

CMSC 754

Computational Geometry¹

David M. Mount
Department of Computer Science
University of Maryland
Fall 2005

¹Copyright, David M. Mount, 2005, Dept. of Computer Science, University of Maryland, College Park, MD, 20742. These lecture notes were prepared by David Mount for the course CMSC 754, Computational Geometry, at the University of Maryland. Permission to use, copy, modify, and distribute these notes for educational purposes and without fee is hereby granted, provided that this copyright notice appear in all copies.

Lecture 10: Voronoi Diagrams and Fortune's Algorithm

Reading: Chapter 7 in the 4M's.

Voronoi Diagrams: Voronoi diagrams are among the most important structures in computational geometry. A Voronoi diagram encodes proximity information, that is, what is close to what. Let $P = \{p_1, p_2, \dots, p_n\}$ be a set of points in the plane (or in any dimensional space), which we call *sites*. Define $\mathcal{V}(p_i)$, the *Voronoi cell* for p_i , to be the set of points q in the plane that are closer to p_i than to any other site. That is, the Voronoi cell for p_i is defined to be:

$$\mathcal{V}(p_i) = \{q \mid \|p_i q\| < \|p_j q\|, \forall j \neq i\},$$

where $\|pq\|$ denotes the Euclidean distance between points p and q . The Voronoi diagram can be defined over any metric and in any dimension, but we will concentrate on the planar, Euclidean case here.

Another way to define $\mathcal{V}(p_i)$ is in terms of the intersection of halfplanes. Given two sites p_i and p_j , the set of points that are strictly closer to p_i than to p_j is just the *open halfplane* whose bounding line is the perpendicular bisector between p_i and p_j . Denote this halfplane $h(p_i, p_j)$. It is easy to see that a point q lies in $\mathcal{V}(p_i)$ if and only if q lies within the intersection of $h(p_i, p_j)$ for all $j \neq i$. In other words,

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

Since the intersection of halfplanes is a (possibly unbounded) convex polygon, it is easy to see that $\mathcal{V}(p_i)$ is a (possibly unbounded) convex polygon. Finally, define the *Voronoi diagram* of P , denoted $\text{Vor}(P)$ to be what is left of the plane after we remove all the (open) Voronoi cells. It is not hard to prove (see the text) that the Voronoi diagram consists of a collection of line segments, which may be unbounded, either at one end or both. An example is shown in the figure below.




Figure 34: Voronoi diagram

Voronoi diagrams have a number of important applications. These include:

Nearest neighbor queries: One of the most important data structures problems in computational geometry is solving nearest neighbor queries. Given a point set P , and given a query point q , determine the closest point in P to q . This can be answered by first computing a Voronoi diagram and then locating the cell of the diagram that contains q . (We have already discussed point location algorithms.)

Computational morphology: Some of the most important operations in morphology (used very much in computer vision) is that of “growing” and “shrinking” (or “thinning”) objects. If we grow a collection of points, by imagining a grass fire starting simultaneously from each point, then the places where the grass fires meet will be along the Voronoi diagram. The *medial axis* of a shape (used in computer vision) is just a Voronoi diagram of its boundary.

Facility location: We want to open a new Blockbuster video. It should be placed as far as possible from any existing video stores. Where should it be placed? It turns out that the vertices of the Voronoi diagram are the points that are locally at maximum distances from any other point in the set, and hence they form a natural set of candidate locations.

Neighbors and Interpolation: Given a set of measured height values over some geometric terrain. Each point has (x, y) coordinates and a height value. We would like to interpolate the height value of some query point that is not one of our measured points. To do so, we would like to interpolate its value from neighboring measured points. One way to do this, called *natural neighbor interpolation*, is based on computing the Voronoi neighbors of the query point, assuming that it has one of the original set of measured points.

Properties of the Voronoi diagram: Here are some observations about the structure of Voronoi diagrams in the plane.

