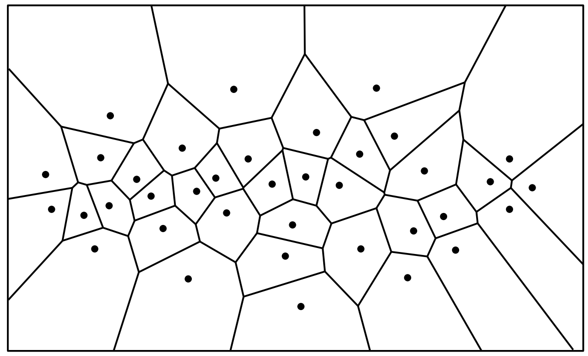

Figure 1. Voronoi Diagram of point set P, denoted $\mathrm{Vor(P)}$

Note: These lecture notes are based on the textbook "Computational Geometry" by Berg et al.*and lecture notes from [3]*

1 Voronoi Diagrams

A Voronoi diagram encodes proximity information, i.e. what is close to what. Let P be a set of points (henceforth sites) in the plane (or more generally in $\mathbb{R}^d$), and for any two sites p, q, let $|p - q| = \left(\sum_{i=1}^{d}(p_i - q_i)^2\right)^{1/2}$ be the Euclidean distance between p and q. Define $\mathcal{V}(p)$, the *Voronoi cell* for p to be the set of points q in the plane that are closer to site p than to any other site $p' \in P$. More formally:

$$\mathcal{V}(p) = \{q \in \mathbb{R}^d : |q - p| < |q - p'|, \forall p' \in P - p\}$$

The union of the closures of the Voronoi cells defines a cell complex called the *Voronoi diagram* of P, denoted Vor(P).

1.1 Convex Polyhedra

The cells of the Voronoi diagram are possibly unbounded convex polyhedra. To see this, fix two sites $p, p' \in P$ and note that the set of points closer to p than p' is equal to an open halfplane, whose bounding hyperplane is the perpendicular bisector of the line segment $\overline{pp'}$. Denote this halfplane $h(p, p')$. Now note that

$$\mathcal{V}(p) = \bigcap_{q \in P - p} h(p, q)$$

Since the intersection of convex spaces is convex and the intersection of polyhedra are polyhedra, $\mathcal{V}(p)$ is a (possibly unbounded) convex polyhedra.

1.2 Applications

Imagine the sites P are post offices and we want to compute **regions best served by each post office**, when cost to visit a site is a linear function of the Euclidean distance to that site.[1] The Voronoi diagram Vor(P) exactly delineates those regions. Moreover, it's not hard to generalize the concept of the Voronoi diagram to the case where costs to sites may be determined by *different* linear

[1]this also works with e.g. supermarkets

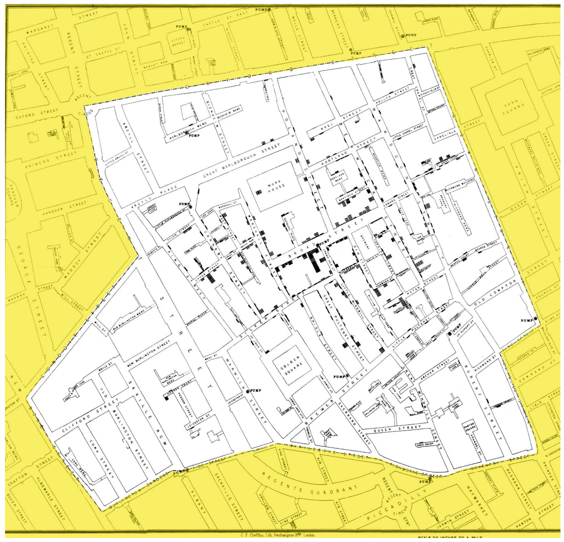

Figure 2. Voronoi Cell of the water pump responsible for the 1854 cholera outbreak in London. Bars represent cholera deaths.

functions of the Euclidean distance. Voronoi diagrams have also been used by "anthropologists to describe **regions of influence of different cultures**; by crystallographers to **explain the structure of certain crystals and metals**; by ecologists to **study competition between plants**; and by economists to **model markets in the U.S. economy**." [1] They were even used by John Snow to **isolate the pump responsible for the 1854 London Cholera outbreak**. (See Figure 2.)

