

CMSC 754: Lecture 13

Delaunay Triangulations: Incremental Construction

Reading: Chapter 9 in the 4M's. This algorithm was presented in L. J. Guibas, D. E. Knuth, and M. Sharir, *Randomized incremental construction of Delaunay and Voronoi diagrams*, *Algorithmica*, 7, 1992, 381–413.

Constructing the Delaunay Triangulation: We will present a simple randomized incremental algorithm for constructing the Delaunay triangulation of a set of n sites in the plane. Its expected running time is $O(n \log n)$ (which holds in the worst-case over all point sets, but in expectation over all random insertion orders). This simple algorithm had been known for many years as a practical solution, but it was dismissed by theoreticians as being inefficient because its worst case running time is $O(n^2)$. When the randomized analysis was discovered, the algorithm was viewed much more positively.

The algorithm is remarkably similar in spirit to the randomized algorithm for trapezoidal map algorithm in that it not only builds the triangulation, but it also provides a point-location data structure for the final triangulation as well. (We will not discuss the point-location data structure explicitly, but it follows from the history-DAG approach used for trapezoidal maps.)

The input consists of a set $P = \{p_1, \dots, p_n\}$ of point sites in \mathbb{R}^2 . As with any randomized incremental algorithm, the idea is to insert sites in random order, one at a time, and update the triangulation with each new addition. The issues involved with the analysis will be showing that, after each insertion, the expected number of structural changes in the diagram is $O(1)$.

As with the incremental algorithm for trapezoidal maps, we need some way of keeping track of where newly inserted sites are to be placed in the diagram. We will store each of the uninserted sites in a *bucket* according to the triangle in the current triangulation that contains it. We will show that the expected number of times that a site is rebucketed throughout the course of the algorithm is $O(\log n)$, which when summed over all the sites leads to a total time of $O(n \log n)$.

Incircle Test: Before presenting the algorithm, we need to introduce the geometric primitives involved in testing whether triangles satisfy the Delaunay condition. Recall that a triangle $\triangle abc$ is in the Delaunay triangulation, if and only if the circumcircle of this triangle contains no other site in its interior. (Recall that we make the general position assumption that no four sites are cocircular.)

How do we test whether a site d lies within the interior of the circumcircle of $\triangle abc$? First off, through the use of orientation tests, we can determine whether d lies in any of the blue or red shaded regions shown in Fig. 1(a). (We will leave this as an exercise.) If d lies in the blue regions labeled “out”, it is definitely not in the circumcircle of $\triangle abc$, and if it lies in the red region labeled “in”, it is definitely inside the circumcircle.

Otherwise, d lies in one of the unshaded regions labeled “?”. You might think that we need to derive the equation for a circle, but there is a remarkably elegant solution to this problem. Observe first that if d lies in any of these regions, the four points are in convex position.




Fig. 1: Incircle test.

Second, by relabeling a , b , and c , we may assume that the counterclockwise order of points is $\langle abcd \rangle$. Under this assumption, we claim that d lies in the circumcircle determined by the $\triangle abc$ if and only if the following determinant is positive (see Fig. 1(b)–(d)).

$$\text{inCircle}(a, b, c, d) \equiv \det \begin{pmatrix} a_x & a_y & a_x^2 + a_y^2 & 1 \\ b_x & b_y & b_x^2 + b_y^2 & 1 \\ c_x & c_y & c_x^2 + c_y^2 & 1 \\ d_x & d_y & d_x^2 + d_y^2 & 1 \end{pmatrix} > 0.$$

This is called the *incircle test*. It is notable that the incircle test in 2-D can be reduced to an orientation test in 3-D, where we have effectively lifted the point onto a paraboloid in 3-space by creating an addition z -coordinate whose value is $x^2 + y^2$.

Deriving the Incircle Test (Optional): We will not prove the correctness of this test, but we will show a somewhat simpler assertion, namely that if the four points are cocircular then the above determinant is equal to zero. (It follows from continuity that as d moves from inside the circle to the outside, the sign of the determinant changes as well.)