Voronoi complex: Clearly the diagram is a cell complex, whose faces are convex polygons. Each point on an edge of the Voronoi diagram is equidistant from its two nearest neighbors p_i and p_j . Thus, there is a circle centered at such a point such that p_i and p_j lie on this circle, and no other site is interior to the circle.




Figure 35: Properties of the Voronoi diagram.

Voronoi vertices: It follows that the vertex at which three Voronoi cells $\mathcal{V}(p_i)$, $\mathcal{V}(p_j)$, and $\mathcal{V}(p_k)$ intersect, called a *Voronoi vertex* is equidistant from all sites. Thus it is the center of the circle passing through these sites, and this circle contains no other sites in its interior.

Degree: If we make the general position assumption that no four sites are cocircular, then the vertices of the Voronoi diagram all have degree three.

Convex hull: A cell of the Voronoi diagram is unbounded if and only if the corresponding site lies on the convex hull. (Observe that a site is on the convex hull if and only if it is the closest point from some point at infinity.) Thus, given a Voronoi diagram, it is easy to extract the convex hull in linear time.

Size: If n denotes the number of sites, then the Voronoi diagram is a planar graph (if we imagine all the unbounded edges as going to a common vertex infinity) with exactly n faces. It follows from Euler's formula that the number of Voronoi vertices is at most $2n - 5$ and the number of edges is at most $3n - 6$. (See the text for details. In higher dimensions the diagram's combinatorial complexity can be as high as $O(n^{\lceil d/2 \rceil})$.)

Computing Voronoi Diagrams: There are a number of algorithms for computing Voronoi diagrams. Of course, there is a naive $O(n^2 \log n)$ time algorithm, which operates by computing $\mathcal{V}(p_i)$ by intersecting the $n - 1$ bisector halfplanes $h(p_i, p_j)$, for $j \neq i$. However, there are much more efficient ways, which run in $O(n \log n)$ time. Since the convex hull can be extracted from the Voronoi diagram in $O(n)$ time, it follows that this is asymptotically optimal in the worst-case.

Historically, $O(n^2)$ algorithms for computing Voronoi diagrams were known for many years (based on incremental constructions). When computational geometry came along, a more complex, but asymptotically superior $O(n \log n)$ algorithm was discovered.

Lecture 11: Variants of Voronoi Diagrams

Reading: This material is not covered in our text.

Voronoi Diagrams: As we saw in the previous lecture, a Voronoi diagram encodes proximity information, that is, what is close to what. Given a finite set $P = \{p_1, p_2, \dots, p_n\}$ of points in space, called *sites*, the Voronoi diagram is the complex that results when space is subdivided according to which site of P is closest. This suggests a number of ways of generalizing this concept. The first arises by considering not just the closest site, but more generally the k -closest sites for some $k \geq 2$. The second arises by considering the closest site but according to other distance functions.

The Voronoi diagram of order k : Let P_k be any subset of size k of P . The Voronoi cell of P_k is defined to be the set of points x in the plane such that all sites of P_k are closer to x than any of the other sites of P . That is, define the *Voronoi cell* of P_k to be

$$\mathcal{V}(P_k) = \{x \in \mathbb{R}^d \mid \forall p \in P_k, \forall q \in (P \setminus P_k), \|xp\| < \|xq\|\},$$

(where the notation $P \setminus Q$ denotes the subset of points of P that are not in Q). Observe that the standard Voronoi diagram is just the Voronoi diagram of order 1.