Nearest Neighbor Queries. A final critical application is fast nearest neighbor queries. After construction and processing of a Voronoi diagram of point set P, we can answer nearest neighbor queries in $O(\log n)$ time. The easiest way to do this is to first construct vertical slabs that are delineated by the vertices in the Voronoi diagram. Sort these slabs by x-coordinate so that in logarithmic time, it is possible to find which slab a new point falls in. Next, sort the line segments that intersect each slab along y-coordinates, so that given the slab that a point falls in, in logarithmic time it is possible to find which cell of the Voronoi diagram that point falls in. Since the number of slabs and line segments are polynomial (See Theorem 2), lookups take logarithmic time.[2]

1.3 Properties

Theorem 1. *Let P be a set of n sites in the plane. If all sites are colinear then Vor(P) consists of $n-1$ parallel lines. Otherwise Vor(P) is connected and its edges are all either line segments or half lines.*

Proof: The first part holds trivially, so assume not all sites are colinear.

We now show that the edges of Vor(P) are either line segments or half-lines. Note that the edges of Vor(P) are parts of straight lines. Assume by way of contradiction that there is an edge e that is a full line. Let e be on the boundary of the Voronoi cells $\mathcal{V}(p_1)$ and $\mathcal{V}(p_2)$ for $p_1, p_2 \in P$. Let p_3 be a point that is not collinear with p_1 and p_2. The bisector of p_2 and p_3 is not parallel to e and so it intersects e. But then the part of e that lies in the interior of $h(p_3, p_2)$ cannot be on the boundary of $\mathcal{V}(p_2)$, since it is closer to p_3 than p_2, a contradiction.

[2]Note that there are more space-efficient ways to do point lookup using something called trapezoidal maps.

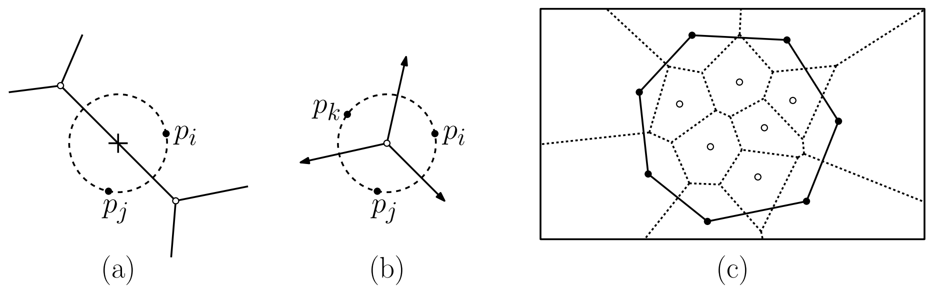

Figure 3. Properties of Voronoi Diagram

Now assume that Vor(P) is not connected. Then, there would be a Voronoi cell $\mathcal{V}(p_1)$ that splits the plan in two. Since Voronoi cells are convex, $\mathcal{V}(p_1)$ would be a vertical strip bounded by two parallel lines. But we just proved that no edge of the Voronoi diagram can be a line, a contradiction. $\square$

Theorem 2. *The number of vertices, faces and edges in Vor(P) are all $O(n)$.*

Proof: Create a planar graph from the Vor(P) by adding an extra vertex v_∞ at "infinity", and connecting all half-infinite edges of Vor(P) to this vertex.

Recall that Euler's formula says that $(V+1)-E+F=2$ (we add one to V because of v_∞) If the number of sites is n (i.e. $|P|=n$) then the number of faces in the Voronoi diagram is n. The degree of each Voronoi vertex is at least 3, since if a Voronoi vertex had degree 2, one of the neighboring Voronoi cells would be concave, which is not possible. Recall that the sum of all degrees is twice the number of edges (handshaking lemma). Thus, $2E \geq 3(V+1)$. So we have

$$\begin{aligned} V+1 &= (2-n)+E \\ &\geq (2-n)+(3/2)(V+1) \end{aligned}$$

From whence, we get $V+1 \leq 2(n-2)$. Using this to bound E, we get that $E \leq 3n-6$. $\square$

Theorem 3. *For the Voronoi Vor(P) of a set of sites P, the following hold.*

1. ***Voronoi vertices.*** *A point q is a vertex of Vor(P) iff its largest empty circle, called $C_P(q)$, contains three or more sites on its boundary.*

2. ***Voronoi edges.*** *The bisector between sites p_1 and p_2 defines an edge of Vor(P) iff there is a point q on the bisector such that $C_P(q)$ contains both p_1 and p_2 on its boundary but no other sites. (Then all such points q are on the Voronoi edge)*

Proof: For the first property, suppose there is a point q such that $C_P(q)$ contains three or more sites on its boundary. Let $p_1, \ldots p_\ell$ be these sites. Since the interior of $C_P(q)$ is empty, and since $p_1, \ldots p_\ell$ are equidistant from q since they lie on the circle centered at q, we know that point q must be on the boundary of each of $\mathcal{V}(p_1), \ldots, \mathcal{V}(p_\ell)$. Hence, point q is a Voronoi vertex of Vor(P).

