of this planar graph is a simple polygon, and the region outside this boundary corresponds to the removed face. Then by (1) we get the formula, remembering that there is one more face (this is why we get 2 instead of 1). □

It is now easy to see that in case (1), the number of edges and faces is a linear function of the number of vertices and boundary edges, and that in case (3), the number of edges and faces is a linear function of the number of vertices.

If there are e_b edges in the boundary and e_i edges not in the boundary, we have

$$3f = e_b + 2e_i,$$

and togeher with

$$v - e_b - e_i + f = 1,$$

we get

$$\begin{aligned} v - e_b - e_i + e_b/3 + 2e_i/3 &= 1, \\ 2e_b/3 + e_i/3 &= v - 1, \end{aligned}$$

and thus, $e_i = 3v - 3 - 2e_b$. Since $f = e_b/3 + 2e_i/3$, we have $f = 2v - 2 - e_b$.

Similarly, since $v' - e' + f' = 2$ and $3f' = 2e'$, we easily get $e = 3v - 6$ and $f = 2v - 4$.

Thus, given a set P of n points, the number of triangles (and edges) for any triangulation of the convex hull of P using the n points in P for its vertices is fixed.

Case (2) is trickier, but it can be shown that

$$v - 3 \leq t \leq (v-1)(v-2)/2.$$

Thus, there can be different numbers of tetrahedra for different triangulations of the convex hull of P.

Remark: The numbers of the form $v - e + f$ and $v - e + f - t$ are called *Euler–Poincaré characteristics*.

They are topological invariants, in the sense that they are the same for all triangulations of a given polytope. This is a fundamental fact of algebraic topology.

We shall now investigate triangulations induced by Voronoi diagrams.

8.3 Delaunay Triangulations

Given a set $P = \{p_1, \ldots, p_n\}$ of n points in the plane and the Voronoi diagram $\mathcal{V}(P)$ for P, we explained in Section 8.1 how to define an (undirected) graph:

The vertices of this graph are the points p_i (each corresponding to a unique region of $\mathcal{V}(P)$), and there is an edge between p_i and p_j iff the regions V_i and V_j share an edge.

The resulting graph turns out to be a triangulation of the convex hull of P having P as its set of vertices. Such a complex can be defined in general.

For any set $P = \{p_1, \ldots, p_n\}$ of n points in $\mathbb{E}^m$, we say that a triangulation of the convex hull of P is *associated with P* if its set of vertices is the set P.

Definition 8.3.1 Let $P = \{p_1, \ldots, p_n\}$ be a set of n points in $\mathbb{E}^m$, and let $\mathcal{V}(P)$ be the Voronoi diagram of P. We define a complex $\mathcal{D}(P)$ as follows:

The complex $\mathcal{D}(P)$ contains the k-simplex $\{p_1, \ldots, p_{k+1}\}$ iff $V_1 \cap \cdots \cap V_{k+1} \neq \emptyset$, where $0 \leq k \leq m$.

The complex $\mathcal{D}(P)$ is called the *Delaunay triangulation of the convex hull of* P.

Thus, $\{p_i, p_j\}$ is an edge iff $V_i \cap V_j \neq \emptyset$, $\{p_i, p_j, p_h\}$ is a triangle iff $V_i \cap V_j \cap V_h \neq \emptyset$, $\{p_i, p_j, p_h, p_k\}$ is a tetrahedron iff $V_i \cap V_j \cap V_h \cap V_k \neq \emptyset$, etc.

For simplicity, we often write $\mathcal{D}$ instead of $\mathcal{D}(P)$. A Delaunay triangulation for a set of twelve points is shown in Figure 8.7.

Actually, it is not obvious that $\mathcal{D}(P)$ is a triangulation of the convex hull of P, but this can be shown, as well as the properties listed in the following lemma.