One way to visualize whether a point x is in $\mathcal{V}(P_k)$ is to imagine growing a circle centered at x until it contains exactly k sites of P . If these k sites are the sites of P_k , then x lies in $\mathcal{V}(P_k)$. Observe that not every k -element subset will have a nonempty Voronoi cell. For example, if the convex hull of three sites $P' = \{p_1, p_2, p_3\}$ contains the site p_4 , then any circle containing these three sites must contain p_4 , implying that the Voronoi cell of P' is empty. The Voronoi diagrams of various orders are shown in Fig. 40

Recall that $h(p, q)$ denotes the bisecting (open) halfspace of points that are strictly closer to p than to q . It is easy to see that

$$\mathcal{V}(P_k) = \bigcap_{\substack{p \in P_k \\ q \in P \setminus P_k}} h(p, q).$$

Since the intersection of convex sets is convex, this implies that $\mathcal{V}(P_k)$ is a convex body, and in fact a convex polyhedron (which may be unbounded). In fact, if all the diagrams of all orders are overlaid upon one another, the result is an arrangement consisting of all the bisectors. (See Fig. 40.)

What is the combinatorial complexity of the Voronoi diagram of order k ? First observe that there are $\binom{n}{k} = \Theta(n^k)$ distinct subsets of size k , and therefore there are at most $O(n^k)$ faces in the Voronoi diagram of order k . However, we mentioned earlier that not all subsets of sites have a nonempty Voronoi cell, and this upper bound

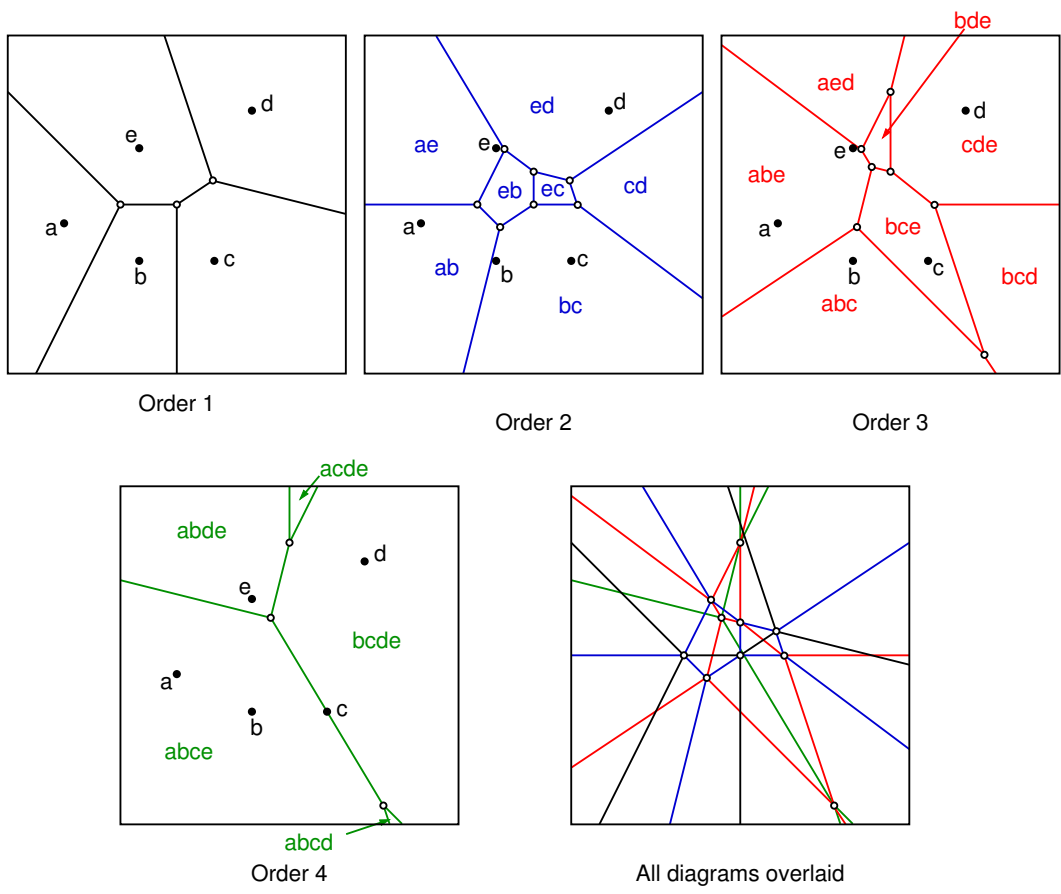


Figure 40: Higher order Voronoi diagrams.

is horribly pessimistic. In the plane it can be shown that the order- k Voronoi diagram complexity has complexity $O(k(n-k))$. The total complexity of all the of the Voronoi diagrams for k ranging from 1 to $n-1$ is $O(n^3)$.

