2.3 Triangulations as simplicial complexes and f-vectors

As we defined triangulations early on this chapter we did not consider the role of the faces of the simplices involved in the triangulation. We introduce now an important language that is general enough for our constructions and explanations later on:

A *simplicial complex* K is a finite family of simplices in $\mathbb{R}^d$ such that

- If $F \in K$ and G is a face of F, then $G \in K$.
- If $F, G \in K$ then $F \cap G$ is a face of both F and G.

The elements of a simplicial complex K will be called *cells*. The zero dimensional polyhedra belonging to K will be called the *vertices* of K and denoted $V(K)$. Note that the set of all faces of the simplices that form a triangulation constitute a simplicial complex.

It is common to specify a simplicial complex by the list of its maximal faces as we did for triangulations. In our discussions sometimes we will need to use this additional structure of triangulations. For a simplicial complex K we denote by $|K|$ the underlying topological space of the simplicial complex. This is the union of the simplices seen as subsets of some Euclidean space. If $|K|$ is homeomorphic to a ball or sphere we see that K is a *simplicial ball* or *simplicial sphere* respectively.

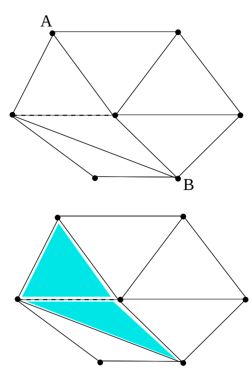

Figure 2.9: The link (two vertices A,B) and the star (two triangles) of an edge

The *boundary* ∂K of a simplicial complex K is the simplicial complex obtained as the union of all faces contained in a $d-2$-face of K that is contained in exactly one facet of K. The boundary of a simplicial ball is a simplicial sphere. The boundary of a simplicial sphere is empty. We say that a simplicial complex F is a *subcomplex* of K if every face of F is a face of K. The boundary of K is a subcomplex of F.

Let K be a simplicial complex. For any $F \in K$, the *star* of F in K, $st(F, K)$, is the subcomplex of K made up of all simplices of K having F as a face plus all their faces. The *link* of F in K is the simplicial complex $lk(F, K) = \{C \in st(F, K) | F \cap C = \emptyset\}$. Note that the link of the $(k-1)$-

dimensional face F, $Lk(F,K)$, is a simplicial complex of dimension $d-k-1$.

If F and G are two simplices with distinct vertices in $\mathbb{R}^d$ such that the totality of these vertices is at most $d+1$ and they are in general position in $\mathbb{R}^d$, then these vertices span a simplex called the *join* of F and G. This will be denoted by $F * G$. We note that the join of a p-simplex with a q-simplex is a $(p+q+1)$-simplex.

For an arbitrary simplicial complex there is the notion of f-*vector*, which has been central in the development of the combinatorial theory of polytopes [101]. For a triangulation T of a d-dimensional point configuration, we define its f-vector to be $f(T) = (f_{-1}(T), f_0(T), f_1(T), f_2(T), \ldots, f_d(T))$, i.e. the integral vector whose $f_i(T) =$ number of i-dimensional simplices inside T. Of course, in this definition $f_{-1}(T) = 1$ for the empty set.

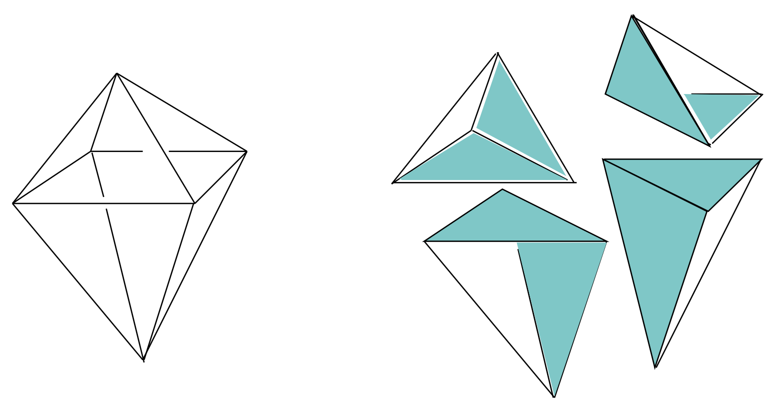

Figure 2.10: The f-vector of the triangulated regular octahedron is $f(T_{octahedron}) = (1, 6, 13, 12, 4)$

What can be said about the f-vector? How do the entries grow? We can begin by studying the entry $f_d(T)$. We call the *size* of a triangulation T the number d-dimensional simplices (top dimensional simplices) used in the triangulation. How large or small can the size be in terms of $f_0(T)$ (the number of vertices)? We begin with some results about the size:

Suppose we have a triangulation of an n point set A in $\mathbb{R}^d$. What is the smallest size of a triangulation? Let t be the number of d-simplices in a triangulation $T = \{S_1, S_2, \ldots, S_t\}$ of A. Think of the following graph: Take a

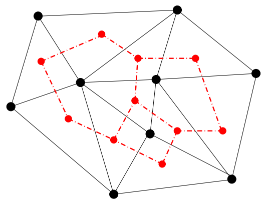

Figure 2.11: The dual graph of a triangulation appear in red

node for each d-simplex and join two of them by an edge if they share a common $d-1$-face. This will be called the *dual graph* of a triangulation for obvious reasons.

Not much is known about dual graphs of triangulations (see exercises) but they are useful to estimate a lower bound for the size. The dual graph of T has t nodes and it must be connected because triangulations are connected sets. The set of vertices A is the union of the set of vertices of the simplices $S_1, \dots, S_t$. Since the graph is connected we can index the simplices such that the dual subgraph associated with a partial list of simplices, say $\{S_1, S_2, \dots, S_m\}$ with $m \leq t$, is connected. If V_m is the set of vertices in the union of $\cup_{i=1}^m S_i$ then V_1 has cardinality $\#V_1 = d+1$ and V_t has cardinality n. Moreover $\#V_m \leq \#V_{m-1} + 1$. Because the new simplex S_m has at least d vertices in common with V_{m-1}. In conclusion,

$$n = \#V_t \leq (d+1) + (t-1) = d + t$$

and we have the following theorem:

Theorem 2.3.1. *The size of a triangulation for an n point set A in $\mathbb{R}^d$ is at least $n-d$. Moreover the equality is achieved precisely if one of the following equivalent conditions occurs:*

- *The dual graph of the triangulation is a tree.*
- *No $(d-2)$-face of the triangulation intersects the interior of the convex hull of A.*

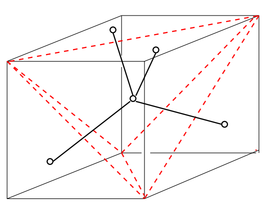

Figure 2.12: A triangulation of the 3-cube with 5 tetrahedra. Its dual graph is a tree

We simply need to explain the details of the second part. Clearly when the dual graph of the triangulation is a tree our labeling of the simplices S_i can be such that $\#V_m = \#V_{m-1}+1$ hence the end result becomes an equality. Conversely, if the dual graph of the triangulation has a cycle and S_{m+1} is the last simplex closing the cycle, then all the vertices in V_{m+1} lie already in V_m, thus inequality will be strict in that case.

We want to see now that having a dual graph which is a tree is equivalent to excluding interior $(d-2)$-faces (interior edges in the case of three dimensions for example). Suppose for a moment the dual graph contains a cycle.

One can imagine the cycle has as nodes the barycenters of the face $S_i \cap S_{i+1}$ and that its edges are the line segments joining successive nodes. If we shrink this cycle it will remain in the interior. At some point it must intersect one of the boundary faces of $S_i \cap S_{i+1}$, thus this face cuts through the interior of $conv(A)$. Conversely, let F be an interior $(d-2)$-face of the triangulation. It belongs to a d-simplex Sa and let p be a point in the intersection of the relative interior of F with the interior of $conv(A)$. Consider a circle C with center in p and lying in the 2-plane perpendicular to the affine hull of F. If the radius is small enough the circle will be contained in $conv(A)$ and its relative interior does not intersect any other $(d-2)$-face besides F. Then the d-simplices intersected by the circle form a a cycle.

What is the largest size of a triangulation? We will see below the answer is an application of the deep upper bound theorem for spheres proved by Stanley in 1980 [82]: The cyclic d-polytope with n vertices, denoted by $C(n,d)$, is the convex hull of n vertices taken from the moment curve $(t, t^2, t^3, \dots, t^d)$. The faces of this polytope form a simplicial d-sphere.

Theorem 2.3.2 (Upper bound theorem). *For any simplicial d-sphere S*

$$f_i(S) \leq f_i(C(n+1, d+1)), \quad 0 \leq i \leq d.$$

where $C(n,d)$ denotes the simplicial complex made of the boundary faces of the cyclic d-polytope with n vertices.