On the other hand, assume point q is a vertex of Vor(P). Then q is incident to at least three edges, and hence incident to at least three Voronoi cells $\mathcal{V}(p_1), \mathcal{V}(p_2), \mathcal{V}(p_3)$, for $p_1, p_2, p_3 \in P$. Voronoi vertex q is equidistant to p_1, p_2, p_3 and there can not be a site closer to q. Hence, $C_P(q)$ is an empty circle containing three or more sites on its boundary. (See Figure 3 (b))

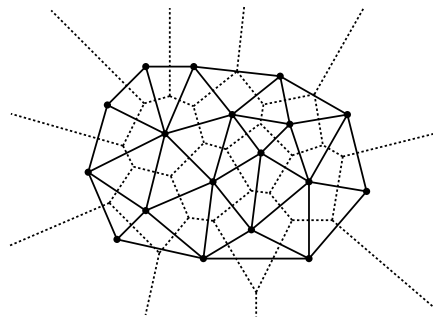

Figure 4. Delaunay triangulation of a set of points (solid lines) and the Voronoi diagram (dashed lines).

For the second property, suppose there is a point q on the bisector between sites p_1 and p_2 such that $C_P(q)$ contains p_1 and p_2 on its boundary but no other sites. Then, $dist(q, p_1) = dist(q, p_2) < dist(q, p_x)$ for any other site $p_x \in P - \{p_1, p_2\}$. Hence q lies on an edge of Vor(P) that is defined by the bisector of p_1 and p_2.

On the other hand, let the bisector of p_1 and p_2 define a Voronoi edge. Then the largest circle of any point q on this edge must contain p_1 and p_2 on its boundary and no other sites. (See Figure 3 (a))

□

1.3.1 Additional Properties

Degree. Three points in the plane define a unique circle. If we make a general position assumption that no four sites are cocircular, then each vertex of the Voronoi diagram is incident to 3 edges. In $\mathbb{R}^d$, the vertex is defined by $d + 1$ points in general position, and the hypersphere centered at the vertex passing through these sites is empty.

Convex Hull. A cell of the Voronoi diagram is unbounded iff the corresponding site is on the convex hull of P. To see this, note that a point is on the convex hull iff it is the closest point from some point at infinity. Thus, given a Voronoi diagram it is easy to compute the convex hull. Later we'll see the reverse is also true. (See Figure 3 (c).)

1.4 Computing the Voronoi Diagram

How can we efficiently compute a Voronoi diagram? There are several direct algorithms, but soon we'll see how we can do so using (higher-dimensional) convex hull algorithms.

2 Delaunay Graphs

We've shown that the Voronoi diagram is a planar subdivision that subdivides the plane into convex (possibly open) polyhedra.

The *Delaunay graph* is the dual[3] of the Voronoi graph defined as follows (See Figure 4)

- For each face of the Voronoi graph, we create a vertex corresponding to the face's site.

[3]This is another type of duality, different from the point/line duality of last lecture

- For each edge of the Voronoi graph lying between sites p and q, we create an edge in the dual connecting the two vertices associated with these sites.
- Thus, each vertex of the Voronoi graph corresponds to a face in the dual.

If no four points are on the same circle (i.e. general position). Then all vertices of the Voronoi have degree 3 and so all faces of the Delaunay graph are triangles. This is why the Delaunay graph is generally called the *Delaunay triangulation*.

2.1 Properties of Delaunay Triangulation

Theorem 4. *The Delaunay triangulation of n points P in the plane, forms a planar graph, where the number of faces, edges and vertices are all $O(n)$.*

Proof: This follows from Theorems 1 and 2 and the duality property between Delaunay graphs and Voronoi graphs. □

In 3-space, the number of tetrahedra in the Delaunay triangulation can range from $O(n)$ to $O(n^2)$. In dimension d, the number of simplices (d dimensional generalization of a triangle) can be $\Theta(n^{\lceil d/2 \rceil})$.

Theorem 5. *Let P be a set of sites in the plane. Then*

1. *Three or more sites $p_1, \ldots, p_\ell \in P$ are vertices of the same face of the Delaunay graph of P iff the circle through $p_1, \ldots, p_\ell$ contains no site of P in its interior*
2. *Two sites $p_1, p_2 \in P$ form an edge of the Delaunay graph of P iff there is a circle that has p_1 and p_2 on its boundary and contains no other sites of P*

Proof: By Theorem 3 (1), we know that a point q is a vertex of Vor(P) iff its largest empty circle $C_p(q)$ contains three or more sites on its boundary. Since Voronoi vertices become faces in the Delaunay triangulation, this translates to part (1) of the theorem statement.

