


FLIP GRAPHS

VIVIAN KUPERBERG

1. TRIANGULATIONS AND DIAGONAL FLIPS

In this class, we're going to talk about triangulations. One setting in which triangulations come up a lot is that of triangulating regular n -gons, which is something like:



Remark 1.1. It's also possible to triangulate point sets that aren't regular n -gons. This is a pretty interesting generalization, and we'll come back to it, but for the moment we'll stick with the n -gons.

Now, let's say we have a set of points and several different ways to triangulate them. We'd like to say that some of our triangulations are closer than others. (Add a picture here)

One way of doing this is an idea of two triangulations differing just in one diagonal. This idea is made rigorous by the definition of a diagonal flip.

Definition 1.2. For an edge pq of a triangulation T contained in a quadrilateral, the *diagonal flip* of the edge consists in removing edge pq and replacing it with the other diagonal of the quadrilateral.

So we have triangulations, and we have these moves, diagonal flips, that let us go from one to another. We can then study the *flip graph* of the set of points.

Definition 1.3. The *flip graph* of an n -gon is the graph whose vertex set is the set of all triangulations of the n -gon, and which has an edge between two vertices if those triangulations can be related by a diagonal flip.

Example 1.4. Go through, with pictures (maybe handout?) the flip-graph of a regular hexagon.

There are various questions we might have about the flip-graph! For starters, is it even connected? In other words, if I have two triangulations of the same n -gon, can I always get from one to the other using diagonal flips? The answer (spoiler alert!) is yes, which is good, because we wanted to the flips to give us a picture of how closely related two triangulations are, so this tells us it gives us a pretty complete picture.

Proposition 1.5. *The flip-graph G_n of an n -gon is connected.*

Proof. Let's number our vertices counterclockwise by $0, 1, \dots, n - 1$. We'll show that any triangulation T is connected via diagonal flips to the triangulation F , a fan at 0, where every diagonal is through 0. Then we can connect any two triangulations T and S by going through our process to get from T to F , and then backwards to get from F to S .

Let $T \neq F$ be any triangulation of the n -gon. Consider the points connected to 0 via diagonals in T . This must include 1 and $n - 1$ via outside edges, so it's some nontrivial subset $S \subseteq \{1, \dots, n - 1\}$. If every vertex is connected to 0, then T would be F , which is not the case. Thus some vertex $1 < j < n - 1$ is missing. Let $i < j$ be the largest vertex before j that is connected to 0, and let $k > j$ be the smallest vertex after j that is connected to 0. Then $(i, 0, j)$ must form a triangle in our triangulation, so (i, j) is a diagonal of T . Flipping (i, j) gives one more diagonal through 0.

In particular, this argument shows that for any triangulation T with fewer than $n - 1$ diagonals through 0, a diagonal flip can be performed that will increase the number of diagonals through 0. We can keep performing this process until it terminates, which must be when we have reached F via diagonal flips. \square

Remark 1.6. Here is one fun application (working out the details is a good homework question).

If we have an associative product operation \cdot and some product $x_0 \cdot x_1 \cdots x_n$, we usually say that we feel comfortable not writing parentheses by associativity. But associativity only tells us that when multiplying three things a, b , and c , $a(bc) = (ab)c$. How do we know that *any* parenthesization of x_0, \dots, x_n is equivalent? There's a correspondence between parenthesizations of x_0, \dots, x_n and triangulations of an $(n + 2)$ -gon (both of which are examples of things counted by *Catalan numbers*). It turns out that performing one operation of the associativity law exactly corresponds to doing a diagonal flip. So the argument that the flip graph is connected is exactly what tells us that associativity of three things is enough to never worry about parentheses.

So, now that we know that the flip-graph is connected, a next question is its diameter. What's the largest possible flip distance between two triangulations of an n -gon? This is known as the *diameter* of G_n .

