**CS368: Geometric Algorithms** Handout # 7
**Design and Analysis**
Stanford University Tuesday, 13 May 1999

Solutions for homework #2: Voronoi and Delaunay diagrams

- **The Common Theory Problems**

### Problem 1. [5 points]

The edges of both the Delaunay and Voronoi diagrams are line segments. Give a simple necessary and sufficient condition on a pair of sites $A$ and $B$ so that $AB$ is a Delaunay edge *and* $AB$ intersects its dual Voronoi edge.

### Solution

We claim that $AB$ is a Delaunay edge that intersects its dual Voronoi edge iff the circle with $AB$ as diameter is site-free. The proof is as follows.

Note that a point $P$ is on the Voronoi edge separating $A$ and $B$ iff the unique circle through $A$ and $B$ centered at $P$ is site-free. Now clearly $AB$ will intersect its dual iff there is a point $P$ on $AB$ that is also on its dual, ie. there is a point $P$ on $AB$ such that the circle through $A$ and $B$ with center $P$ is site-free. The only such circle (through $A$ and $B$ with a center on $AB$) is the circle with $AB$ as the diameter, which proves the result.

### Problem 2. [10 points]

We mentioned in class that a triangulation of a set $S$ on $n$ sites in the plane is a Delaunay triangulation if and only if every edge passes the *InCircle* test with respect to its two adjacent triangles. This gives a linear-time algorithm to verify that a triangulation is Delaunay, and it also suggests the following algorithm to fix it up (if it's not). Start with any triangulation of the $n$ sites. If an edge fails the *InCircle* test, then swap it with the other diagonal of the quadrilateral formed by the two adjacent triangles (this edge must pass the test). Make this idea into a rigorous algorithm and prove its correctness. Prove that your algorithm always terminates in $O(n^2)$ steps. (*Open Problem:* Can this method be parallelized in an interesting way? What processor/time bounds can you get?)

### Solution

A triangulation is as defined in class: every face (except possibly the unbounded face) is a triangle, so there are always $2(n-1)-h$ bounded triangles, where $h$ is the size of the convex hull.

Given an initial triangulation $\mathcal{T}_0$, all of its edges not on the convex hull are suspect. Mark the suspect edges and keep them in a queue. At each step, we take an edge from the queue and check it using the *InCircle* test. If the edge needs to flip, we flip it and

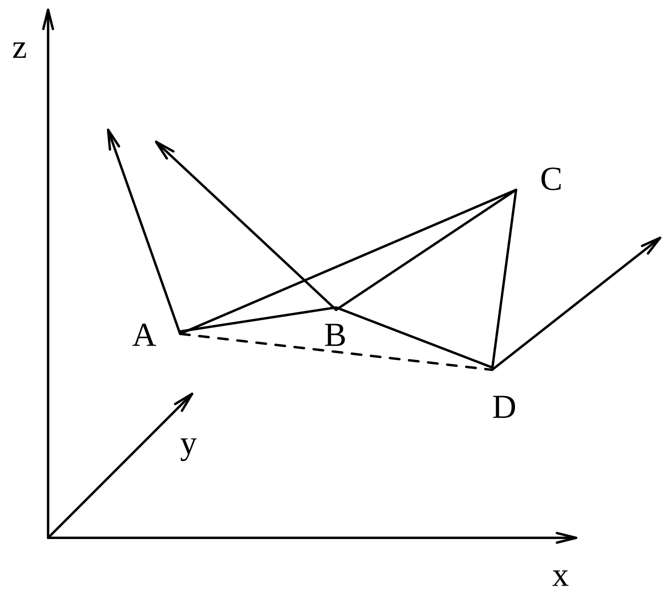


Figure 1: An edge flip under the lifting map.

add any previously unsuspect neighboring edges to the queue (up to four). Run this algorithm until the queue is empty. Note this algorithm uses space $O(n)$, since the queue stores at most one copy of each edge in the current triangulation. It is also clear that we always have triangulations during the whole process.