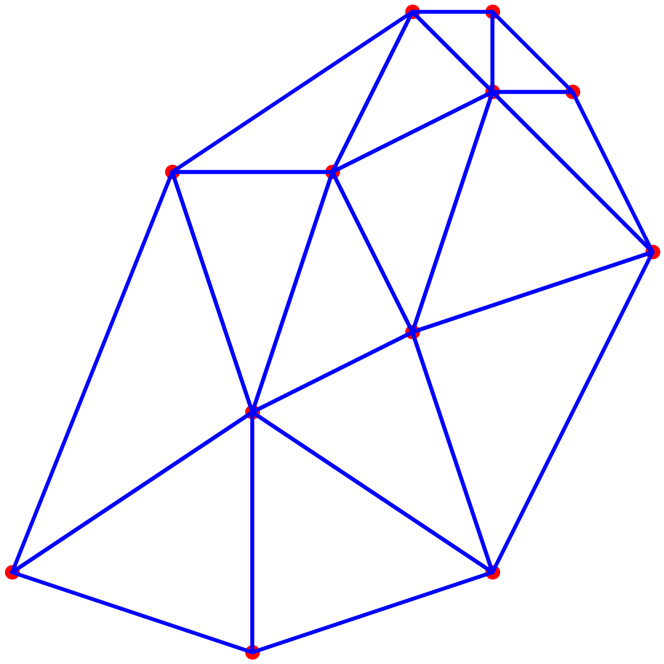

Figure 8.7: A Delaunay triangulation

Lemma 8.3.2 *Let $P = \{p_1, \ldots, p_n\}$ be a set of n points in $\mathbb{E}^m$, and assume that they are in general position. Then the Delaunay triangulation of the convex hull of P is indeed a triangulation associated with P, and it satisfies the following properties:*

(1) The boundary of $\mathcal{D}(P)$ is the convex hull of P.

(2) A triangulation T associated with P is the Delaunay triangulation $\mathcal{D}(P)$ iff every $(m-1)$-sphere $S(\sigma)$ circumscribed about an m-simplex σ of T contains no other point from P (i.e., the open ball associated with $S(\sigma)$ contains no point from P).

The proof can be found in Risler [**?**] and O'Rourke [**?**].

In the case of a planar set P, it can also be shown that the Delaunay triangulation has the property that it maximizes the minimum angle of the triangles involved in any triangulation of P. However, this does not characterize the Delaunay triangulation.

Given a connected graph in the plane, it can also be shown that any minimal spanning tree is contained in the Delaunay triangulation of the convex hull of the set of vertices of the graph (O'Rourke [**?**]).

We will now explore briefly the connection between Delaunay triangulations and convex hulls.

8.4 Delaunay Triangulations and Convex Hulls

We will see that given a set P of points in the Euclidean space $\mathbb{E}^m$ of dimension m, we can “lift” these points onto a paraboloid living in the space $\mathbb{E}^{m+1}$ of dimension $m+1$, and that the Delaunay triangulation of P is the projection of the downward-facing faces of the convex hull of the set of lifted points.

This remarkable connection was first discovered by Brown [**?**], and refined by Edelsbrunner and Seidel [**?**].

For simplicity, we consider the case of a set P of points in the plane $\mathbb{E}^2$, and we assume that they are in general position.

Consider the paraboloid of revolution of equation
$z = x^2 + y^2$.

A point $p = (x, y)$ in the plane is lifted to the point $l(p) = (X, Y, Z)$ in $\mathbb{E}^3$, where $X = x$, $Y = y$, and $Z = x^2 + y^2$.

The first crucial observation is that a circle in the plane is lifted into a plane curve (an ellipse).

The intersection of the cylinder of revolution consisting of the lines parallel to the z-axis and passing through a point of the circle C with the paraboloid $z = x^2 + y^2$ is a planar curve (an ellipse).

We can compute the convex hull of the set of lifted points. Let us focus on the downward-facing faces of this convex hull.

Let $(l(p_1), l(p_2), l(p_3))$ be such a face. The points p_1, p_2, p_3 belong to the set P.

The circle C circumscribed about p_1, p_2, p_3 lifts to an ellipse passing through $(l(p_1), l(p_2), l(p_3))$.

We claim that no other point from P is inside the circle C.

Therefore, we have shown that *the projection of the part of the convex hull of the lifted set $l(P)$ consisting of the downward-facing faces is the Delaunay triangulation of P.*

Figure 8.8 shows the lifting of the Delaunay triangulation shown earlier.

Another example of the lifting of a Delaunay triangulation is shown in Figure 8.9.

The fact that a Delaunay triangulation can be obtained by projecting a lower convex hull can be used to find efficient algorithms for computing a Delaunay triangulation. It also holds for higher dimensions.