Theorem 1.7. *The diameter D_n of G_n satisfies*

$$D_n \leq 2n - 10 + \lfloor \frac{12}{n} \rfloor.$$

Proof. Just like our proof above with connectivity, this proof relies on flipping to and from star triangulations, where all diagonals go through one vertex.

Let T_1 and T_2 be two triangulations of an n -gon. For a point x on our n -gon, the degree $d_1(x)$ is the number of diagonals incident to x in T_1 , and analogously with d_2 and T_2 . For $i = 1, 2$, if $d_i(x) < n - 3$, then as in our previous argument the degree of x can be increased by an appropriate flip on T_i . The number of flips required to get from T_i to S_x , the star through x , is $n - 3 - d_i(x)$, based on just counting diagonals through S_x . Thus going from T_1 to T_2 through S_x takes $2n - 6 - d_1(x) - d_2(x)$ flips. Thus, the distance from T_1 to T_2 is bounded above by

$$D_n \leq \min_x (2n - 6 - d_1(x) - d_2(x)) = 2n - 6 - \max_x (d_1(x) + d_2(x)).$$

So we'd like to find a *lower* bound on $\max_x(d_1(x) + d_2(x))$ in order to get an *upper* bound on D_n . We can get this from taking an average of $d_1(x) + d_2(x)$ over all x . The sum of all degrees is twice the number of diagonals, which is $n - 3$, so

$$\begin{aligned} \sum_x d_1(x) + \sum_x d_2(x) &= 2n - 6 + 2n - 6 \\ \Rightarrow \frac{1}{n} \sum_x (d_1(x) + d_2(x)) &= \frac{1}{n}(4n - 12). \end{aligned}$$

Some vertex x must have an above average value of $d_1(x) + d_2(x)$, so $\max_x(d_1(x) + d_2(x)) \geq 4 - \lfloor \frac{12}{n} \rfloor$. Thus

$$D_n \leq 2n - 6 - \max_x(d_1(x) + d_2(x)) \leq 2n - 6 - 4 + \lfloor \frac{12}{n} \rfloor,$$

which is the desired result. □

So now the question is, is this bound tight? It seems like we didn't do very much work. Also, we guessed a silly way to get from one triangulation to another. If the bound is tight, we're saying that for some triangulations, this silly way is the best possible. However, the bound is indeed tight for $n \leq 18$. For $n \leq 8$, this is doable by hand (hw problem?). Surprisingly enough, the bound of $2n - 10$ is *also* known to be tight for certain large values of n (and in particular, infinitely often). The proof uses some hyperbolic geometry (no kidding!) and gets a bit gnarly, but we'll spend the next day or two going over an outline.

2. FROM TRIANGULATIONS TO POLYTOPES

We'll be working through an outline of the following result.

Theorem 2.1. *For large enough n , the diameter D_n of the flip-graph G_n is $2n - 10$. In particular, $D_n \geq 2n - 10$.*

We begin by the observation that if T_1 and T_2 are a pair of triangulations maximizing the flip-distance $\text{dist}(T_1, T_2)$, then they must have no diagonal in common. There can never be a reason for a common diagonal to be flipped, so if T_1 and T_2 shared a diagonal, we could flip a shared diagonal in T_1 to get T'_1 and T_2 , that are at least one more step away from each other. From now on, let T_1 and T_2 be a pair of triangulations with no shared diagonals.

We're then going to construct a tetrahedrated polyhedron from a flip path from T_1 to T_2 as follows. (Draw draw draw draaaaaw here picture with hexagon and tiles). We're going to glue the two triangulations T_1 and T_2 along their boundary to create a decoration on the boundary of a sphere. Now, each flip is going to become a tetrahedron. How? Well, the flip operation consists of taking a parallelogram containing a diagonal. Then we can imagine drawing the old diagonal faintly ('cause it's old), and the new diagonal firmly on top...and hey, we've just drawn a tetrahedron!