If this algorithm terminates then we have the Delaunay triangulation, since all edges pass the *InCircle* test. Note that we don't need to worry about flipping the diagonal of a non-convex quadrilateral, since these quadrilaterals always pass the *InCircle* test.

All that remains is to show the $O(n^2)$ time bound (which implies termination). Note that the total number of edges ever added to the queue is at most $O(n)$ for the initial suspects plus four for each edge flip. Since each step of the algorithm removes an edge from the queue, it is enough to bound the number of edge flips. We give two arguments that $O(n^2)$ edge flips are used: the first shows that no edge is ever reappears once flipped, and the other uses a potential function method.

**First Method.** To see that no edge recurs, consider the sites projected up on our favorite paraboloid $(x, y) \mapsto (x, y, x^2 + y^2)$. A given triangulation defines a (nonconvex) surface of triangular facets on these sites. Consider an edge $BC$ between faces $ABC$ and $BCD$ which fails the *InCircle* test (see figure 1). Since $D$ lies in the circumcircle of $ABC$, it lies below the plane through $ABC$ under the lifting map, and thus there is a concave angle between the faces $ABC$ and $BCD$. When we flip the edge $BC$, we add volume $ABCD$ to the region enclosed from below by the lifted triangulation, and the edge $BC$ becomes hidden in the interior. Since each edge flip only adds to the enclosed region, the edge $BC$ can never appear on the boundary again.

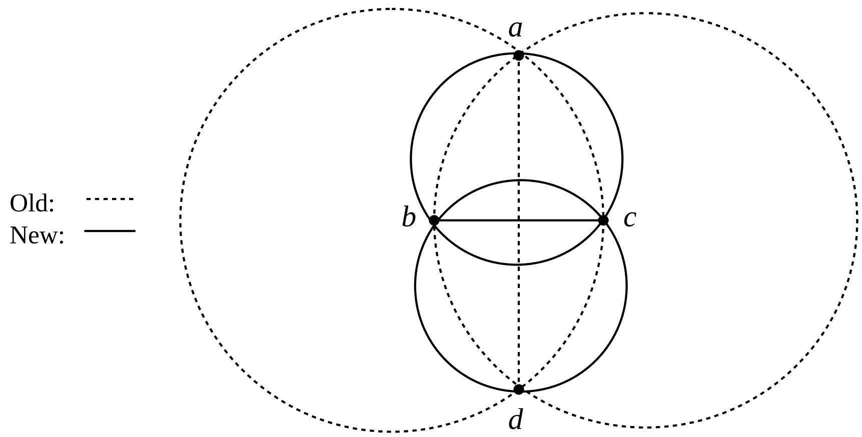


Figure 2: The function $\phi$ decreases with each flip.

**Second Method.** Now we present the potential function argument. Fix the set of sites. Define a potential function $\phi$ on triangulations $\mathcal{T}$:

$$\phi(\mathcal{T}) = \sum_{\triangle \in \mathcal{T}} w(\triangle)$$

where $w(\triangle)$ is the weight of the triangle $\triangle$, i.e. the number of sites in the interior of its circumcircle.

Any one triangle can have weight at most $n - 3$; since there are at most $2(n - 2)$ triangles in a triangulation, we have $\phi(\mathcal{T}_0) \leq 2(n-2)(n-3) = O(n^2)$. On the other hand $\phi \geq 0$, where the only triangulation that achieves this is the Delaunay triangulation. Now it suffices to prove the following.

**Lemma 1.** *The function $\phi$ decreases by at least two on each edge flip.*

**Proof:** Suppose we have a suspect edge $ad$ between triangles $bad$ and $cad$, and it fails the *InCircle* test, so that we flip it to create edge $bc$ (see figure 2). The flip destroys the two old triangles and creates the two new triangles $abc$, $bcd$. Then it suffices to prove

$$w(abc) + w(bcd) \leq w(bad) + w(cad) - 2.$$

This will follow from the following three claims:

(a) every site in a new circle is in an old circle;

(b) every site in both new circles is in both the old circles;

(c) sites $b$ and $c$ each appear in one of the old circles but in neither of the new circles (this accounts for the $-2$ term).

Of these, (c) follows immediately from the definition of the *InCircle* test.