Farthest-Site Voronoi diagram: An interesting and important special case arises when $k = n - 1$. In this case, each Voronoi cell corresponds to having the same $n - 1$ closest sites, that is, having the same farthest site. For this reason, the Voronoi diagram of order $n - 1$ is called the *farthest-site Voronoi diagram*. (Actually, it is more often called the *furthest-site Voronoi diagram*, but this is not really proper English.)

Clearly the farthest-site Voronoi diagram consists of at most n cells (each corresponding to the points that are farthest away from each of the n sites). However, not all the sites are nonempty. For example, in the example above, where p_4 was in the convex hull of $P' = \{p_1, p_2, p_3\}$, there is no point in the plane that is farthest away from p_4 . It is not hard to prove that a point has a nonempty farthest-site Voronoi cell if and only if the site is a Vertex of the convex hull of P .

Different Distance Functions: The other common way of designing variants on Voronoi diagrams is by altering the notion of distance. One way to do this is by altering the underlying metric. The definition of Voronoi diagram generalizes readily to any metric. There are many ways in which to define a metric. An important class of metrics are the *Minkowski metrics*. The L_k Minkowski distance between points p and q is defined to be

$$\left(|p_1 - q_1|^k + |p_2 - q_2|^k + \dots + |p_d - q_d|^k\right)^{\frac{1}{k}}.$$

Observe that the L_2 distance is just the familiar *Euclidean distance*. The L_1 distance is called the *Manhattan distance* (or *chessboard distance*) because it corresponds to distance along a horizontal/vertical grid. In the limit as k tends to ∞ , this approaches the *max distance* (or *box distance*). The L_∞ distance can be formally defined to be

$$\max(|p_1 - q_1|, |p_2 - q_2|, \dots, |p_d - q_d|).$$

Each of these distance metrics is characterized by a distinct unit ball. For the L_2 distance the unit metric ball is based on a circle (hypersphere in dimension d). For the L_∞ distance the unit metric ball is based on axis-parallel square (hypercube in dimension d). For the L_1 distance the unit metric ball is a 45° rotation of a square (2^d -sided polytope that generalizes an octahedron).

Weighted Point Sets: Sometimes it is useful to increase or decrease the weight of influence of each site in the diagram. This suggests assigning weights to the sites, and then basing the diagram on the notion of weighted distance. Here are some natural ways of modifying the Voronoi diagram based on weighted points.

Additive weights: Each site is assigned a weight w_i and the distance from a point q to site p_i is $\|qp_i\| - w_i$. In the plane the diagram's edges are arcs of hyperbolas.

Multiplicative weights: Each site is assigned a weight λ_i and the distance from a point q to site p_i is $\lambda_i \|qp_i\|$. In the plane the diagram's edges are arcs of circles. This is also called the *Apollonius diagram*.

Power diagram: Each site is assigned a weight w_i and the distance from a point q to site p_i is $\|qp_i\|^2 - w_i^2$. In the plane the diagram's edges are straight line segments.

Lecture 12: Delaunay Triangulations

Reading: Chapter 9 in the 4M's.

Delaunay Triangulations: Recently we discussed the topic of Voronoi diagrams. Today we consider the related structure, called a *Delaunay triangulation* (DT). The Voronoi diagram of a set of sites in the plane is a planar subdivision, that is, cell complex. The *dual* of such subdivision is another subdivision that is defined as follows. For each face of the Voronoi diagram, we create a vertex (corresponding to the site). For each edge of the

Voronoi diagram lying between two sites p_i and p_j , we create an edge in the dual connecting these two vertices. Finally, each vertex of the Voronoi diagram corresponds to a face of the dual.