By Theorem 3 (2), we also know that the bisector between sites p_1 and p_2 defines an edge of Vor(P) iff there is a point q on the bisector such that there is a circle with p_1 and p_2 on its boundary but containing no other sites. Since Voronoi edges associated with bisectors between p_1 and p_2 become Delaunay edges connecting p_1 and p_2, this translates to part (2) of the theorem statement. □

Theorem 6. *Three sites in P form a Delaunay triangle if and only if no other site of P lies in the closed circumcircle defined by the three sites.*

Proof: Consider three sites $p_1, p_2, p_3 \in P$. By Theorem 5 (1) they are the only sites on the same face of the Delaunay graph iff the closed circle through p_1, p_2, p_3 contains no other site in P. If only these three sites are on the same face, then clearly all edges exist between them, and so they form a Delaunay triangle. □

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

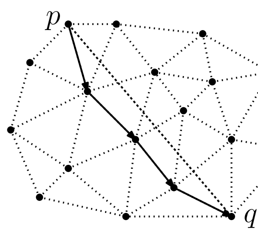

Figure 5. Spanner property of Delaunay triangulation.

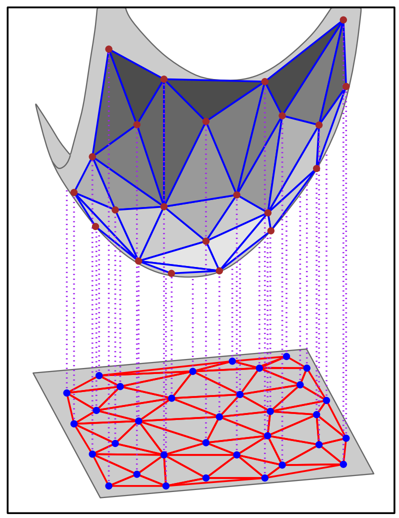

Figure 6. The Delaunay triangulation and Convex Hull.

2.2 Applications and Additional Properties

Spanner Properties. The length of the shortest path between two points through edges in the planar Delaunay triangulation is not too much longer than the Euclidean distance between these two points. In particular, the increase in distance by using only edges in the Delaunay triangulation is at most $4\pi\sqrt{3}/9 \approx 2.418$. This is proven in a paper by Keil and Gutwin.

Maximizing Angles. Among all triangulations, the Delaunay maximizes the minimum angle. This is useful because in many applications, we want to avoid skinny triangles, for better interpolation. In fact a stronger statement holds. Among all triangles with the smallest angle, the Delaunay triangulation maximizes the second smallest angle, and so on.

3 Delaunay to Convex Hull

Let Ψ be the paraboloid $z = x^2 + y^2$. For any point $p = (p_x, p_y) \in \mathbb{R}^2$, define the *vertical projection* (or *lifted image*) of p onto Ψ to be point $p^{\uparrow} = (p_x, p_y, p_x^2 + p_y^2)$ in $\mathbb{R}^3$

Given a set of points in the plane, P, let $P^{\uparrow}$ be the projection of every point in P onto Ψ. Let the *lower convex hull* of $P^{\uparrow}$ be the part of the convex hull visible to an observer at $z = -\infty$.

The following theorem is illustrated in Figure 6.

Theorem 7. *Let P be any set of points in the xy-plane. Then the projection of the lower convex hull of $P^{\uparrow}$ back onto the plane is the Delaunay triangulation of P.*

Proof: We must show the Delaunay condition(from Theorem 6), and the lower convex hull condition are equivalent.
Delaunay condition. Three sites in P form a Delaunay triangle if and only if no other site of P lies within the circumcircle defined by the sites.
Lower Convex hull condition. Three points $p^{\uparrow}, q^{\uparrow}, r^{\uparrow} \in P^{\uparrow}$ form a face of the lower convex hull of $P^{\uparrow}$ if and only if no other point of $P^{\uparrow}$ lies below the plane passing through $p^{\uparrow}, q^{\uparrow}, r^{\uparrow}$.

To do so, first consider an arbitrary plane in $\mathbb{R}^3$ that is tangent to Ψ at some point (a, b) in the xy-plane. Note that for the equation $z = x^2 + y^2$, $\frac{\partial z}{\partial x} = 2x$ and $\frac{\partial z}{\partial y} = 2y$. At the tangent point $(a, b, a^2 + b^2)$, these evaluate to $2a$ and $2b$, so the plane passing through that point has the form

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

To solve for γ, we use the fact that the plane passes through $(a, b, a^2 + b^2)$ to get that

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

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

$$z = 2ax + 2by - (a^2 + b^2) \tag{1}$$

Now if we shift the plane upwards by some positive amount h^2, we get the plane

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

The intersection of this with Ψ is

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

which after rearrangement is:

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

Thus, the intersection of an arbitrary lower halfspace with Ψ, when projected onto the x, y-plane is the interior of a circle!