To prove (a), it suffices to show the interior of the new circle $abc$ is completely covered by the old circles $bad$ and $cad$ (an analogous argument will then apply to $bcd$). We know that $d$ is outside of circle $abc$, and furthermore for a swap to be possible we must have $abdc$ a convex quadrilateral, i.e. $d$ lies somewhere between the rays $\vec{ab}$ and $\vec{ac}$. Since two distinct circles can intersect in at most twice, circle $bad$ intersects circle $abc$ at only $a$ and $b$, and $bad$ covers all of $abc$ except for a lune (a "grin" shape) between $a$ and $b$, and furthermore this lune is concave *away* from the center of $abc$ (this follows from the convexity of $abdc$). Similarly $acd$ covers all of $abc$ except a lune between $a$ and $c$. Since these lunes are concave away from the center of $abc$, they do not overlap (except at site $a$), so the interior of $abc$ is covered by the old circles.

To prove (c), we need to show that the intersection of the new circles $abc$ and $bcd$ (a convex "lemon" lune between $b$ and $c$) is contained in both the old circles. This would follow from showing that it is disjoint from the lunes considered in the previous paragraph; the argument is similar and omitted. □

**Problem 3. [10 points]**

Second-order Voronoi diagram:
Given $n$ sites in the plane, suppose we partition the plane according the nearest *and* the second-nearest site. Thus all points in the same region have the same nearest and second-nearest neighbor. This is sometimes called the *second-order* Voronoi diagram. Show that it is also of size $O(n)$ and can be computed in time $O(n \log n)$.

**Solution**

Let VD1 refer to the first-order Voronoi diagram and let VD2 be the second-order diagram. Clearly VD2 is a subdivision of VD1, since we can form VD1 by merging all faces of VD2 with the same nearest point. Consider the VD1 region $R_x$ of a site $x$. To compute the subdivision of $R_x$, we need to know which site is second-nearest for each $q \in R_x$. An apparent method is to compute the Voronoi diagram of the remaining sites, restricting in $R_x$. Naively, this leads to $O(n \log n)$ time bound for each face. Fortunately, the following lemma says that we only need to consider those neighbors of $x$ in VD1.

**Lemma 2.** *If $x$ and $y$ are nearest and second-nearest sites for a point $q$, then $xy$ is a Delaunay edge (ie. the VD1 regions for $x$ and $y$ are adjacent).*

**Proof:** The circle through $y$ with center $q$ contains only the site $x$. By contracting this circle around $y$ until it hits $x$ (maintaining the tangent slope at $y$), we get a site-free circle through $x$ and $y$, indicating that $xy$ is Delaunay. □

By the lemma, we can first compute the Voronoi diagram of the sites, then compute the Voronoi diagram of the neighbors of $x$, and clip it by $R_x$ to obtain the subdivision of $R_x$ in VD2. The computation of VD1 can be done in $O(n \log n)$ time. For a site $x_i$, with degree $d_i$ in the Delaunay triangulation, the Voronoi diagram, $V_i$, of the neighbors of $x_i$ can be computed in time $O(d_i \log d_i)$. To clip $V_i$ inside $R_{x_i}$, we can compute the intersection between each edge of $V_i$ and the boundary of $R_{x_i}$, which can be done in $O(\log d_i)$ time as $R_{x_i}$ is a convex face with $d_i$ sides. Therefore, the latter two steps cost $\sum_i O(d_i \log d_i) = O(n \log n)$ time.

The space requirement is bounded by the sum of the complexities of all the Voronoi diagrams computed, which is $O(n) + O(\sum_i d_i) = O(n)$.

**Remark.** We can also prove that VD2 has linear size directly. Any vertex $v$ of VD2 must arise from one of two possibilities: either there are three sites which are all nearest to $v$, or there are three sites which are all second-nearest to $v$ (with a single nearest site).

In the first case the vertex corresponds to three sites with a site-free circumcircle (the usual Delaunay triangle), and in the second case the vertex corresponds to three sites whose circumcircle contains exactly one interior site.