The resulting dual graph is a planar subdivision. Assuming general position, the vertices of the Voronoi diagram have degree three, it follows that the faces of the resulting dual graph (excluding the exterior face) are triangles. Thus, the resulting dual graph is a triangulation of the sites, called the *Delaunay triangulation*. This is illustrated in the figure below.




Figure 41: The Delaunay triangulation of a set of points (solid lines) and the Voronoi diagram (broken lines).

Delaunay triangulations have a number of interesting properties, that are consequences of the structure of the Voronoi diagram.

Convex hull: The boundary of the exterior face of the Delaunay triangulation is the boundary of the convex hull of the point set.

Circumcircle property: The circumcircle of any triangle in the Delaunay triangulation is empty (contains no sites of P).

Proof: This is because the center of this circle is the corresponding dual Voronoi vertex, and by definition of the Voronoi diagram, the three sites defining this vertex are its nearest neighbors.

Empty circle property: Two sites p_i and p_j are connected by an edge in the Delaunay triangulation, if and only if there is an empty circle passing through p_i and p_j .

Proof: If two sites p_i and p_j are neighbors in the Delaunay triangulation, then their cells are neighbors in the Voronoi diagram, and so for any point on the Voronoi edge between these sites, a circle centered at this point passing through p_i and p_j cannot contain any other point (since they must be closest). Conversely, if there is an empty circle passing through p_i and p_j , then the center c of this circle is a point on the edge of the Voronoi diagram between p_i and p_j , because c is equidistant from each of these sites and there is no closer site. Thus the Voronoi cells of two sites are adjacent in the Voronoi diagram, implying that there edge is in the Delaunay triangulation.

Closest pair property: The closest pair of sites in P are neighbors in the Delaunay triangulation.

Proof: Suppose that p_i and p_j are the closest sites. The circle having p_i and p_j as its diameter cannot contain any other site, since otherwise such a site would be closer to one of these two points, violating the hypothesis that these points are the closest pair. Therefore, the center of this circle is on the Voronoi edge between these points, and so it is an empty circle.

If the sites are not in general position, in the sense that four or more are cocircular, then the Delaunay triangulation may not be a triangulation at all, but just a planar graph (since the Voronoi vertex that is incident to four or more Voronoi cells will induce a face whose degree is equal to the number of such cells). In this case the more appropriate term would be *Delaunay graph*. However, it is common to either assume the sites are in general position (or to enforce it through some sort of symbolic perturbation) or else to simply triangulate the faces of degree four or more in any arbitrary way. Henceforth we will assume that sites are in general position, so we do not have to deal with these messy situations.

Maximizing Angles and Edge Flipping: Another interesting property of Delaunay triangulations is that among all triangulations, the Delaunay triangulation maximizes the minimum angle. This property is important, because it implies that Delaunay triangulations tend to avoid skinny triangles. This is useful for many applications where triangles are used for the purposes of interpolation.

In fact a much stronger statement holds as well. Among all triangulations with the same smallest angle, the Delaunay triangulation maximizes the second smallest angle, and so on. In particular, any triangulation can be associated with a sorted *angle sequence*, that is, the increasing sequence of angles $(\alpha_1, \alpha_2, \dots, \alpha_m)$ appearing in the triangles of the triangulation. (Note that the length of the sequence will be the same for all triangulations of the same point set, since the number depends only on n and h .)

Theorem: Among all triangulations of a given planar point set, the Delaunay triangulation has the lexicographically largest angle sequence, and in particular, it maximizes the minimum angle.