Now if we project points p,q and r from the xy-plane onto Ψ, the points $p^{\uparrow}, q^{\uparrow}, r^{\uparrow}$ define a plane. Since $p^{\uparrow}, q^{\uparrow}, r^{\uparrow}$ are in the intersection of this plane and Ψ, the original points p, q and r lie on the circumference of the unique circle passing through p, q and r. Thus, any separate point $s \in P$ lies within this circle iff its projection $s^{\uparrow}$ onto Ψ lies in the lower halfspace of the plane passing through $p^{\uparrow}, q^{\uparrow}, r^{\uparrow}$. (See Figure 7.) □

4 Voronoi to Convex Hull

What about Voronoi Diagrams? From the above, we can use convex hull algorithms to solve Delaunay, and then use graph duality to go from Delaunay to Voronoi. But there's a direct connection between convex hull and Voronoi, and it's worth learning since it can help us with Voronoi variants.

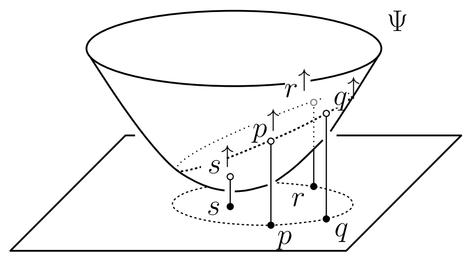

Figure 7. Planes and Circles.

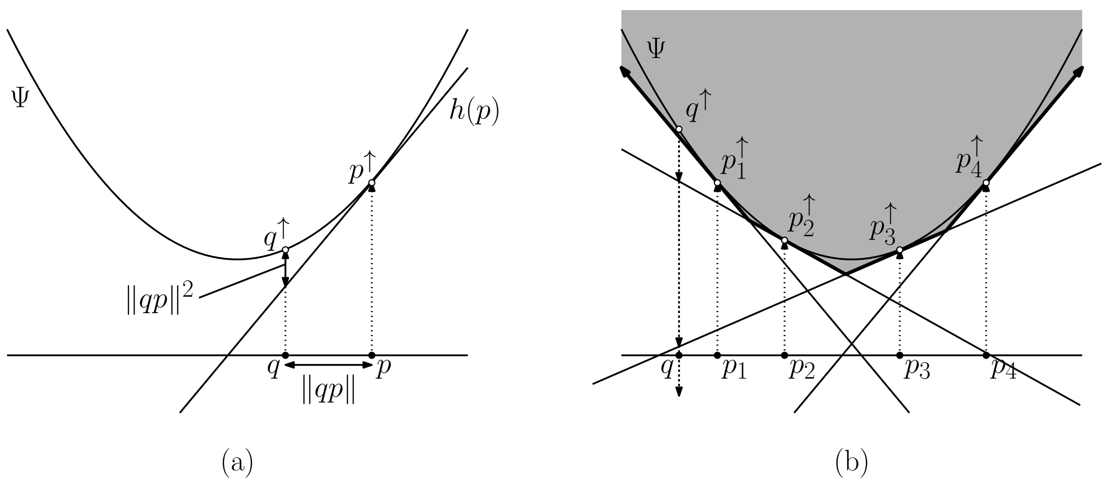

Figure 8. (a) Vertical distance from $q^\uparrow$ to $h(p)$ equals $||qp||^2 (i.e. |q-p|)^2$; (b) Thus, a vertical ray from $q^\uparrow$ intersects the planes in $H(P)$ in the same order as distances from q to points in P in the xy-plane.

Lemma 1. *Consider any two points p and q in the xy-plane, and let $h(p)$ be the tangent plane to Ψ passing through $p^\uparrow$. Then the vertical distance between $q^\uparrow$ and $h(p)$ is the squared distance from q to p.*

Proof: For any point $p = (a, b)$ in the plane, recall from Equation 1, that the tangent plane to Ψ passing through $p^\uparrow$ is

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

Let $h(p)$ be this plane. Now consider an arbitrary point $q = (q_x, q_y)$ in the xy-plane. Then $q^\uparrow = (q_x, q_y, q_z = q_x^2 + q_y^2)$. Thus the vertical distance from the point $q^\uparrow$ to $h(p)$ is

$$\begin{aligned} q_z - (2aq_x + 2bq_y - (a^2 + b^2)) &= (q_x^2 + q_y^2) - (2aq_x + 2bq_y - (a^2 + b^2)) \\ &= (q_x - a)^2 + (q_y - b)^2 \\ &= |q - p|^2 \end{aligned}$$