Recall the discussion in class regarding the 'scope' of a triangle (namely, the number of sites included in its circumcircle). We proved that $|T_{\leq k}| = \Theta(n(k+1)^2)$ where $T_{\leq k}$ represents the number of triangles with scope at most $k$. In particular, $T_0 + T_1 = T_{\leq 1} = \Theta(n)$. Thus the second-order Voronoi Diagram has a linear number of vertices, and thus linear complexity overall (it is planar).

**Remark.** Using the arrangement of tangent planes to a paraboloid, we may think of VD2 as the second layer down in this arrangement. That is, if we look from above with X-ray vision strong enough to see through exactly one plane, we will see the second order Voronoi diagram.

- **The Additional Theory Problems**

**Problem 4. [10 points]**

Davenport-Schinzel sequences of order 2 and triangulations:
Let $P$ be any convex polygon with $n$ vertices. A triangulation of $P$ is a collection of $n-3$ non-intersecting chords connecting pairs of vertices of $P$ and partitioning $P$ into $n-2$ triangles. Set up a correspondence between such triangulations and $DS(n-1,2)$ sequences, as follows. Number the vertices $1, 2, \ldots, n$ in their order along $\partial P$. Let $T$ be a given triangulation. Include in $T$ the edges of $P$ too. For each vertex $i$, let $T(i)$ be the sequence of vertices $j < i$ connected to $i$ in $T$ and arranged in *decreasing* order, and let $U_T$ be the concatenation of $T(2), T(3), \ldots, T(n)$.

(a) Show that $U_T$ is a $DS(n-1,2)$ sequence of maximum length.

(b) Show that any $DS(n-1,2)$-sequence of maximum length can be realized in this manner, perhaps with an appropriate renumbering of its symbols.

(c) Use (a) and (b) to show that the number of different $DS(n, 2)$ sequences of maximum length is $\frac{1}{n}\binom{2n-2}{n-1}$ (where two sequences are different if one cannot obtain one sequence from the other by renumbering its symbols).

**Solution**

(a) First, we show that $U_T$ is a $DS(n-1, 2)$ sequence. (1) Every two adjacent symbols in $U_T$ are distinct. Because the polygon edge $(i-1, i)$ is always in the triangulation and all the symbols in $T(i)$ are in decreasing order, the first symbol in $T(i)$ is $i-1$, which is greater than all of those in $T(i-1)$. Together with the fact that each symbol appears at most once within each $T(i)$, we have that any two consecutive symbols in $U_T$ are distinct. (2) There is no subsequence with the pattern of $a, b, a, b$. Otherwise, suppose there is a subsequence $i, j, i, j$, which is from the blocks $T(i_1), T(j_1), T(i_2), T(j_2)$ where $i_1 \leq j_1 \leq i_2 \leq j_2$. There are two cases: (i) $i < j$. Because all the symbols within one block are in decreasing order, we have $i_2 < j_2$. In addition, $j < j_1 \leq i_2$. Thus, both the edges $(i, i_2)$ and $(j, j_2)$ are in the triangulation where $i < j < i_2 < j_2$. This is impossible because those two edges cross. (ii) $i > j$. By a similar argument, we would have that the edges $(j, j_1)$ and $(i, i_2)$, where $j < i < j_1 < i_2$, are in the triangulation. This is also impossible. In addition, it is clear that there are at most $n-1$ symbols in $U_T$. Therefore $U_T$ is a $DS(n-1, 2)$ sequence. Notice also that the length of $U_T$ is $2n-3$ because each edge in the triangulation corresponds to exactly one occurrence of a symbol. Therefore, $U_T$ is a $DS(n-1, 2)$ sequence of maximum length.

(b) For a $DS(n-1, 2)$ sequence $S$ of maximum length, we break it into blocks right before wherever a symbol appears at the first time. Then, $S$ is broken into $n-1$ blocks since each symbol occurs in $S$ at least once. Let the blocks be $T(2), \ldots, T(n)$. Renumber the first symbol of $T(i)$ to $i-1$. Now, create an edge set $E$ as $E = \{(i, j) |\ i \text{ appears in the block } T(j)\}$.