So, we have a base tiling T_1 . Then every flip is a tetrahedral tile we can glue on the previous tetrahedron. At every step, all faces have to perfectly match up. When we reach T_2 , we glue on the triangulation T_2 as the top face. Since T_1 and T_2 share no diagonals, Then we can imagine inflating our flat picture like a beach ball to get this polyhedron filled with tetrahedra. Let P_G be this polyhedron.

Black Box 2.2. If we inflate the ball so that all the tetrahedra are convex, we get something with rigid edges that looks like a polyhedron. In fact, this is always a convex polyhedron.

Assume that our flip sequence consists of t flips, with tetrahedra τ_1, \dots, τ_t corresponding to each flip. Then

$$\text{vol}(P_G) = \sum_{i=1}^t \text{vol}(\tau_i).$$

OK, so we want a lower bound on the number t of tetrahedra in our path. If we knew that some value V_Δ were an upper bound on the volume of a tetrahedron that can be inscribed in the polytope P_G , then we'd have

$$\text{vol}(P_G) \leq \sum_{i=1}^t V_\Delta,$$

implying that

$$t \geq \frac{\text{vol}(P_G)}{V_\Delta},$$

which would give us exactly the kind of bound we want. However, at the moment this seems pretty bad. We don't know what our polytope in \mathbb{R}^3 really looks like, and we could hypothetically inscribe a big tetrahedron that would have volume close to that of P_G itself. In that case, this would say that $t \geq 1$, which is not very helpful.

By one interpretation, the problem that we're running up against is that in Euclidean geometry, we can't bound the volume of tetrahedra very well. So, why don't we switch to working in a geometry where we can?

3. HYPERBOLIC GEOMETRY

Classically, the idea behind hyperbolic geometry is as follows. Instead of taking the axiom, as we do in Euclidean geometry, that given a line and a point not on the line, there exists exactly one line through the point parallel to the given line, we take the axiom that there are many lines through the point parallel to the given line. This leads to lots of strange properties! For example, and what we most need, triangles in hyperbolic 2-space and tetrahedra in hyperbolic 3-space have uniformly bounded volume, i.e. I can tell you a number c such that you cannot draw a triangle or a tetrahedron in hyperbolic space with volume greater than c . For tetrahedra, we won't prove this fact, but let's look a little bit at how hyperbolic space works in two dimensions, to give us an idea of why this might be true. There are many ways to model hyperbolic space; let's start with the *Poincaré disk*.


facts about hyperbolic space: Poincaré disk model, intuition that triangles have bounded volume, mention that optimal triangles have points on the sphere at infinity, half-space model and equivalence, discuss the point at infinity (TODO put this in)

So at this point, the goal becomes to construct a hyperbolic polyhedron whose boundary consists of gluing two triangulations T_1 and T_2 just like our P_G above, and which has large volume. Then we know that if we built a flip path from T_1 to T_2 , it would need some minimum number of flips because we need a minimum number of tetrahedra. Again, we're going to simplify things a bit. With a fairly straightforward construction, we can do pretty well. If you make the construction significantly uglier and more complicated, you can get the actual bound - if you're excited about that, talk to me at TAU!

Proposition 3.1. *For $k \in \mathbb{N}$ and $n = k^2 + 1$, there are triangulations T_1, T_2 of an n -gon with flip distance*

$$\text{dist}(T_1, T_2) \geq 2n - 4\sqrt{n} + O(1).$$

Proof. Consider the section of the triangular grid below.



For parallelogram sidelength k , this has $k^2 + 1$ points. We can fill it with ideal tetrahedra, with each tetrahedron having one vertex at v_∞ and base one triangle in our grid. There are $2(k - 1)^2$ of these triangles in the grid, so the volume of this polytope is $2(k - 1)^2 V_0$, with V_0 the maximal volume of a tetrahedron in hyperbolic space.