Before getting into the proof, we should recall a few basic facts about angles from basic geometry. First, recall that if we consider the circumcircle of three points, then each angle of the resulting triangle is exactly half the angle of the minor arc subtended by the opposite two points along the circumcircle. It follows as well that if a point is inside this circle then it will subtend a larger angle and a point that is outside will subtend a smaller angle. This in the figure part (a) below, we have $\theta_1 > \theta_2 > \theta_3$.




Figure 43: Angles and edge flips.

We will not give a formal proof of the theorem. (One appears in the text.) The main idea is to show that for any triangulation that fails to satisfy the empty circle property, it is possible to perform a local operation, called an *edge flip*, which increases the lexicographical sequence of angles. An edge flip is an important fundamental operation on triangulations in the plane. Given two adjacent triangles $\triangle abc$ and $\triangle cda$, such that their union forms a convex quadrilateral $abcd$, the edge flip operation replaces the diagonal ac with bd . Note that it is only possible when the quadrilateral is convex. Suppose that the initial triangle pair violates the empty circle condition, in that point d lies inside the circumcircle of $\triangle abc$. (Note that this implies that b lies inside the circumcircle of $\triangle cda$.) If we flip the edge it will follow that the two circumcircles of the two resulting triangles, $\triangle abd$ and $\triangle bcd$ are now empty (relative to these four points), and the observation above about circles and angles proves that the minimum angle increases at the same time. In particular, in the figure above, we have

$$\phi_{ab} > \theta_{ab} \quad \phi_{bc} > \theta_{bc} \quad \phi_{cd} > \theta_{cd} \quad \phi_{da} > \theta_{da}.$$

There are two other angles that need to be compared as well (can you spot them?). It is not hard to show that, after swapping, these other two angles cannot be smaller than the minimum of $\theta_{ab}, \theta_{bc}, \theta_{cd}$, and θ_{da} . (Can you see why?)

Since there are only a finite number of triangulations, this process must eventually terminate with the lexicographically maximum triangulation, and this triangulation must satisfy the empty circle condition, and hence is the Delaunay triangulation.

Lecture 29: Delaunay Triangulations and Convex Hulls

Reading: O'Rourke 5.7 and 5.8.

Delaunay Triangulations and Convex Hulls: At first, Delaunay triangulations and convex hulls appear to be quite different structures, one is based on metric properties (distances) and the other on affine properties (collinearity, coplanarity). Today we show that it is possible to convert the problem of computing a Delaunay triangulation in dimension d to that of computing a convex hull in dimension $d + 1$. Thus, there is a remarkable relationship between these two structures.

We will demonstrate the connection in dimension 2 (by computing a convex hull in dimension 3). Some of this may be hard to visualize, but see O'Rourke for illustrations. (You can also reason by analogy in one lower dimension of Delaunay triangulations in 1-d and convex hulls in 2-d, but the real complexities of the structures are not really apparent in this case.)

The connection between the two structures is the *paraboloid* $z = x^2 + y^2$. Observe that this equation defines a surface whose vertical cross sections (constant x or constant y) are parabolas, and whose horizontal cross sections (constant z) are circles. For each point in the plane, (x, y) , the *vertical projection* of this point onto this paraboloid is $(x, y, x^2 + y^2)$ in 3-space. Given a set of points S in the plane, let S' denote the projection of every point in S onto this paraboloid. Consider the *lower convex hull* of S' . This is the portion of the convex hull of S' which is visible to a viewer standing at $z = -\infty$. We claim that if we take the lower convex hull of S' , and project it back onto the plane, then we get the Delaunay triangulation of S . In particular, let $p, q, r \in S$, and let p', q', r' denote the projections of these points onto the paraboloid. Then $p'q'r'$ define a *face* of the lower convex hull of S' if and only if $\triangle pqr$ is a triangle of the Delaunay triangulation of S . The process is illustrated in the following figure.