To prove that $E$ forms a triangulation of a convex $n$-gon, we verify the following facts. (1) All the polygon edges are in $E$. Obviously, the edges $(i, i+1)$, for all $1 \leq i \leq n-1$, are in $E$ by the construction. The edge $(1, n)$ is in $E$ as well because if $S$ is of maximum length, then the last symbol of $S$ is the same as the first one. Otherwise we could add the first symbol to the end of $S$, getting a longer $DS$-sequence. Since the first symbol in $S$ is renumbered to 1, we have that 1 is in the block $T(n)$. That is, the edge $(1, n)$ is in $E$. (2) If $i < j < k < l$, then either $(i, k)$ or $(j, l)$ cannot be in $E$. If this were not true, then $i$ would appear in $T(k)$ and $j$ in $T(l)$. Because $i$ is also in $T(i+1)$, and $j$ in $T(j+1)$, we would have had a subsequence of $i, j, i, j$, from the blocks $T(i+1), T(j+1), T(k), T(l)$ (this statement remains valid when $k = j+1$ because $j$ is the first symbol in $T(j+1)$), contradicting the fact that $S$ is a $DS(n-1, 2)$ sequence. Therefore, there is no proper crossing among the edges in $E$. In addition, there are $2n-3$ edges because the length of $S$ is $2n-3$. Thus, all the edges in $E$ form a triangulation of an $n$-vertices convex polygon.

Now, we proceed to prove that the symbols in $T(i)$ are in decreasing order after renumbering. Otherwise, suppose $i < j$ and $i$ appears before $j$ in $T(k)$. Then, we would have a subsequence of $i, j, i, j$, from the blocks $T(i+1), T(j+1), T(k)$, a contradiction.

Therefore, after appropriate renumbering, any $DS(n-1, 2)$ sequence of maximum length can be realized in the described manner.

(c) By (a) and (b), the number of different $DS(n, 2)$ sequences of maximum length is the same as the number of different ways to triangulate a simple polygon with $n+1$ vertices. By Handout #15 in the course reader, it is also the number of different binary trees with $n-1$ nodes, which is $\frac{1}{n}\binom{2n-2}{n-1}$ (refer to Section 2.3.4.4 in Knuth, volume I). □

**Problem 5. [10 points]**

We have discussed in class the lifting map $\lambda(x, y) : (x, y) \mapsto (x, y, x^2 + y^2)$ from points in the $xy$-plane to points on the paraboloid of revolution $z = x^2 + y^2$. As we will also see, the downwards-looking faces of the convex hull of the lifted images of a collection of sites in the $xy$-plane correspond to the Delaunay diagram of the sites. What do the upward-looking faces correspond to? Is there an analogous Voronoi diagram? How fast can it be computed?

**Solution**

For the lifting $p$ of a site $A$, consider a plane in space above $p$. Then the intersection of this plane with the paraboloid, when projected down to the $xy$-plane, must contain $A$, according to the established connection between incircle-ness in the plane and above-ness in the lifting space. Therefore, if we take any upper triangular facet of the 3-D convex hull, then the plane through the facet is above all the other points, and projects down to a circle in the $xy$-plane through three of the sites and surrounding all the others.

Hence the upward-looking faces project down to a triangulation of the (two-dimensional) convex hull of the original sites, such that each triangle contains all the other sites in its circumcircle. Let's call this the Upper Delaunay Triangulation (UDT).

Since an $O(n \log n)$ time algorithm exists to compute the convex hull in three dimensions, we may compute the UDT in the same time (compute the entire convex hull, then throw away the downward-looking facets and project). Or alternatively, we can use the divide and conquer technique similar to computing Voronoi diagrams.

We claim the UDT is dual to the Farthest-Point Voronoi Diagram (FVD), which is defined by dividing the plane up into regions according to which site is the furthest. Similar to Voronoi diagrams, for a site $A$, its FVD region is the intersection of the halfspaces bounded by the bisectors between $A$ and every other site (the only difference is that we intersect those halfspaces which do not contain $A$). Therefore, each FVD region is a convex face. Two FVD regions of sites $A, B$ are adjacent if and only if there is a point in the plane that is equidistant to $A, B$ and closer to all other sites.