Lemma 2.3.3. *Suppose K is a pure simplicial $(d-1)$-complex and F is a face of dimension $k-1$ inside K. The following relations hold:*

$$f_j(K) = f_j(K - F) + f_{j-k}(lk(F,K)), \text{ for } -1 \leq j \leq d-1.$$

Here is the reason this is true: Suppose G is a face of K of dimension j. if F is contained in G then there is a $j-k$-face H in $lk(F,K)$ such that $G = F \cup H$. Otherwise, if G is an element of $K - F$. In either case the presence of G was counted on one summand of the formula. Now we are ready to present the upper bound on the the size of triangulations. This result is also known to be tight.

Corollary 2.3.4. *The size of a triangulation of a point set* A *with* n *vertices inside* $\mathbb{R}^d$ *is bounded above by* $f_d(C(n+1, d+1)) - (d+1)$. *In other words The largest size of a triangulation is asymptotically* $O(n^{\lceil (d+1)/2 \rceil})$. *Similar bounds hold for* i*-th entry of the* f*-vector of a triangulation.*

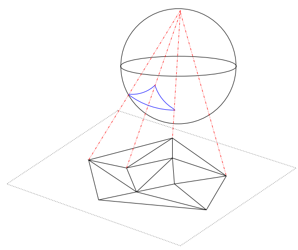

Figure 2.13: "printing" a triangulation in the plane on the surface of the sphere

In order to see this we embed the triangulation in question inside a simplicial d-sphere. Embed the points of A inside $\mathbb{R}^{d+1}$ by positioning the points in the plane $x_{d+1} = \epsilon$. Let S^d be the unit sphere in $\mathbb{R}^{d+1}$ with center at the origin. Project the set A from the origin onto the surface of S^d. The simplices of T transfer to be "spherical simplices". This is some kind of "reverse stereographic projection" as shown in Figure 2.13

Finally take the south pole v of the sphere and join it (over the surface of the sphere) with the boundary $(d-1)$-faces of the original triangulation. The resulting complex is a simplicial d-sphere T'. This is a very useful trick, because it allows us to transfer results about simplicial spheres to results about triangulations of point sets. For example, f-vector of T' can be recovered from the f-vector of T and the f-vector of ∂T. More precisely, we claim that for a d-dimensional triangulation and its associated d dimensional sphere T'.

$$f_j(T') = f_j(T) + f_j - 1(\partial T), \quad -1 \leq j \leq d.$$

For this, note the way we constructed the simplicial complex T', $T = T' - v$ and $lk(v, T') = \partial T$. Thus the formula we stated follows from Lemma 2.3.3. Finally, from the upper bound theorem, we have that size of T' is bounded above by size of T plus the number of boundary $(d-1)$-faces of T. That last number is at least $d+1$ clearly. Hence the first part of the result follows. The second part follows from the specific values of f-vectors of cyclic polytopes.

We conclude this section stating one of the most important formulas involving the f-numbers, the *the Euler relation*. Proofs of this result can be found in many sources in algebraic topology. We will use Euler's formula heavily when we study the space of planar triangulations (in particular how they are connected). More equations of this kind exist (more about this later!).

Lemma 2.3.5. *Suppose* K *is a* $d-1$*-simplicial complex associated to a triangulation of a point set in* $\mathbb{R}^{d-1}$*. Then* $\sum_{j=1}^{d}(-1)^{d-j}f_{j-1}(K) = (-1)^{d-1}$

2.4 Flips and the Graph of Triangulations

We are about to introduce local operations that, from an initial triangulation, produce other "neighboring" triangulations. To define the operations we look at triangulations of the smallest non-trivial sets of points the circuits. The structure of a triangulation as simplicial complex plays a role in the formal definition of the operation too.

We say a subset of points $Z \subset A$ is a *circuit* of A if any proper subset Z' is affinely independent but Z is affinely dependent. This notion comes from the theory of matroids (see [11]). We observe that, up to real scalar multiple, there is a unique real affine relation among the elements of Z. We can decompose Z into those elements Z_+ that appear with positive coefficient in the unique affine relation, and $Z_- = Z \backslash Z_+$. Geometrically, the circuits correspond to the minimal *Radon partitions* of the configuration. A Radon partition consists of disjoint subsets of Z_+ and Z_- that satisfy $\mathrm{relint}(\mathrm{conv}(Z_+)) \cap \mathrm{relint}(\mathrm{conv}(Z_-)) \neq \emptyset$.

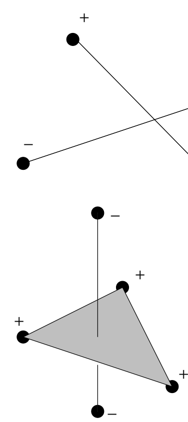

Figure 2.14: Two circuits represented as Radon partitions

Given a circuit $Z \subset A$ we define two triangulations $k_+(Z)$ and $k_-(Z)$ of $\mathrm{conv}(Z)$ as follows: $k_+(Z)$ as the collection of simplices $\{Z - \{p\} | p \in Z_+\}$ and $k_-(Z) = \{Z - \{p\} | p \in Z_-\}$.