The question is, why does this work? To see why, we need to establish the connection between the triangles of the Delaunay triangulation and the faces of the convex hull of transformed points. In particular, recall that

Delaunay condition: Three points $p, q, r \in S$ form a Delaunay triangle if and only if the circumcircle of these points contains no other point of S .

Convex hull condition: Three points $p', q', r' \in S'$ form a face of the convex hull of S' if and only if the plane passing through p', q' , and r' has all the points of S' lying to one side.




Figure 100: Delaunay triangulations and convex hull.

Clearly, the connection we need to establish is between the emptiness of circumcircles in the plane and the emptiness of halfspaces in 3-space. We will prove the following claim.

Lemma: Consider 4 distinct points p, q, r, s in the plane, and let p', q', r', s' be their respective projections onto the paraboloid, $z = x^2 + y^2$. The point s lies within the circumcircle of p, q, r if and only if s' lies on the lower side of the plane passing through p', q', r' .

To prove the lemma, first consider an arbitrary (nonvertical) plane in 3-space, which we assume is tangent to the paraboloid above some point (a, b) in the plane. To determine the equation of this tangent plane, we take derivatives of the equation $z = x^2 + y^2$ with respect to x and y giving

$$\frac{\partial z}{\partial x} = 2x, \quad \frac{\partial z}{\partial y} = 2y.$$

At the point $(a, b, a^2 + b^2)$ these evaluate to $2a$ and $2b$. It follows that the plane passing through these point has the form

$$z = 2ax + 2by + \gamma.$$

To solve for γ we know that the plane passes through $(a, b, a^2 + b^2)$ so we solve giving

$$a^2 + b^2 = 2a \cdot a + 2b \cdot b + \gamma,$$

Implying that $\gamma = -(a^2 + b^2)$. Thus the plane equation is

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

If we shift the plane upwards by some positive amount r^2 we get the plane

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

How does this plane intersect the paraboloid? Since the paraboloid is defined by $z = x^2 + y^2$ we can eliminate z giving

$$x^2 + y^2 = 2ax + 2by - (a^2 + b^2) + r^2,$$

which after some simple rearrangements is equal to

$$(x - a)^2 + (y - b)^2 = r^2.$$

This is just a circle. Thus, we have shown that the intersection of a plane with the paraboloid produces a space curve (which turns out to be an ellipse), which when projected back onto the (x, y) -coordinate plane is a circle centered at (a, b) .

Thus, we conclude that the intersection of an arbitrary lower halfspace with the paraboloid, when projected onto the (x, y) -plane is the interior of a circle. Going back to the lemma, when we project the points p, q, r onto the paraboloid, the projected points p', q' and r' define a plane. Since p', q' , and r' lie at the intersection of the plane and paraboloid, the original points p, q, r lie on the projected circle. Thus this circle is the (unique) circumcircle passing through these p, q , and r . Thus, the point s lies within this circumcircle, if and only if its projection s' onto the paraboloid lies within the lower halfspace of the plane passing through p, q, r .




Figure 101: Planes and circles.

Now we can prove the main result.

Theorem: Given a set of points S in the plane (assume no 4 are cocircular), and given 3 points $p, q, r \in S$, the triangle $\triangle pqr$ is a triangle of the Delaunay triangulation of S if and only if triangle $\triangle p'q'r'$ is a face of the lower convex hull of the projected set S' .

From the definition of Delaunay triangulations we know that $\triangle pqr$ is in the Delaunay triangulation if and only if there is no point $s \in S$ that lies within the circumcircle of pqr . From the previous lemma this is equivalent to saying that there is no point s' that lies in the lower convex hull of S' , which is equivalent to saying that $p'q'r'$ is a face of the lower convex hull. This completes the proof.

In order to test whether a point s lies within the circumcircle defined by p, q, r , it suffices to test whether s' lies within the lower halfspace of the plane passing through p', q', r' . If we assume that p, q, r are oriented counterclockwise in the plane this reduces to determining whether the quadruple p', q', r', s' is positively oriented, or equivalently whether s lies to the left of the oriented circle passing through p, q, r .