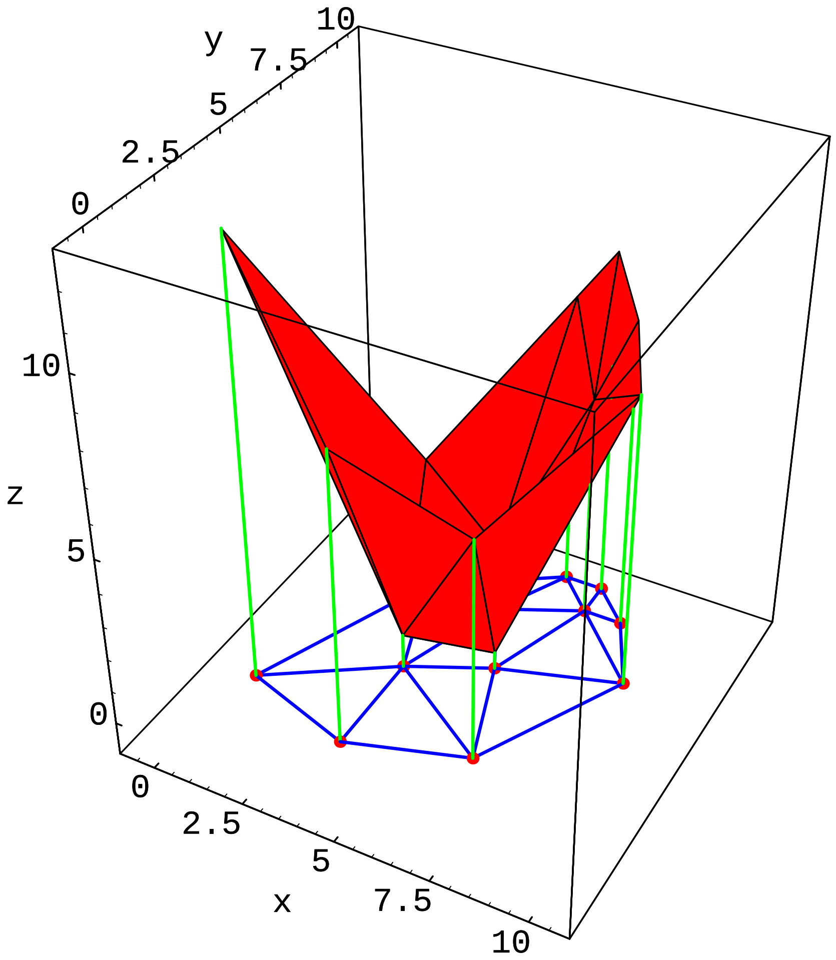

Figure 8.8: A Delaunay triangulation and its lifting to a paraboloid

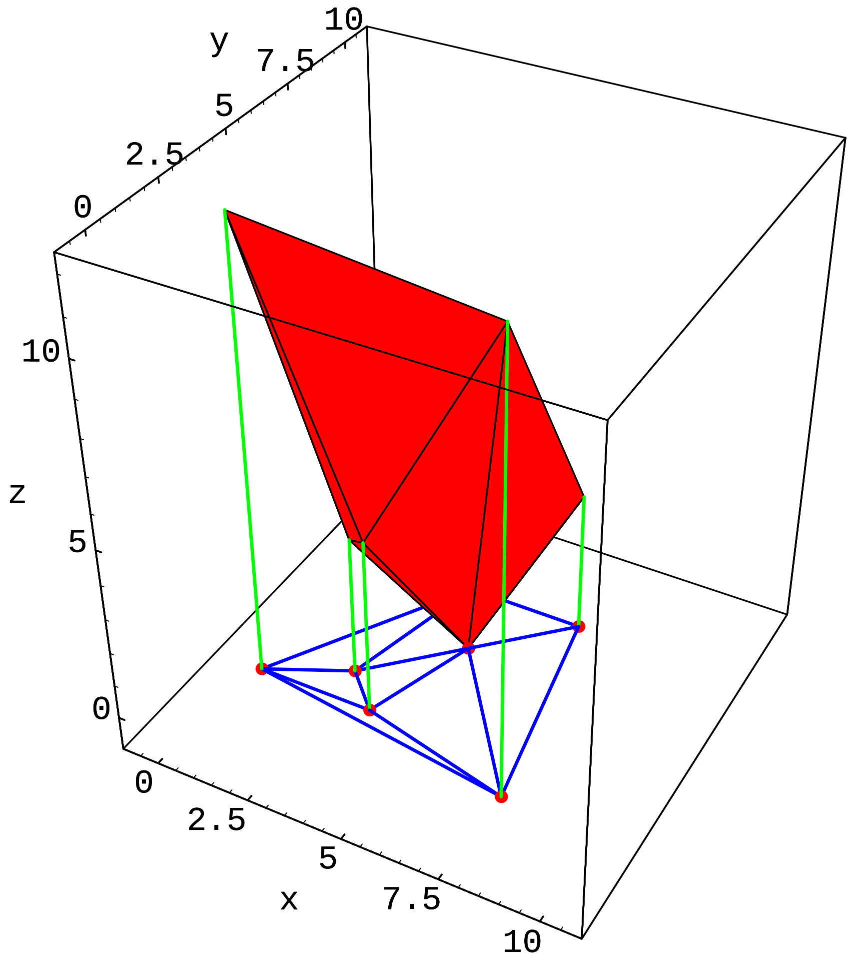

Figure 8.9: Another Delaunay triangulation and its lifting to a paraboloid

The Voronoi diagram itself can also be obtained from the lifted set $l(P)$.