Suppose that a , b , c , and d are all cocircular then there exists a center point $q = (q_x, q_y)$ and a radius r such that

$$(a_x - q_x)^2 + (a_y - q_y)^2 = r^2,$$

and similarly for the other three points. (We won't compute q and r , but merely assume their existence for now.) Expanding this and collecting common terms we have

$$\begin{aligned} 0 &= (a_x^2 + a_y^2) - 2q_x a_x - 2q_y a_y + (q_x^2 + q_y^2 - r^2) \\ &= (-2q_x) a_x + (-2q_y) a_y + 1 \cdot (a_x^2 + a_y^2) + (q_x^2 + q_y^2 - r^2) \cdot 1. \end{aligned}$$

If we do the same for the other three points, b , c , and d , and express this in the form of a matrix, we have

$$\begin{pmatrix} a_x & a_y & a_x^2 + a_y^2 & 1 \\ b_x & b_y & b_x^2 + b_y^2 & 1 \\ c_x & c_y & c_x^2 + c_y^2 & 1 \\ d_x & d_y & d_x^2 + d_y^2 & 1 \end{pmatrix} \begin{pmatrix} -2q_x \\ -2q_y \\ 1 \\ q_x^2 + q_y^2 - r^2 \end{pmatrix} = 0.$$

In other words, there exists a linear combination of the columns of the 4×4 matrix that is equal to the zero vector. We know from linear algebra that this is true if and only if the determinant of the matrix is zero.

Incremental update: It will be convenient to assume that each newly added point lies within some triangle of the triangulation. This will not be true when points are added that lie outside the convex hull of the current point set. To satisfy this, we will start by adding three bogus *sentinel points* that will form an infinitely large triangle that contains all the points. After the final triangulation is completed, we will remove these sentinel points and their incident triangles. (In our trapezoidal map algorithm, this is analogous to putting all the segments in an enclosing rectangle.)¹ We won't show this triangle in our figures, but imagine that it is there nonetheless.

We permute the sites in random order and insert one by one. When a new site p is added, we find the triangle $\triangle abc$ of the current triangulation that contains this site (we will see how later), insert the site into this triangle, and join the site to the three surrounding vertices (see Fig. 2(a)). This creates three new triangles incident to p , $\triangle pab$, $\triangle pbc$, and $\triangle pca$. For each, we check the vertex of the triangle that lies on the *opposite side* of the edge that does not include p . (If there is no such vertex, because this edge is on the convex hull, then we are done.) If this vertex fails the incircle test, we swap the edge. This replaces one triangle that was incident to p with two new triangles. We repeat the same test process recursively with these triangles (see Fig. 2(b)).




Fig. 2: Delaunay point insertion.

The algorithm for the incremental algorithm is shown in the code block below, and an example is presented in Fig. 3. The current triangulation is kept in a global data structure. The edges should be thought of as pointers to entries in the DCEL representation.

As you can see, the algorithm is very simple. There are only two elements that have not been shown are the implementation. The first is the update operations on the data structure for the simplicial complex. These can be done in $O(1)$ time each on any reasonable representation (a DCEL, for example). The other issue is locating the triangle that contains p . We will discuss this below.

Local vs. Global Delaunay: There is one major issue in establishing the correctness of the algorithm. When we performed empty-circle tests, we applied the empty circle tests only for the newly created triangles containing the site p , and then only for sites that lay on the opposite side of an edge of each such triangle.

Why this works is related to an important issue in Delaunay triangulations. We know from the empty circumcircle condition that in a Delaunay triangulation, the circumcircle of every triangle is empty of other sites. This suggests two different criteria for testing whether a triangulation is Delaunay:

¹Some care must be taken in the construction of this enclosing triangle. It is not sufficient that it simply contains all the points. It should be so large that the vertices of the triangle do not lie in the circumcircles of any of the triangles of the final triangulation. Our book suggests a symbolic alternative, which is more reliable.