This leads to the incircle test we presented last time.

$$\text{in}(p, q, r, s) = \det \begin{pmatrix} p_x & p_y & p_x^2 + p_y^2 & 1 \\ q_x & q_y & q_x^2 + q_y^2 & 1 \\ r_x & r_y & r_x^2 + r_y^2 & 1 \\ s_x & s_y & s_x^2 + s_y^2 & 1 \end{pmatrix} > 0.$$

Voronoi Diagrams and Upper Envelopes: We know that Voronoi diagrams and Delaunay triangulations are dual geometric structures. We have also seen (informally) that there is a dual relationship between points and lines in the plane, and in general, points and planes in 3-space. From this latter connection we argued that the problems of computing convex hulls of point sets and computing the intersection of halfspaces are somehow “dual” to one another. It turns out that these two notions of duality, are (not surprisingly) interrelated. In particular, in the

same way that the Delaunay triangulation of points in the plane can be transformed to computing a convex hull in 3-space, it turns out that the Voronoi diagram of points in the plane can be transformed into computing the intersection of halfspaces in 3-space.

Here is how we do this. For each point $p = (a, b)$ in the plane, recall the tangent plane to the paraboloid above this point, given by the equation

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

Define $H^+(p)$ to be the set of points that are above this halfplane, that is, $H^+(p) = \{(x, y, z) \mid z \geq 2ax + 2by - (a^2 + b^2)\}$. Let $S = \{p_1, p_2, \dots, p_n\}$ be a set of points. Consider the intersection of the halfspaces $H^+(p_i)$. This is also called the *upper envelope* of these halfspaces. The upper envelope is an (unbounded) convex polyhedron. If you project the edges of this upper envelope down into the plane, it turns out that you get the Voronoi diagram of the points.

Theorem: Given a set of points S in the plane (assume no 4 are cocircular), let H denote the set of upper halfspaces defined by the previous transformation. Then the Voronoi diagram of H is equal to the projection onto the (x, y) -plane of the 1-skeleton of the convex polyhedron which is formed from the intersection of halfspaces of S' .

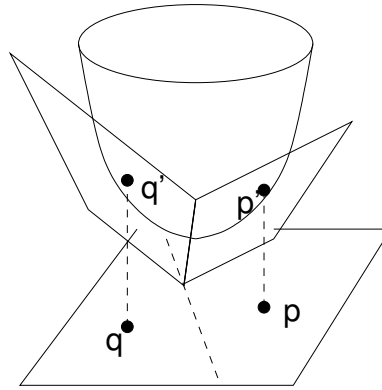


Figure 102: Intersection of halfspaces.

It is hard to visualize this surface, but it is not hard to show why this is so. Suppose we have 2 points in the plane, $p = (a, b)$ and $q = (c, d)$. The corresponding planes are:

$$z = 2ax + 2by - (a^2 + b^2) \quad \text{and} \quad z = 2cx + 2dy - (c^2 + d^2).$$

If we determine the intersection of the corresponding planes and project onto the (x, y) -coordinate plane (by eliminating z from these equations) we get

$$x(2a - 2c) + y(2b - 2d) = (a^2 - c^2) + (b^2 - d^2).$$

We claim that this is the perpendicular bisector between (a, b) and (c, d) . To see this, observe that it passes through the midpoint $((a + c)/2, (b + d)/2)$ between the two points since

$$\frac{a + c}{2}(2a - 2c) + \frac{b + d}{2}(2b - 2d) = (a^2 - c^2) + (b^2 - d^2).$$

and, its slope is $-(a - c)/(b - d)$, which is the negative reciprocal of the line segment from (a, b) to (c, d) . From this it can be shown that the intersection of the upper halfspaces defines a polyhedron whose edges project onto the Voronoi diagram edges.