Thus every way to fill this polytope with tetrahedra requires at least $2(k - 1)^2 = 2k^2 - 4k + 2 = 2n - 4\sqrt{n} + O(1)$ tetrahedra. This is exactly what we want if we can show that there exist T_1, T_2 triangulations such that the polytope above is the flip polytope from one to the other.

But to do this, we really need to draw a Hamiltonian cycle, or a path going through all vertices exactly once, then joining to its end. As one example, we can zigzag through the grid, then connect to and from v_∞ at the end. This cycle has two sides, an “outside” and an “inside,” each of which is one triangulation, just as desired. \square

With a different, more complicated choice of polytope, one can prove:

Theorem 3.2 (Sleator, Tarjan, Thurston). *For sufficiently large n , the diameter of the flip-graph of an n -gon is $2n - 10$.*

Their polytopes are constructions based on subdividing faces of an icosahedron (draw a picture), then sometimes cutting out faces to get the right number n of vertices (draw a minipicture of this too). They then show that the tetrahedralization of these polytopes requiring the least number of tetrahedra is of cone type, so it must have $2n - 10$ idealized tetrahedra, so this is an appropriate large hyperbolic polytope.

This concludes our sketched discussion. We wanted to provide a lower bound on the diameter of the flip-graph, and our strategy, surprisingly enough, consisted of taking the problem up a dimension by making tetrahedrated polyhedra out of a path in the

flip-graph, then translating the problem once again to hyperbolic geometry so that we could use volume bounds, then finding special hyperbolic polyhedra that reached the bounds correctly. Whew!

4. ARBITRARY POINTSETS

Okay, but let's say we've decided now that we're really good at triangulating n -gons, and we want to triangulate an arbitrary finite set of points \mathcal{P} in the plane. We'll assume that \mathcal{P} is a "general" set of points, which basically means that if we encounter problems because the points are arranged in a very specific configuration, we can assume that that configuration doesn't occur. For example, if \mathcal{P} is a set of colinear points, we cannot triangulate \mathcal{P} , but if you randomly pick points in the plane, it would be very surprising if they were colinear, so we will assume that this is not the case. Similarly, we'll assume that no four of our points lie on a circle.


Definition 4.1. A *triangulation* of a set of points \mathcal{P} in the plane is a maximal non-crossing geometric graph with vertex set \mathcal{P} . In other words, it is a set of line segments between points of \mathcal{P} so that no two lines cross and we can't add any more line segments to preserve this property.

Remark 4.2. The outside edges will always be present in any triangulation. If they're not there, there's nothing stopping us from adding them.

First of all, it's worth noting that the flip-graph of a general pointset can be much gnarlier than that of an n -gon. For example, we're used to seeing triangulations where every diagonal is contained in a convex quadrilateral and thus flippable, but in this settings there are cases where this doesn't hold! (Draw star trek logo). So now we have a new question: what are the maximal and minimal degrees in a flip-graph on n vertices?

4.1. Maximal degrees. For the maximal question, we may feel like we can do pretty well with an n -gon, where we can flip $n - 3$ diagonals out of $2n - 3$ total. However, in some cases we can do better!

Let T be the triangulation of the following figure.



Each line is a slightly bent chain of three edges, so only the five hull edges are non-flippable. By iterating this, for any $n \cong 1 \pmod{5}$, we can construct a triangulation on n points with $n - 5$ flippable edges. We can also get $n \cong 0 \pmod{5}$ by removing the central point. Question to think about that I don't know the answer to: what about for other large n ? Can certainly get pretty close, to $n - O(1)$.

Proposition 4.3. *Let S be a triangulation of $n \geq 5$ points with a convex hull of only 3 or 4 edges. Then S has at least 6 non-flippable edges.*

Proof. See handwritten notes for Viv's lecture reference (regular case) / homework problem for the nonregular case. \square

4.2. Minimal degrees. For the minimum degree in a flip-graph, we can get a precise answer.

Proposition 4.4. *Any triangulation T of a set \mathcal{P} of n points in the plane contains at least $\frac{n}{2} - 2$ flippable edges.*

We'll start with this lemma.