However, this time, we need to consider tangent planes to the paraboloid at the lifted points.

It is fairly obvious that the tangent plane at the lifted point $(a, b, a^2 + b^2)$ is

$$z = 2ax + 2by - (a^2 + b^2).$$

Given two distinct lifted points $(a_1, b_1, a_1^2 + b_1^2)$ and $(a_2, b_2, a_2^2 + b_2^2)$, the intersection of the tangent planes at these points is the line of equation

$$(b_1 - a_1)x + (b_2 - a_2)y = (b_1^2 + b_2^2)/2 - (a_1^2 + a_2^2)/2.$$

This is precisely the equation of the bisector line of the two points (a_1, b_1) and (a_2, b_2).

Therefore, *if we look at the paraboloid from* $z = +\infty$ *(with the paraboloid transparent), the projection of the tangent planes at the lifted points is the Voronoi diagram*!

It should be noted that the "duality" between the Delaunay triangulation, which is the projection of the convex hull of the lifted set $l(P)$ viewed from $z = -\infty$, and the Voronoi diagram, which is the projection of the tangent planes at the lifted set $l(P)$ viewed from $z = +\infty$, is reminiscent of the polar duality with respect to a quadric.

The reader interested in algorithms for finding Voronoi diagrams and Delaunay triangulations is referred to O'Rourke [**?**], Preparata and Shamos [**?**], Boissonnat and Yvinec [**?**], de Berg, Van Kreveld, Overmars, and Schwarzkopf [**?**], and Risler [**?**].

We conclude our brief presentation of Voronoi diagrams and Delaunay triangulations with a short section on applications.

8.5 Applications of Voronoi Diagrams and Delaunay Triangulations

The examples below are taken from O'Rourke [**?**]. Other examples can be found in Preparata and Shamos [**?**], Boissonnat and Yvinec [**?**], and de Berg, Van Kreveld, Overmars, and Schwarzkopf [**?**].

The first example is the *nearest neighbors* problem. There are actually two subproblems: *Nearest neighbor queries* and *all nearest neighbors*.

The nearest neighbor queries problem is as follows: Given a set P of points and a query point q, find the nearest neighbor(s) of q in P.

This problem can be solved by computing the Voronoi diagram of P and determining in which Voronoi region q falls.

This last problem, called *point location*, has been heavily studied (see O'Rourke [**?**]).

The all neighbors problem is as follows: Given a set P of points, find the nearest neighbor(s) to all points in P.

This problem can be solved by building a graph, the *nearest neighbor graph*, for short *nng*. The nodes of this undirected graph are the points in P, and there is an arc from p to q iff p is a nearest neighbor of q or vice versa. Then it can be shown that this graph is contained in the Delaunay triangulation of P.

The second example is the *largest empty circle*.

Some practical applications of this problem are to locate a new store (to avoid competition), or to locate a nuclear plant as far as possible from a set of towns.

More precisely, the problem is as follows. Given a set P of points, find a largest empty circle whose center is in the (closed) convex hull of P, empty in that it contains no points from P inside it, and largest in the sense that there is no other circle with strictly larger radius.

The Voronoi diagram of P can be used to solve this problem. It can be shown that if the center p of a largest empty circle is strictly inside the convex hull of P, then p coincides with a Voronoi vertex.

However, not every Voronoi vertex is a good candidate. It can also be shown that if the center p of a largest empty circle lies on the boundary of the convex hull of P, then p lies on a Voronoi edge.

The third example is the *minimum spanning tree*.

Given a graph G, a minimum spanning tree of G is a subgraph of G that is a tree, contains every vertex of the graph G, and minimizes the sum of the lengths of the tree edges.

It can be shown that a minimum spanning tree is a subgraph of the Delaunay triangulation of the vertices of the graph. This can be used to improve algorithms for finding minimum spanning trees, for example Kruskal's algorithm (see O'Rourke [**?**]).

We conclude by mentioning that Voronoi diagrams have applications to *motion planning*.

For example, consider the problem of moving a disk on a plane while avoiding a set of polygonal obstacles. If we "extend" the obstacles by the diameter of the disk, the problem reduces to finding a collision–free path between two points in the extended obstacle space.

One needs to generalize the notion of a Voronoi diagram. Indeed, we need to define the distance to an object, and medial curves (consisting of points equidistant to two objects) may no longer be straight lines.

A collision–free path with maximal clearance from the obstacles can be found by moving along the edges of the generalized Voronoi diagram.

This is an active area of research in robotics. For more on this topic, see O'Rourke [**?**].