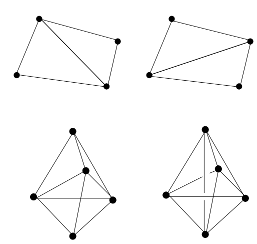

Figure 2.15: And their associated triangulations

Lemma 2.4.1. *If* Z *is a circuit, then*

1. *The collections of simplices* $k_+(Z) = \{Z - \{p\} | p \in Z_+\}$ *and* $k_-(Z) = \{Z - \{p\} | p \in Z_-\}$ *are triangulations.*

2. *These two are the only triangulations possible.*

Proof. We verify the first point: Pick a point $a \in Z$, subdivide $\mathrm{conv}(Z)$ with a placing $s_a^+(\mathrm{conv}(Z))$. We observe that if a belongs to Z_+, then the placing triangulation $s_a^+(conv(Z))$ is precisely $k_+(Z)$. When a belongs to Z_- we obtain $k_-(Z)$. A simplex whose vertices belong to Z must be a simplex of either $k_+(Z)$ or $k_-(Z)$. On the other hand any maximal simplex $F \in k_+(Z)$ intersects improperly with all the maximal simplices of $k_-(Z)$. This is the case because a pair of simplices $F \in k_+(Z)$ and $G \in k_-(Z)$ contains the Radon partition associated with the circuit Z. In conclusion the only candidates we have for triangulations of Z are precisely $k_+(Z)$ and $k_-(Z)$. □

Let T be a triangulation of A and $Z \subset A$ a circuit. We say that Z is *flippable* in T if the following conditions are satisfied:

1. One of the triangulations $k_+(Z)$ or $k_-(Z)$ is a subcomplex of K.

2. Let $F_1, F_2, \ldots, F_r$ be the maximal dimensional simplices of $k_+(Z)$ (similarly $k_-(Z)$), then the link of F_i in K is the same simplicial complex L for all $i = 1, \ldots, r$.

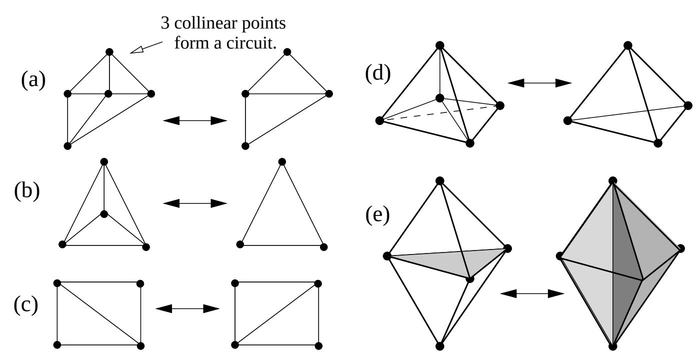

Figure 2.16: Examples of geometric bistellar flips.

Now we are ready for a very important definition: Observe that if a triangulation K is supported on the circuit Z, then we obtain a new triangulation of A, as follows: replace all the joins $F * G$ with $F \in k_+(Z)$ and $G \in L$, with the joins $F' * G$ with $F' \in k_-(Z)$ and $G \in L$. We denote this new triangulation by $\mathrm{flip}_Z(K)$.

This operation of changing the triangulation from K to $\mathrm{flip}_Z(K)$, or vice versa, is called a *flip* or *geometric bistellar operation*. The conditions for doing the flip are topological (e.g. link condition) and geometric (circuit Z is triangulated). This provides us with the required notion of adjacency for triangulations: The graph of triangulations of a point configuration A is the graph G_A whose vertices are the distinct triangulations of A. Two of its vertices are adjacent if the two triangulations are supported on a common circuit. Equivalently, two triangulations are adjacent if one can be obtained from the other by a flip.