(See Figure 8(a)). □

Lemma 2. *Given a set of points $P = \{p_1, \ldots p_n\}$ in the xy-plane, let $H(P) = \{h(p) : p \in P\}$. Then for any point q in the xy-plane, a vertical ray directed downwards from $q^\uparrow$ intersects the planes of $H(p)$ in the same order as the distances of q from the points in P.*

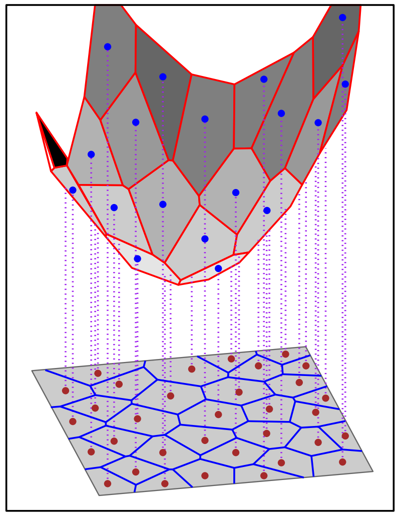

Figure 9. The Upper Envelope of hyperplanes tangent to $P^{\uparrow}$ equals the Voronoi diagram when projected onto the xy-plane

Proof: This follows immediately from Lemma 1. (Figure 8(b) illustrates the lemma.) □

Theorem 8. *Given a set P of points in the xy-plane, let $U(P)$ be the upper envelope of the tangent hyperplanes passing through each point $p^{\uparrow}$ for $p \in P$. Then the Voronoi diagram of P is equal to the vertical projection on the xy-plane of the boundary complex of $U(P)$*

Proof: Consider the upper envelope $U(P)$ of $H(P)$. This is an unbounded convex polytope, whose vertical projection covers the entire xy-plane. Now label every face of this polytope with the point $p \in P$ whose plane $h(p)$ defines the face. Then, by Lemma 2, the site p is closest to every point in the vertical projection of this face onto the plane. Thus, when $U(P)$ is projected on the xy-plane, it exactly gives the Voronoi diagram of P. (Figure 9). □

4.1 Higher-Order Voronoi Diagrams and Arrangements

An *order-k* Voronoi diagram is a subdivision of the plane into regions, each associated wiht a subset of k sites that are the k nearest neighbors of any point in the region. For example, when $k = 2$, each cell of an order 2 Voronoi diagram is associated with two sites (p_i, p_j), and the cell contains all points whose 2 closest sites are p_i and p_j. We next show that all order k Voronoi diagrams can be generated via projections onto Ψ.

Lemma 3. *Let P be a set of points in the xy-plane, and let $H(P) = \{h(p) : p \in P\}$ be the set of hyperplanes defined above. Let A be an arrangement of these hyperplanes and let $\mathcal{L}_k(A)$ be the k-th level of the arrangement of A. Then the vertical projection of $\mathcal{L}_k(A)$ onto the xy-plane equals the order k Voronoi diagram of P.*

Proof: Recall that in 2 dimensions, for any arrangement, A and any $k \in [1, n]$, the k-th level of the arrangement, $\mathcal{L}_k(A)$, consists of the line segments that have exactly k lines lying on or above them. Similarly, in 3 dimensions, for any arrangement A, and any $k \in [1, n]$, the k-th level of the arrangement, $\mathcal{L}_k(A)$, consists of the faces of the arrangement that have exactly k hyperplanes lying on or above them. By Lemma 2, the level k of the arrangement of $H(P)$, when projected vertically onto the xy-plane is exactly the order-k Voronoi diagram. (See Figure 10.) □