Lemma 4.5. *Let T be a triangulation of a set \mathcal{P} of n points. Let γ be the number of points of \mathcal{P} contained in its convex hull. Then T has $3n - 3 - \gamma$ edges and $2n - 2 - \gamma$ bounded triangular faces.*

Proof. This follows from the idea of *Euler characteristic*, which we're using so narrowly that we'll barely speak of it. If G is any planar graph with n vertices, e edges, and f faces,

then $n - e + f = 1$. This can be shown by induction. If we start with one vertex, then $e = f = 0$, so $n - e + f = 1$. Then we can either add a vertex and an edge, or add an edge between two existing vertices, which also adds a bounded face. Either move keeps the value $n - e + f$ constant, so it must always be 1.

Thus $n - e + f = 1$. But also, each face is a triangle with three edges and each edge save for the outer γ edges is contained in two faces, so the number of pairs (F, E) with F a bounded face and E an edge in F is counted both by $3f$ and by $2e - \gamma$, so $3f = 2e - \gamma$. Solving both these equations for e and f in terms of n and γ gives $e = 3n - 3 - \gamma$ and $f = 2n - 2 - \gamma$. \square

Now we prove the proposition.

Proof. Certainly, the γ edges of the convex hull cannot be flipped. Every non-flippable interior edge e has one endpoint, say u , so that the sum of the two angles adjacent to e at u is equal to or exceeds π . In other words, it's a star trek shape in some direction. Then orient each non-flippable edge e away from u .

An important aspect of this orientation is that if two edges are oriented away from the same point u , they must share an angle at u , assuming (due to general position) that we never have perfectly right angles. Thus a point u can have at most three outward oriented edges, and this maximum can only be attained by points of degree 3 in T . Let η_i be the number of interior points with i outward edges, for $0 \leq i \leq 3$. Then

$$\eta_3 + \eta_2 + \eta_1 \leq n - \gamma,$$

since the left hand side represents a subset of all interior points.

Let's call a corner u of a triangle Δ in T the *root* of Δ if the two edges of Δ incident to u are both outward oriented at u . This implies that every triangle of T has at most one root. If u has out-degree two, then exactly one triangle is rooted at u , since the edges must share an angle. If u has out-degree three, then there are exactly three triangles rooted at u , namely the three triangles intersecting u . In all other cases, u is not a root. Counting roots of triangles on one side and all triangles of T on the other, we get

$$3\eta_3 + \eta_2 \leq 2n - 2 - \gamma.$$

Thus adding $3/2$ our first inequality with $1/2$ our second gives

$$3\eta_3 + 2\eta_2 + \frac{3}{2}\eta_1 \leq \frac{5}{2}n - 2\gamma - 1.$$

The number of interior unflippable edges is given by $3\eta_3 + 2\eta_2 + \eta_1$. Thus the above inequality implies that there are at most $\frac{5}{2}n - 2\gamma - 1$ such edges. Together with the hull edges, this gives at most $\frac{5}{2}n - \gamma - 1$ unflippable edges. Thus the number of flippable edges of T must be at least


$$(3n - 3 - \gamma) - \left(\frac{5}{2}n - \gamma - 1\right) = \frac{1}{2}n - 2,$$

as desired. \square

This bound can be attained! It is sharp. One family of extremal examples is as follows: let T be a triangulation of a set of m points in convex position. We can take the corners of a regular m -gon, or points that are a little off from these if we want them to be in general position. T has $m - 3$ interior edges, with $m - 2$ triangles. Subdivide each triangle with a new point connected to the three corners (draw a picture), which gives a triangulation T^*

of the set of $n = 2m - 2$ points. The flippable edges are the $m - 3 = \frac{n}{2} - 2$ interior edges of T , and none of the newly added edges.

Another family is exhibited in the following picture.



To see that the flip-graph is connected, we're going to define *Voronoi regions* and *Delaunay Triangulations*.