Randomized Incremental Delaunay Triangulation Algorithm

```

Insert(p) {
  Find the triangle  $\triangle abc$  containing  $p$ 
  Insert edges  $pa$ ,  $pb$ , and  $pc$  into triangulation
  SwapTest(ab) // check/fix the surrounding edges
  SwapTest(bc)
  SwapTest(ca)
}

SwapTest(ab) {
  if ( $ab$  is an edge on the exterior face) return
  Let  $d$  be the vertex to the right of edge  $ab$ 
  if (inCircle( $b, p, a, d$ )) { //  $d$  violates the incircle test
    Flip edge  $ab$  // replace  $ab$  with  $pd$ 
    SwapTest(ad) // check/fix the new suspect edges
    SwapTest(db)
  }
}
    
```




Fig. 3: Delaunay point insertion.

Global Delaunay: The circumcircle of each triangle $\triangle abc$ contains no other site d . (Fig. 4(a) shows a violation.)

Local Delaunay: For each pair of neighboring triangles $\triangle abc$ and $\triangle acd$, d lies outside the circumcircle of $\triangle abc$. (Fig. 4(b) shows a violation.)




Fig. 4: Global- and local-Delaunay conditions.

Clearly, if a triangulation is globally Delaunay it is locally Delaunay. Our incremental algorithm only checks the local-Delaunay condition, however. Could it be that a triangulation might satisfy the condition locally, but fail to satisfy it globally (see Fig. 4(c))? Delaunay proved, however, that the two conditions are in fact equivalent:

Delaunay's Theorem: A triangulation is globally Delaunay iff it is locally Delaunay.

Proof: (Sketch) The global to local implication is trivial, so it suffices to prove that local implies global. Consider any triangle $\triangle abc$ of a locally Delaunay triangulation, and let d be the remaining vertex of neighboring triangle that lies on the opposite side of edge bc . We assert that if d lies outside the circumcircle of $\triangle abc$, then no other site can lie within this circumcircle.

A formal justification will take too much work, so we'll just consider a limited scenario, which illustrates the key idea. Suppose that d is outside the circumcircle of $\triangle abc$ (the blue circle Fig. 4(d)) but (to the contrary) the vertex e opposite the edge cd lies within this circumcircle (see Fig. 4(d)). Consider the circumcircle of $\triangle cde$ (the red circle Fig. 4(d)). By an elementary (but somewhat tedious) analysis of the configuration of these points, it follows that b lies within this circumcircle. Since b is in the neighboring triangle to $\triangle cde$, this implies that the triangulation is *not* locally Delaunay, which yields the contradiction.

Because the algorithm checks that all the newly created triangles are locally Delaunay, the algorithm's correctness follows as a direct consequence.

Running-Time Analysis: To analyze the expected running time of algorithm we will establish two bounds, each averaged over all possible insertion orders. With the addition of each site:

- (1) $O(1)$ structural changes are made to the triangulation (in expectation), and
- (2) $O(\log n)$ time is spent determining which triangle contains each newly inserted site (in expectation).

These bounds depend *only* on the insertion order, not the distribution of the sites.

Bounding the Structural Changes: We argue first that the expected number of edge changes with each insertion is $O(1)$ by a simple application of backwards analysis. First observe that (assuming general position) the structure of the Delaunay triangulation is independent of the insertion order of the sites so far. Thus, any of the existing sites is equally likely to have been the last site to be added to the structure.

Suppose that some site p was the last to have been added. How much work was needed to insert p ? Observe that the initial insertion of p involved the creation of three new edges, all incident to p . Also, whenever an edge swap is performed, a new edge is added to p . These are the only changes that the insertion algorithm can make. Therefore the total number of changes made in the triangulation for the insertion of p is proportional to the *degree* of p after the insertion is complete (see Fig. 5). Although any one vertex may have a very high degree, we will exploit the fact that in a planar graph, the average vertex degree is just a constant.




Fig. 5: Number of structural changes is equal to p 's degree after insertion (three initial edges and three edge flips).