Now suppose that two sites $A$ and $B$ are adjacent in the UDT. It means that there is a circle through $A$ and $B$ centered at a point $p$ that contains all the other sites. This implies that $p$ is equidistant from $A$ and $B$ and closer or equidistant to all other sites. It then serves as a witness of that the FVD regions of $A$ and $B$ are adjacent. Conversely, if the FVD regions of $A, B$ are adjacent, there must be a point $p$ equidistant to $A, B$ and closer to all other sites. We can then draw a circle centered at $p$ and passing through $A, B$. This circle should contain all the other points in the interior. Suppose that $C$ is the first site the circle encounters by pushing the circle along one direction. Then the triangle $ABC$ contains all the other sites and therefore a triangle in the UDT.

**Remark.** Unlike in the Voronoi diagram case, not every site has non-empty FVD region — only those points on the convex hull can have non-empty FVD regions.

**Problem 6. [15 points]**

We want to do an analysis of the randomized incremental algorithm for Delaunay triangulations discussed in class, but based on the appearance and disappearance of Delaunay edges during the process, rather than that of triangles.

We defined the *weight* of a triangle $\Delta$ to be the number of sites inside the circumcircle of $\Delta$ and related in the class analysis this weight with the probability that $\Delta$ would ever appear as a Delaunay triangle during the random process. How should we define the corresponding notion of the weight of an edge $e = AB$? Below we define one possibility, but feel free to explore your own definition, as long as it leads to the same eventual result.

One way is to let the weight of an edge $AB$ be $w$ if there is a circle through $A$ and $B$ which contains exactly $w$ other sites to the left of $AB$ and $w$ sites to the right. Show that this notion of weight is well defined. Prove an edge of weight $w$ arises as Delaunay at some point of the random process with probability at most $4/(w+1)(w+2)$.

By emulating the argument given in class for triangles, show the combinatorial-geometric result that in any collection of $n$ sites, the number of edges of weight at most $w$ is $O(n(k+1))$.

Briefly outline how combining these results (using the summation-by-parts trick shown in class) allows us to conclude that the expected number of edges that arise during the random process is $O(n)$ (of course, this also follows immediately from the result for triangles ...).

Given an edge $AB$ in a group of $n$ sites, how fast can you calculate its weight? How fast can you calculate the *minimum* number of other sites whose deletion would make $AB$ a Delaunay edge? Are those two quantities related?

**Solution**

Since the number of triangles that ever appear during the incremental construction is $O(n)$, it follows that the total number of edges that ever appear as Delaunay edges is also linear. Following the approach provided in the problem, we can have a refined analysis of

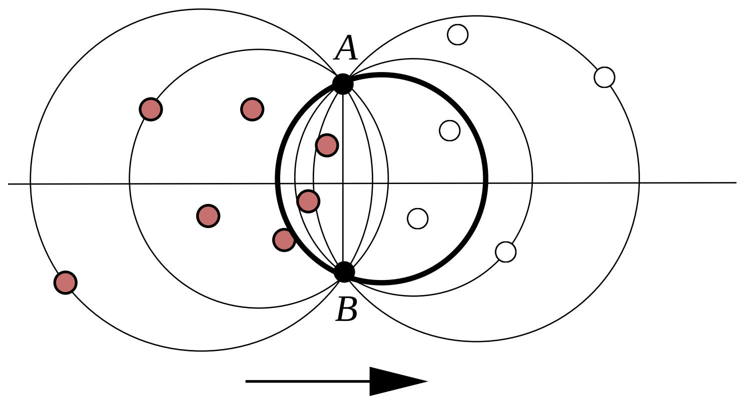


Figure 3: In the above picture, the red points are drawn in solid and the blue ones are hollow. It can be seen that $r(-\infty) = 7$, $b(+\infty) = 5$, and $w = 2$ (the bold circle contains the same number of red and blue points).

the edge probabilities. This refined analysis is interesting in its own right and allows us to obtain bounds in certain situations for which a triangle-based analysis is not sufficient.