5. DELAUNAY TRIANGULATIONS

Definition 5.1. Let \mathcal{P} be a set of n points in \mathbb{R}^2 . For $p \in \mathcal{P}$, the *Voronoi region* $V(p)$ of p is the set of all points x that are at least as close to p as to any other point in \mathcal{P} . In other words,

$$V(p) = \{x \in \mathbb{R}^2 \mid \|x - p\| \leq \|x - q\| \text{ for all } q \in \mathcal{P}\}.$$

Draw a picture here!

Lemma 5.2. For \mathcal{P} a set of points in \mathbb{R}^2 and $p \in \mathcal{P}$, $V(p)$ is a possibly infinite convex polygonal region, i.e. an intersection of halfplanes.

Proof. Let $q \in \mathcal{P}$ be any other point. Then let $V_q(p)$ be defined by

$$V_q(p) = \{x \in \mathbb{R}^2 \mid \|x - p\| \leq \|x - q\|\}.$$

Then $V_q(p)$ is a halfspace. (Draw a picture).

Also, $V(p) = \bigcap_{q \in \mathcal{P} \setminus \{p\}} V_q(p)$, so $V(p)$ is a possibly infinite convex polygonal region. \square

Consider the set of all Voronoi regions; every point $x \in \mathbb{R}^2$ has some closest point in \mathcal{P} , so it belongs to some Voronoi regions. Thus the Voronoi regions divide up all of \mathbb{R}^2 , like this: (Draw a picture)

Definition 5.3. For $p, q \in \mathcal{P}$, if $V(p)$ and $V(q)$ share an edge, we call p and q *Delaunay neighbors*. The *Delaunay triangulation* of a point set \mathcal{P} is the graph obtained by connecting all pairs of Delaunay neighbors by straight edges.

Proposition 5.4. *The Delaunay triangulation is a triangulation for a general set of points.*

Proof. Let T be the Delaunay “triangulation,” which we don’t yet know is a triangulation. Voronoi edges correspond to points that are equally close to the two closest points in our set; Voronoi vertices correspond to points that are equally close to the three or more closest points in our set. If a face in our Delaunay configuration is *not* a triangle, this corresponds to four or more points being equally close to some point in the plane. But then those four or more points must all lie on a circle centered at our Voronoi point, which doesn’t happen because our points are in general position.

Also, no two Delaunay edges will cross. If they do, we have crossing Voronoi lines (draw a picture!!! above, but especially here), which contradicts our definition of Voronoi region.

This (explain this more) shows that the Delaunay triangulation is in fact a triangulation. \square

The idea of thinking of Delaunay triangles as lying on circles is quite useful. Let’s look a little bit more at circles in our pointset.

Definition 5.5. A circle C is an *empty circle* for \mathcal{P} if there is no point of \mathcal{P} in the interior of C .

For a Delaunay triangle p, q, r , there is a corresponding Voronoi vertex $v = V(p) \cap V(q) \cap V(r)$. As in the argument above, the largest empty circle with center v has p, q , and r on the boundary, so it must be the unique circle through p, q , and r , i.e. the circumcircle for the triangle p, q, r . Meanwhile, if p, q, r are three points of \mathcal{P} with an empty circumcircle, then the center of that circle is a Voronoi vertex in $V(p) \cap V(q) \cap V(r)$, so p, q , and r form a Delaunay triangle. This gives the following.


Fact 5.6. Let $\mathcal{P} \subseteq \mathbb{R}^2$ be a set of n points in general position. A triple $p, q, r \in \mathcal{P}$ is a Delaunay triangle if and only if the circumcircle of p, q, r is empty.

So intuitively, the triangles in the Delaunay triangulation are the most “logical” triangles; it is more likely to have smaller and equilateral-ish triangles than long and thin triangles, although this is certainly not always true.