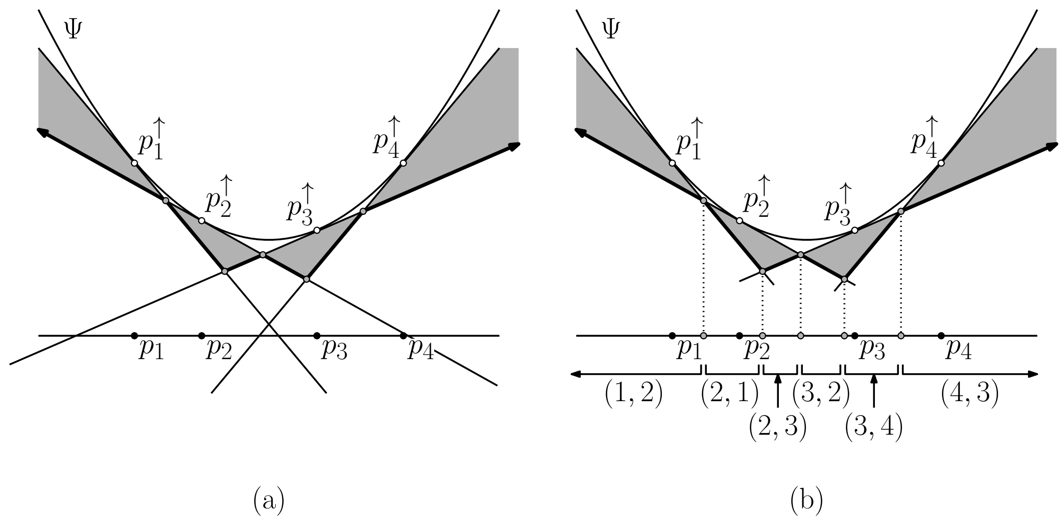

Figure 10. (a) Planes in $H(P)$ and the 2-level of their arrangement in bold; (b) The projection of the 2-level of this arrangement back onto the xy-plane. Note that this gives the order-2 Voronoi diagram of P.

Note that as shown in Figure 10, the projection actually gives a refinement of the order-2 Voronoi diagram because it distinguishes between e.g. the cells $(1, 2)$ and $(2, 1)$, depending on which of the two sites is closer.

Also note that the lower envelope of $H(P)$ is the order-n Voronoi diagram. This is called the farthest-point Voronoi diagram since each cell is defined by the farthest site from that cell.

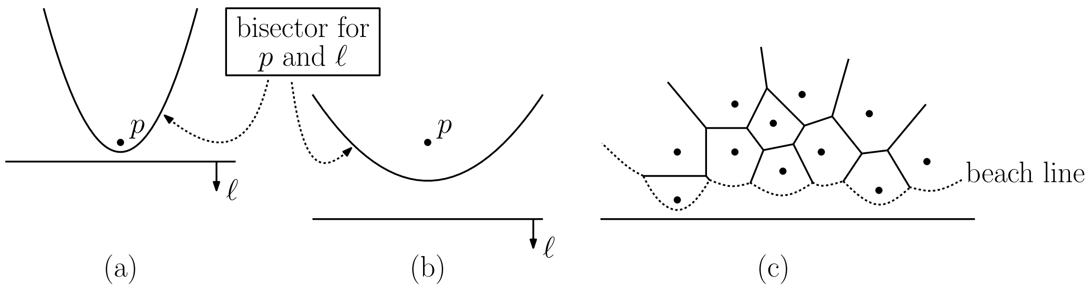

Figure 11. (a) and (b) parabolas defining the points closer to p than ℓ; (c) the "beach line" is the lower envelope of the parabolas

5 Fortune's Algorithm for Voronoi Algorithm

The above projection onto a paraboloid approach gives us one way to compute Voronoi diagrams (projection plus half-plane intersection in 3d). Another interesting and direct way is to use Fortune's algorithm (much of the discussion and figures in this section are from [2]). Fortune's algorithm is another sweep line method. One benefit of Fortune's algorithm is that it is generalizable in ways orthogonal to convex hull algorithms. In particular, it can more directly handle weighted Voronoi diagrams in 2D. It is a $O(n \log n)$ algorithm, which is optimal by a reduction from sorting. Clearly, by Duality, Fortune's algorithm can also be used to efficiently compute the Delaunay triangulation.

Basic idea:

- Horizontal line sweeps the sites from top to bottom
- Maintain portion of diagram due to sites above the sweep line
- Which points are closer to the sites above the sweep line than to the line itself? Set of parabolic arcs form a beach line that bounds the location of all such points

Also:

1. Points where two parabolas intersect are called "break points"
2. The break points trace out Voronoi edges

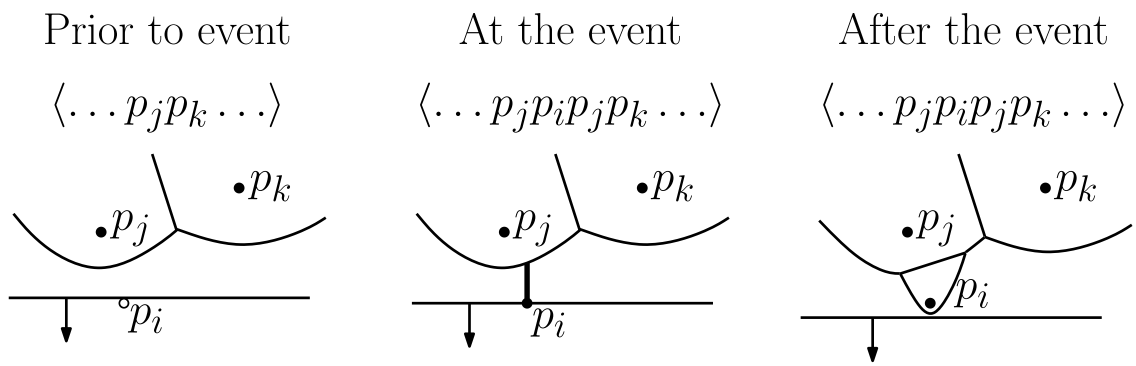

Figure 12. Site Event for site p_i. At the top is the sorted list of parabolas.