For an edge $AB$, color all the other points by red or blue according to whether they are to the left or to the right of the line $AB$ — assign a point red color if it is to the left or blue otherwise. For a circle $O$ passing through $A, B$, let $r(O), b(O)$ denote the number of red and blue points contained in the interior of $O$, respectively.

Now, consider the family of circles that pass through $A, B$ and imagine that the center moves on the bisector between $A, B$, say the $x$-axis, from $-\infty$ to $+\infty$. Define a function $r(t), b(t)$ whose values are $r(O), b(O)$, respectively, where $O$ is the circle centered at point $(t, 0)$ and passing through $A, B$. When $t$ goes from $-\infty$ to $+\infty$, $r(t)$ decreases monotonically and $b(t)$ increases monotonically — when the circle passes a red point, $r(.)$ increases by one; when it passes a blue point, $b(.)$ decreases by one (see figure 3). In other words, the function $f(t) = r(t) - b(t)$ decreases monotonically (by one at each breaking point), and $f(-\infty) \geq 0, f(+\infty) \leq 0$. Therefore, there must exist $t_0$ so that $f(t_0) = 0$, i.e. $r(t_0) = b(t_0)$. Furthermore, for any $t < t_0$, since $r(t) \geq r(t_0), b(t) \leq b(t_0)$, for $r(t) = b(t)$, it must be the case that $r(t) = r(t_0)$ and $b(t) = b(t_0)$. The same analysis applies to $t > t_0$. Thus, the weight is uniquely defined for each edge.

Notice that $0 \leq w \leq \lfloor n/2 \rfloor$, and that edges of zero weight are true Delaunay edges. The concept of edge weight turns out to be the key tool in our edge-based analysis, because it provides a canonical classification of edges in which most edges have a fairly large weight, and it is related to the probability that an edge can ever appear as a Delaunay edge.

**Lemma 3.** *The probability that edge $AB$ appears as a Delaunay edge during a random incremental procedure is at most $4/(w + 1)(w + 2)$, when $w$ is the weight of $AB$.*

**Proof.** Suppose the set of red and blue points realizing the weight of $AB$ are $S_r, S_b$, where $\mid S_r \mid = \mid S_b \mid = w$. Observe that if there are both points in $S_r$ and points in $S_b$

presenting, $AB$ cannot be a Delaunay edge. Thus, $A, B$ must before all the points in $S_r$ or all the points in $S_b$ in the random sequence in the insertion. The probability for this to happen is at most $2\frac{2!w!}{(w+2)!} = \frac{4}{(w+1)(w+2)}$. □

.

Let $E_w$ denote the set of edges $AB$ whose weight is $w$, and let $E_{\leq k} = \cup_{w=0}^{k} E_w$. Thus, $E_0 = E_{\leq 0}$ is the set of Delaunay edges. When $w \geq \lfloor n/2 \rfloor$, the set $E_w$ is empty and $E_{\leq w}$ contains all $\binom{n}{2}$ edges.

**Lemma 4.** *The number $\|E_{\leq k}\|$ of edges $AB$ whose weight is at most $k$ is $O(n(k+1))$.*

**Proof.** The proof is analogous to that of the corresponding lemma for triangles. We draw a random sample $\mathcal{R}$ by choosing $r$ of the given points, where $r$ is a parameter to be chosen later. We now let $D(\mathcal{R})$ denote the set of Delaunay edges of the sample $\mathcal{R}$. The expected number of such edges is

$$\mathbf{E}\left[\|D(\mathcal{R})\|\right] \geq \sum_{w=0}^{k} \sum_{AB \in E_w} \mathbf{Prob}\left[AB \in D(\mathcal{R})\right], \tag{1}$$

arguing as before. We also know that the total number of Delaunay edges in the sample is always at most $3r$.