So now we have a new “special” triangulation, and once again our strategy for proving that the flip-graph is connected is going to be quantifying how far away we are from that triangulation. Given a triangulation T of \mathcal{P} , we can measure how far it is from the Delaunay triangulation by counting points contained in circumcircles of triangles in the triangulation. Then, we reduce this number by performing *Lawson flips*.

Definition 5.7. Let T be a triangulation of \mathcal{P} containing triangles p, q, r and p, q, s such that s is in the interior of the circumcircle of p, q, r . Then p, q, r, s form a convex quadrilateral, and so we can flip the edge p, q . In this case, the edge p, q is called a *weak edge* of the triangulation T , and the flip is known as a *Lawson flip*.

See the picture below.



Proposition 5.8. *Let T be an arbitrary triangulation of a set \mathcal{P} of n points in general position. Any algorithm that repeatedly performs Lawson flips will reach the Delaunay triangulation of \mathcal{P} with at most $\binom{n}{2}$ flips.*

Note that this will mean that the Delaunay triangulation is the *only* triangulation of \mathcal{P} that has no weak edge. This also implies that the flip-graph is connected with diameter at most $n^2 - n$. This is not as good a bound as we had for the n -gon case - in fact it's pretty bad - but it is what we have.

We're going to prove this by lifting our triangulation into space! Imagine the triangulation as living on a table. We'll drag each point up to a different height. Namely, let's drag a point $p = (p_1, p_2) \in \mathcal{P}$ to the point $\hat{p} = (p_1, p_2, p_1^2 + p_2^2)$, so that all of our points lie on the paraboloid $z = x^2 + y^2$, and each point is moving straight up. We'll lift our triangulation T of \mathcal{P} by lifting each triangle p, q, r to the spatial triangle (not on the paraboloid) with corners $\hat{p}, \hat{q}, \hat{r}$. The crucial property of this lifting is as follows:

Lemma 5.9. *A point s is in the interior of the circumcircle of p, q, r if and only if \hat{s} lies below the plane spanned by $\hat{p}, \hat{q}, \hat{r}$.*

This proof is a bit of linear algebra (specifically, determinants). If you'd like to treat this as a black box, that is OK.

Proof. To prove this, let's go up a dimension! For $a = (a_1, a_2, a_3) \in \mathbb{R}^3$, let $a^+ = (a_1, a_2, a_3, 1)$ in \mathbb{R}^4 .

Fact 5.10 (Linear Algebra Fact). Four points $a, b, c, d \in \mathbb{R}^3$ lie in the same plane if and only if the determinant $|a^+, b^+, c^+, d^+|$ vanishes. If the determinant is negative instead of vanishing, then standing at d and looking at the triangle (a, b, c) , the vertices are listed in counterclockwise order.

Let (p, q, r) be a counterclockwise triangle in \mathbb{R}^2 . Consider the mapping $\varphi_{p,q,r} : \mathbb{R}^2 \rightarrow \mathbb{R}$ given by $\varphi_{p,q,r}(s) = \det|\hat{p}^+, \hat{q}^+, \hat{r}^+, \hat{s}^+|$. This determinant vanishes if the lifted point \hat{s} is in the plane spanned by $\hat{p}, \hat{q}, \hat{r}$. Otherwise, the sign tells us whether \hat{s} lies above or below the plane. We'd like to show that $\varphi_{p,q,r}(s) = 0$ if and only if s is on the circumcircle C of p, q, r , and that $\varphi_{p,q,r}(s) > 0$ if and only if s is in the interior of C . Let m be the center of C , and let

$\psi_{p,q,r} = \varphi_{p-m,q-m,r-m}$ be the mapping $\mathbb{R}^2 \rightarrow \mathbb{R}$ given by

$$\psi_{p,q,r}(s) = \det|p \hat{-} m^+, q \hat{-} m^+, r \hat{-} m^+, \hat{s}^+|.$$