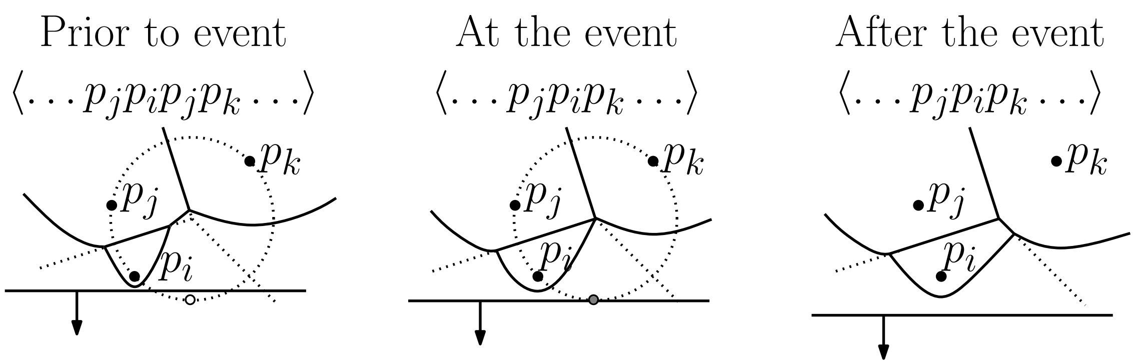

Figure 13. Voronoi vertex event

5.1 The Two Events

We maintain a sorted list (along x-axis) of parabolas that compose the lower envelope of the beach. Then there are two types of events that can occur as the sweep-line moves downward.

- **Site event.** The sweep-line intersects a new site in P (Figure 12). We do binary search to determine where to insert this new point in the beach line. The point initially forms a vertical line on the beach. This is the only way new sites (and their parabolas) can be inserted into the beach line, and it increases the size by at most 2. Thus, our beach line list (top of figure) is $O(n)$.

- **Voronoi vertex event.** The sweep-line dips below the circumcircle of three adjacent sites, p_i, p_j, p_k on the beach line(Figure 13). At this event, a Voronoi vertex is created in the diagram, and the consecutive triple $p_i p_j p_k$, is replaced with $p_i p_k$ in the beach line list.

5.2 Sweeping

There are three key data structures we use in the algorithm.

Partial Voronoi Diagram. We can store this as usual as a planar graph with doubly-linked lists of edges around the faces. Since the graph is partial, we connect all currently incomplete edges in the Voronoi diagram to a special vertex at infinity.

Beach Line. This is the sorted list of sites that form the arcs in the beach line. We don't need to explicitly store the parabolas. The key search operation is placing a newly intersected site's arc into the beach line.

How to do this? Between any pair of sites p_i and p_j there is break-point on the sweep-line, which is the center of a circle on the sweep line that has p_i and p_j on its borders. We can dynamically find this break-point in constant time, determine if the new site falls to the right or left of it, and then proceed with our binary search to find where the new site falls in the beach line in $O(\log n)$ time.

Event Queue. The event queue is a priority queue where we can insert and delete events, with keys equal to their y-coordinates. Initially, all sites are inserted into the priority queue. Additionally, for every consecutive triple $p_i p_j p_k$ currently on the beach line, if the bottom endpoint of their circumcircle is below the current sweep-line, we store a Voronoi vertex event in the priority queue with key equal to the smallest y value of this circle. We remove this event if sites $p_i p_j p_k$ ever cease to be consecutive.

Analysis. Each event takes $O(\log n)$ time to pull it out of the queue and process it. The size of all data structures are $O(n)$ and there are $O(n)$ total events processed. So total time is $O(n \log n)$

References

[1] David Austin. Voronoi diagrams and a day at the beach. *http://www.ams.org/samplings/feature-column/fcarc-voronoi*, 2006.

[2] Allen Miu. Lecture 7: Voronoi Diagrams, 2001. http://nms.csail.mit.edu/~aklmiu/6.838/L7.pdf.

[3] David Mount. Computational Geometry. http://www.cs.umd.edu/class/fall2016/cmsc754/Lects/cmsc754-fall16-lects.pdf, 2016.