Suppose $AB$ is an edge of $E_w$. If $w$ is fairly small, the edge $AB$ is pretty likely to be Delaunay, provided that points $A$ and $B$ are in the sample and the sample isn't too large, because the absence of only a few points can guarantee that $AB$ will be Delaunay. Indeed, the definition of weight tells us that there is a set $\mathcal{P}_{AB}$ of $2w$ points, $w$ to the left of $AB$ and $w$ to the right, for which some circle through $A$ and $B$ encloses precisely these $2w$ points. Thus we have the lower bound

$$\mathbf{Prob}\left[AB \in D(\mathcal{R})\right] \geq \mathbf{Prob}\left[\{A, B\} \subseteq \mathcal{R} \subseteq \mathcal{P} \setminus \mathcal{P}_{AB}\right] = \binom{n-2w-2}{r-2} \Big/ \binom{n}{r}.$$

Putting these estimates together yields

$$\begin{aligned} 3r \geq \mathbf{E}\left[\|D(\mathcal{R})\|\right] &\geq \sum_{w=0}^{k} \binom{n-2w-2}{r-2} \|E_w\| \Big/ \binom{n}{r} \\ &\geq \sum_{w=0}^{k} \binom{n-2k-2}{r-2} \|E_w\| \Big/ \binom{n}{r} = \binom{n-2k-2}{r-2} \|E_{\leq k}\| \Big/ \binom{n}{r}. \end{aligned}$$

In other words, we have derived an upper bound on $\|E_{\leq k}\|$ that depends on the sample size $r$:

$$\|E_{\leq k}\| \leq \frac{3r\binom{n}{r}}{\binom{n-2k-2}{r-2}} = \frac{3n\binom{n-1}{r-1}}{\binom{n-2k-2}{r-2}} = 3n\, B(n-1, 2k+1, r-2, 1).$$

An analysis of $B(N, M, s, t)$ tells us that the best bound is obtained by setting $r = \lfloor n/(2k+1) \rfloor + 1$; for that choice we obtain $\|E_{\leq k}\| = O(n(k+1))$. The actual implied $O$-constant we get from this calculation is about $6e \approx 16$. □

An analog of the triangle theorem for the total number of edges can now easily be proven. The overall $O$-constant thus obtained is about 65. This constant could be significantly reduced by slight refinements of the analysis.

As the above discussion makes clear, the weight of an edge $e = AB$ can be easily obtained once we know the ordering of all the other sites $C$ according to where the center of the circumcircle of $ABC$ lies along the horizontal bisector of $AB$. Both this ordering and the weight can be easily computed in $O(n \log n)$ time. Actually, we can compute the weight in linear time as follows.

To compute the weight, consider the center points of circles $ABC$ for all the other sites $C$ in the set. Suppose their $x$-coordinates (reminder: the bisector of $AB$ is the $x$-axis) in increasing order are $t_1, t_2, \cdots, t_{n-2}$. If there are $l$ red points, we claim that $r(t) - b(t) = l - s$, for $t \in (t_{s-1}, t_s)(t_0 = -\infty)$, for any $1 \leq s \leq n-2$. This follows from the facts that we have observed earlier: 1) for all $t \in (t_i, t_{i+1})$, $r(t) - b(t)$ remains the same, 2) $r(t) - b(t)$ decreases by one when $(t, 0)$ is the center point of a circumcircle of $ABC$, and 3) $r(-\infty) - b(-\infty) = l$. If we take $s = l$, we have that $r(t) = b(t)$ for $t \in (t_{s-1}, t_s)$. A linear time algorithm therefore works as follows: compute all the $t$'s, pick the $l$-th element according to the $x$-coordinates ordering , and count the number of red/blue points inside the circle centered at $(t, 0)$ and passing through $A, B$. The linear running time of our algorithm is based on the existence of linear time selection algorithm.

For an edge $e$ with weight $w$, suppose that $S_r, S_b$ are the red and blue sites, respectively, in the circle that contains the same number of red and blue sites. It is clear that for $e$ to be a Delaunay edge, we have to at least delete either all the sites in $S_r$ or all the sites in $S_b$. This requires to delete $w$ sites. On the other hand, if we delete all the sites in $S_r \cup S_b$, $e$ becomes a Delaunay edge as the circle becomes empty. Thus, it is sufficient to delete $2w$ sites for $e$ to be a Delaunay edge. □