By linearity of the determinant, $\psi_{p,q,r}(s-m) = \varphi_{p,q,r}(s)$. However, if $p = (p_1, p_2)$ and likewise with q, r, s, m , and if ρ is the radius of C , then $(p_1 - m_1)^2 + (p_2 - m_2)^2 = \rho^2$ and the same for q and r (draw a picture). Thus

$$\psi_{p,q,r}(s-m) = \det \begin{bmatrix} p_1 - m_1 & p_2 - m_2 & \rho^2 & 1 \\ q_1 - m_1 & q_2 - m_2 & \rho^2 & 1 \\ r_1 - m_1 & r_2 - m_2 & \rho^2 & 1 \\ s_1 - m_1 & s_2 - m_2 & (s_1 - m_1)^2 + (s_2 - m_2)^2 & 1 \end{bmatrix}.$$

If $s \in C$, then $(s_1 - m_1)^2 + (s_2 - m_2)^2 = \rho^2$, and the last two columns are multiples of each other, so this value is 0. If not, then $\varphi_{p,q,r}$ must be nonzero. Thus $\varphi_{p,q,r}$ has constant sign on interior points of C , so we can pick any to test in order to show that $\varphi_{p,q,r}$ is positive on interior points. Let's pick m , since that is a particularly nice value.

$$\varphi_{p,q,r}(m) = \psi_{p,q,r}(0) = \det \begin{bmatrix} p_1 - m_1 & p_2 - m_2 & \rho^2 & 1 \\ q_1 - m_1 & q_2 - m_2 & \rho^2 & 1 \\ r_1 - m_1 & r_2 - m_2 & \rho^2 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

But this is just ρ^2 times the determinant of the points p^+, q^+, r^+ (say a bit more with picture), and this determinant is positive since p, q, r is a counterclockwise triangle. This completes the proof. \square

Lawson flips suffice proof. Let T be a triangulation with a weak edge p, q , and let $T \rightarrow T_f$ be the Lawson flip replacing p, q by r, s . Then the lifted triangulations \hat{T} and \hat{T}_f enclose the tetrahedron $\hat{p}, \hat{q}, \hat{r}, \hat{s}$, by the same inflate-a-beach ball argument that we used before. By the lemma, \hat{T} contains the two upper triangles $\hat{p}, \hat{q}, \hat{r}$ and $\hat{p}, \hat{q}, \hat{s}$ of the tetrahedron, and \hat{T}_f the lower triangles $\hat{p}, \hat{r}, \hat{s}$ and $\hat{q}, \hat{r}, \hat{s}$. Thus the surface \hat{T}_f is below the surface \hat{T} , and the edge \hat{p}, \hat{q} of \hat{T} is above \hat{T}_f . Then Lawson flips consistently lower a sequence of surfaces, and we can never return to an edge that has been flipped away. Since each edge can be flipped at most once, there are at most as many Lawson flips as there are possible edges on n points, namely $\binom{n}{2}$.

We now need only show that the Delaunay triangulation is the unique triangulation with no weak edges, so that the process terminates in the Delaunay triangulation. If T is not Delaunay, then there exists some triangle p, q, r with a circumcircle C containing a point $s \in \mathcal{P}$ in its interior. By relabeling, assume that p, q and r, s are the diagonals of the quadrilateral p, q, r, s (picture!!!). By the lemma, the segment \hat{r}, \hat{s} lies below \hat{p}, \hat{q} . Thus \hat{T} is not convex, so it contains a non-convex edge \hat{a}, \hat{b} . This edge must be part of triangles a, b, c and a, b, d , but since \hat{a}, \hat{b} is non-convex, the point \hat{d} is below the plane spanned by $\hat{a}, \hat{b}, \hat{c}$. Thus a, b is a weak edge, so it allows a Lawson flip. \square

As a corollary, we've shown that the Delaunay triangulation is just the vertical projection of the lower convex hull of the point set lifted to the paraboloid.

Hopefully I'll stop here - if not, I can talk about secondary polytopes (but probably handwritten notes).