To perform the backwards analysis, we consider the situation after the insertion of the i th site. Let d_i be a random variable that indicates the degree of the newly inserted site in our randomized algorithm. Let $P_i = \{p_1, \dots, p_i\}$ denote the first i sites to be inserted. Although the diagram depends on which particular i sites are in this subset, our analysis will not. For $1 \leq j \leq i$, let $\deg(p_j)$ denote the degree of site p_j in triangulation $DT(P_i)$ just after the i th insertion.

Because the diagram does not depend on the insertion order, each of the sites of P_i has an equal probability of $\frac{1}{i}$ of being the last site to be inserted. Recall that (by Euler's formula), the triangulation has at most $3i$ edges. It is easy to see that the sum of vertex degrees is equal to twice the total number of edges (since each edge is counted twice), that is, $6i$. We conclude that expected value of d_i , denoted $E[d_i]$, satisfies:

$$E[d_i] = \frac{1}{i} \sum_{j=1}^i \deg(p_j) \leq \frac{6i}{i} = 6.$$

Therefore, by the magic of backwards analysis, the expected number of structural changes following the insertion of the i th site is, in expectation, just 6.

Bounding the Location Cost: The second aspect of the expected-case running time is the cost of determining which triangle contains each newly created site. As mentioned earlier, we will employ a bucketing approach, as we did with the trapezoidal-map algorithm. Think of each triangle of the current triangulation as a *bucket* that holds the sites that lie within this triangle and have yet to be inserted (see Fig. 6(a)). When a new site p is inserted, a number of

old triangles are deleted (shaded red in Fig. 6(a)) and a number of new triangles are created (shaded blue in Fig. 6(b)). All the points in the buckets of the old triangles need to be moved into the associated new triangle. This process is called *rebucketing*.




Fig. 6: Rebucketing points after inserting site p .

For the sake of simplifying the analysis, let us assume that the cost of rebucketing a single point during a single insertion is $O(1)$. (The issue is that the cost of rebucketing depends on the degree of the newly inserted site. In the previous section we showed that the average degree is a constant, so this assumption is not unreasonable.) We will show through a backwards analysis that, in expectation, any fixed site is rebucketed $O(\log n)$ times.

Let us fix a site $q \in P$. Consider the situation just after the insertion of the i th site. We may assume that q has not yet been inserted, since otherwise its rebucketing cost is zero after the i th insertion. For $1 \leq i \leq n$, let $X_i(q)$ denote the random event that q is moved to a new triangle after the i th insertion, and let $\text{Prob}(X_i(q))$ denote the probability of this event. Letting $B(q)$ denote the average number of times that q is rebucketed throughout the algorithm, we have

$$B(q) \leq \sum_{i=1}^n \text{Prob}(X_i(q)).$$

To bound $\text{Prob}(X_i(q))$, let Δ be the triangle containing q after the i th insertion. As observed above, after we insert the i th site, all the newly created triangles are incident to this new site. Thus, Δ would have come into existence as a result of the last insertion if and only if one of its three incident vertices happened to be the last to be inserted (see Fig. 6(c)). Since Δ is incident to exactly three sites, and every site is equally likely to be the last inserted, it follows that the probability that Δ came into existence is $\frac{3}{i}$. (We are cheating a bit here by ignoring the three initial sites at infinity.) Therefore, $\text{Prob}(X_i(q)) \leq \frac{3}{i}$.

From this, it follows that the expected number of times that the point q is rebucketed is

$$B(q) \leq \sum_{i=1}^n \frac{3}{i} = 3 \sum_{i=1}^n \frac{1}{i}.$$

Recall that $\sum_{i=1}^n \frac{1}{i}$ is the Harmonic series, and for large n , its value is very nearly $\ln n$. Thus we have

$$B(q) \leq 3 \cdot \ln n = O(\log n).$$

Although the diagram depends on the order in which the sites have been added, this bound does not. Summing over all the n sites, it follows that the total time spent rebucketing all the points is $\sum_{i=1}^n B(p_i) = O(n \log n)$.