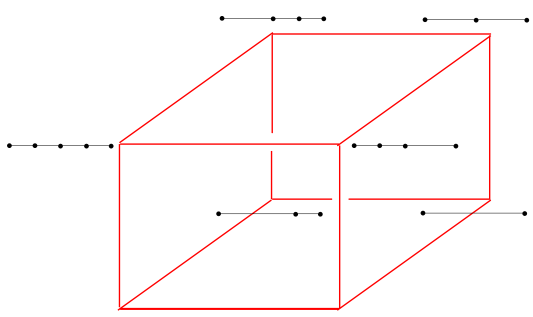

Figure 2.17: The graph of triangulations of a 5 point set in the real line

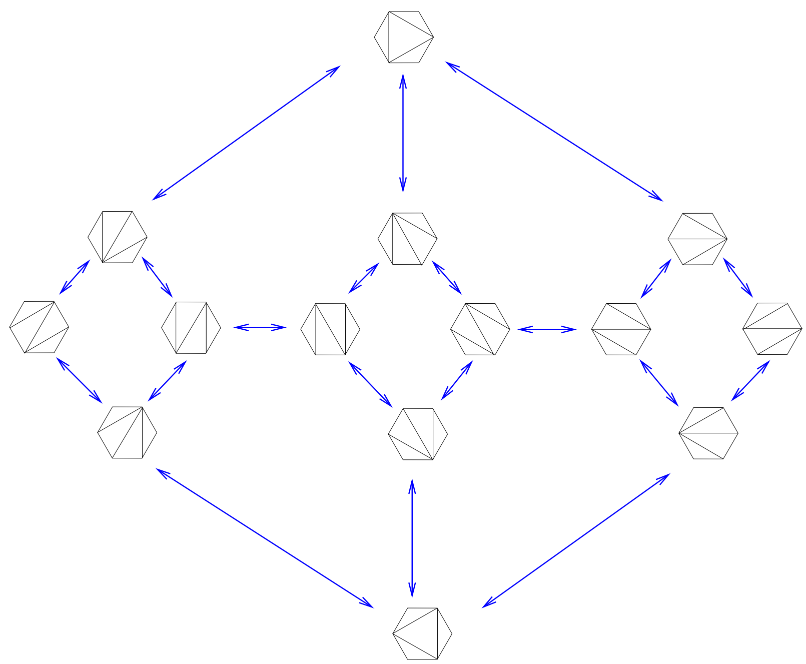

Figure 2.18: The Graph of flips for a hexagon, with one edge missing.

3.2 How many Triangulations are there?

Given a point set one would like to know how many planar triangulations are there. This turns out to be a very difficult question. There is only one family in the plane for which we can give an exact answer: this is when the points are vertices of a convex polygon. We present two proofs of the following count:

Theorem 3.2.1. *The number of triangulations of a convex n-gon is $\frac{1}{n-1}\binom{2n-4}{n-2}$, the Catalan numbers. Therefore the number of triangulations is of the order $O(2^{2n})$.*

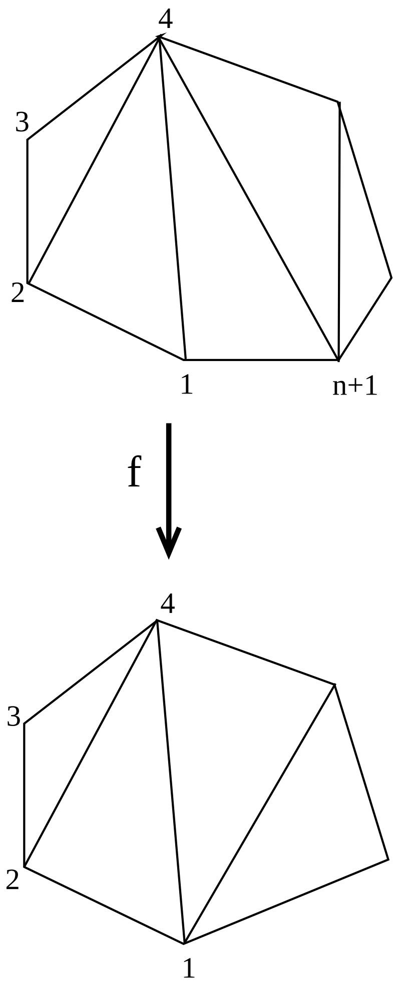

Figure 3.1: The contracting map.

Proof. 1) Assume the vertices of the n-gon are labeled from 1 to n in clockwise order. Denote by T_n all triangulations of an n-gon. Denote by $R(n)$ its cardinality, i.e. the number of ways to triangulate an n-gon. There is a nice surjective map f from T_{n+1} onto T_n. A triangulation in T_{n+1} is mapped to a triangulation in T_n obtained by contracting the boundary edge $\{1, n+1\}$ (see Figure **??**).

First of all, observe that for a triangulation $t \in T_n$ the cardinality of $f^{-1}(t)$ equals precisely the number $\deg_1(t)$ of arcs touching vertex 1 in t. This is true because each edge incident to 1, and point 1 itself, can be "doubled" to become the new point $n+1$ and the edge $\{1, n+1\}$. Note also if t_1, t_2 are two triangulations the inverse images $f^{-1}(t_1)$ and $f^{-1}(t_2)$ are disjoint. Therefore $R(n+1) = \sum_{s=1}^{n} s\#\{t \in T_n | \deg_1(t) = s\}$. Dividing by $R(n)$ we obtain that $R(n+1)/R(n)$ equals the expected degree of vertex 1 on a random triangulation. This expected value is easy to compute, since the sum of the degrees at a node at a triangulation in T_n, equals twice the number of edges in the triangulation, which equals $2(2n-3)$. The average degree on a node is then $2(2n-3)/n$. This is the expected degree on node 1.

Now we are left with solving the recursion $R(n+1) = (2(2n-3)/n)R(n)$. It is interesting to note that this is a different recursion as the one we obtain later in the second proof. We know that $R(3) = 1$, thus we can easily proof the stated formula is satisfied by $R(n)$ inductively. Indeed, if we assume now that $R(n) = \frac{1}{n-1}\binom{2n-4}{n-2}$, from the relation we see that

$$R(n+1) = (2(2n-3)/n)\frac{1}{n-1}\binom{2n-4}{n-2} = \frac{2(2n-3)!}{n(n-1)(n-2)!(n-2)!} =$$

$$\frac{2(2n-2)!}{(2n-2)n(n-1)!(n-2)!} = \frac{(2n-2)!}{n(n-1)!(n-1)!} = \frac{1}{n}\binom{2n-2}{n-1}$$

This ends the first proof □

2)

Consider an n-gon P with a distinguished edge the base $\{1, n\}$. In a triangulation of P the base is a side of one of the triangles, say $\{1, k, n\}$, and this triangular region divides P into two disjoint convex polygonal regions $S_1$with k vertices and S_2 with $n-k+1$ vertices for some $k = 2, ...n-1$. See Figure 3.2. The rest of the triangulation of P is complete with some triangulation of the polygons S_1, S_2. Now, S_1 can be triangulated into $R(k)$ ways and S_2 into $R(n-k+1)$ ways. Hence for a given choice of vertex k containing the edge $\{1, n\}$, there are $R(k+1)R(n-k)$ ways of triangulating P thus we have

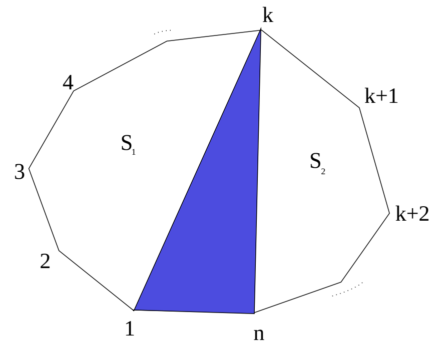

Figure 3.2: Setting up another recursion for $R(n)$.

$$R(n) = R(2)R(n-1)+R(3)R(n-2)+R(4)R(n-3)+\cdots+R(n-1)R(2)$$

.

Now we need to solve this recurrence relation.

We set $R(2) = 1$ and let $a_{n-1} = R(n)$. Now we use the method of formal power series: We have a series $F(x) = \sum_{n=1}^{\infty} a_n x^n$. Now multiplying F by itself we get

$$F^2(x) = (a_1)^2x^2+(a_1a_2+a_2a_1)x^3+\cdots+(a_1a_{n-1}+a_2a_{n-3}+\cdots+a_{n-1}a_1)x^{n-1}+\ldots$$

but by our recurrence relation $a_{n-1} = a_1a_{n-1} + a_2a_{n-3} + \cdots + a_{n-1}a_1$ so we have $F^2 = F - a_1x = F - x$ so the power series is the solution of the quadratic equation $F^2-F+x = 0$. It has two solutions

$$F(x) = (1 + \sqrt{1-4x})/2 \quad \text{or} \quad (1 - \sqrt{1-4x})/2,$$

but we know that $F(0) = 0$ so the right solution is the second one. Finally we need to prove that the formula for $R(n)$ is the right one: Recall Newton's binomial theorem we have

$$(1+z)^{1/2} = \sum_{k=0}^{\infty} \binom{1/2}{k} z^k$$

where the binomial

$$\binom{1/2}{k} = \frac{[1/2(1/2-1)(1/2-2)\dots(1/2-k+1)]}{k!} = \frac{(-1)^{k-1}}{k2^{2k-1}}\binom{2k-2}{k-1}$$

.

Applying this to $z = -4x$ we get the desired result (don't forget to shift again by 1, because $a_n = R(n+1)$).

□

Now we wish to estimate the number of triangulations for an arbitrary point set in the plane.

Theorem 3.2.2. *There is a constant* c *such that every point set in the plane with* n *points has at most* 2^{cn} *triangulations.*

One interesting open problem is to determine the exact constant c that sets the upper bound. How big can c be? The first proof that the number of triangulations of a planar point set is $2^{\Theta(n)}$ is in [1]. Upper bounds of $173\,000^n$, $7\,187.52^n$ and $276.75^{n+O(\log(n))}$ were given respectively in [**?**], [**?**] and [**?**]. The best upper bound known for c is that of F. Santos and R. Seidel []. The precise statement of our upper bound is:

Theorem 3.2.3. *The number of triangulations of* A *is bounded above by*

$$\frac{59^v \cdot 7^b}{\binom{v+b+6}{6}},$$

where v *and* b *denote the numbers of* interior *and* boundary *points of* A, *respectively, meaning by this points of* A *lying in the interior and the boundary of* $\mathrm{conv}(A)$.

As for lower bounds of the value of the constant c, we have already if all the points in A are vertices of its convex hull then the number of triangulations is the well-known Catalan number $\frac{1}{n-1}\binom{2n-4}{n-2} = \Theta(4^n n^{-\frac{3}{2}})$. If A consists of two concave chains of points of the same size facing each other then its number of triangulations is $\Theta(8^n n^{-\frac{7}{2}})$ [32].

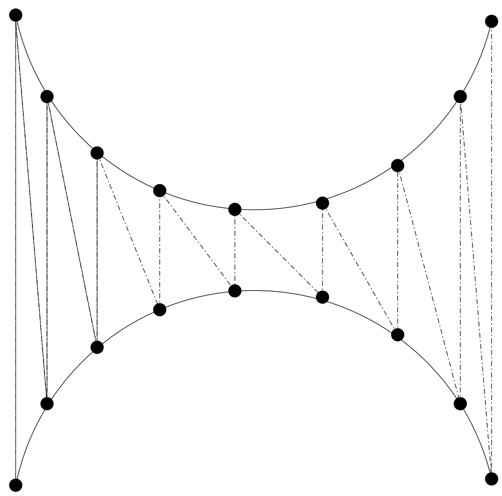

Figure 3.3: A point set with many triangulations (and a partial triangulation of it. Can this construction be improved?

3.4 All Planar Triangulations are Connected by Flips

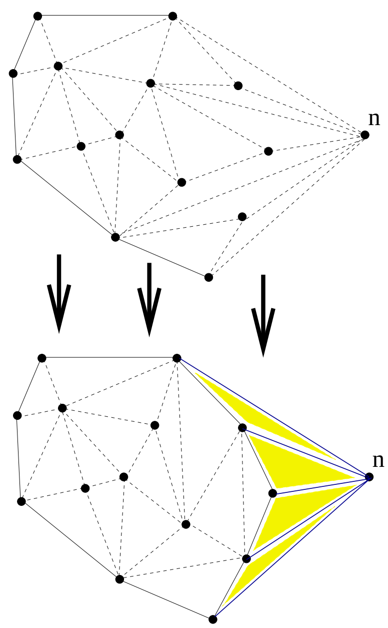

Figure 3.4: The desired transformation for T_i

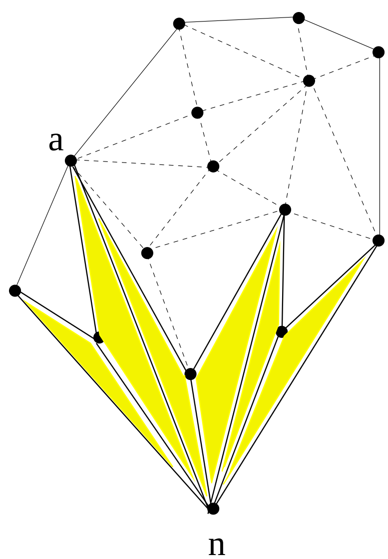

Figure 3.5: The link of n and flippable inner diagonal

We introduced earlier the flip operation between pairs of triangulations. It is a natural question to ask whether any pair of triangulations is connected by a finite sequence of flips. For the longest time we did not know whether this was always the case until Francisco Santos constructed a disconnected example in 1999. It was the end of a long research effort that extended for at least ten years. We will take a close look at Santos' construction, in a simpler version, in Chapter 7. Here we show that in two dimensions the connectivity holds. We show several proofs of this fact, and then illustrate what goes wrong with them already in dimension three!

Theorem 3.4.1. *Every pair of triangulations of a dimensional point configuration is connected by a sequence of finitely many flips.*

Our first proof proceed via induction on n, the number of points in the configuration. It is clear that for three or four points the statement is true. We assume it is true for triangulations of point sets with $n-1$ or less points. Next suppose we have a point set A with n points and two of its triangulations T_1 and T_2. We can label the points of A in such a way that the n-th point is a vertex of the convex hull of A. By induction hypothesis, all the triangulations of $A - \{n\}$ are connected by flips. So our strategy is to slowly transform T_i until it becomes a triangulation T_i' of A that is the union of a triangulation of $A - \{n\}$ and triangles that connect vertex n to its boundary. This is represented in Figure **??**

Note that, by the induction hypothesis T_1', T_2' are indeed connected by flips, once we have transformed T_1, T_2 into the stated shape we will be done. First we can assume without loss of generality that T_i uses all points of A already, else we can flip in the missing points. Now the next step is to look at the link of n (see Figure **??**).

This is a triangulated (not necessarily convex) polygon $P = \text{link}(n, T)$. From now on we start flipping inside P each diagonal that starts at n and belongs to a convex quadrilateral. At each step we reduce the area of P. Since the area is finite we must reach a point when we cannot flip anymore. This means T_i has been taken to the desired shape.

Another proof proceeds in a different way! Given a set of points, $A = \{a_1, \ldots, a_n\}$ in $\mathbb{R}^2$. In a triangulation T of A, an edge $a_i a_j$ is *locally Delaunay* if the edge is in the convex hull or for the triangles $a_i a_j a_k$, $a_i a_j a_l$ that contain the edge, the angles at a_k and a_l add to at most 180 degrees. This is equivalent to say that the circle passing through one of the triangles leaves the other point outside or on the boundary. If T uses all the points of A as vertices and every edge of T is locally Delaunay we say it is a *Delaunay triangulation* of the point set A.

Because of the above definition it is naturally suggested that one can compute the Delaunay triangulation by first computing any triangulation then flipping the edges that are not locally Delaunay until we have no such bad edges left. Can one succeed with this process? For planar triangulations this is indeed possible and it follows from the next lemma. Already for 3-dimensional configurations this process gets stuck in triangulations that are not even regular.

Lemma 3.4.2. *Lift the points A into the paraboloid with equation $z = x^2 + y^2$, by mapping $a_i = (a_{i1}, a_{i2})$ to the three dimensional point $(a_{i1}, a_{i2}, (a_{i1})^2 + (a_{i2})^2)$, then the projection of the lower convex hull of the lifted points decomposes $\mathrm{conv}(A)$ into a union of disjoint convex polygons. A triangulation of A is Delaunay if and only if it triangulates each of these polygons.*

Proof. : If the points a_i, a_j, a_k lie in counterclockwise order along the circle $C(i,j,k)$ they define, then the point a_l is inside the circle $C(i,j,k)$ if and only if

$$\det \begin{bmatrix} a_{i1} & a_{i2} & {a_{i1}}^2 + {a_{i2}}^2 & 1 \\ a_{j1} & a_{j2} & {a_{j1}}^2 + {a_{j2}}^2 & 1 \\ a_{k1} & a_{k2} & {a_{k1}}^2 + {a_{k2}}^2 & 1 \\ a_{l1} & a_{l2} & {a_{l1}}^2 + {a_{l2}}^2 & 1 \end{bmatrix} > 0$$

To see this observe that for any non-vertical plane $z = rx + sy + t$ its intersection with the paraboloid $z = x^2 + y^2$ when projected down into the xy plane is a circle (to see this complete the squares and you get that it is a circle with center at $(r/2, s/2)$ and radius $(r^2/4 + s^2/4 + t)^{1/2})$. With respect to the lifting the point a_l is inside $C(i, j, k)$ if and only if its lifted image $(a_{l1}, a_{l2}, (a_{l1})^2+(a_{l2})^2)$ is below the plane defined by the liftings of a_i, a_j, a_k. Finally the determinant is six times the volume of the tetrahedron with vertices on the four lifted points and it is positive if and only if this is the case. □

We have an important corollary.

Corollary 3.4.3. *Any triangulation of the set of points* A*, that uses all of the points, one can find in a finite sequence of edge flips that transforms it into a Delaunay triangulation. A quadratic number of flips suffice. In particular the graph of triangulations of a planar point set is connected.*

Proof. In the three dimensional lifting picture we have presented above we see that an edge flip can be obtained by gluing a tetrahedron underneath two triangles that share a concave dihedral angle. A sequence of flips generates a tetrahedrization of the polyhedron that lies bounded by the initial triangulation T and the final Delaunay triangulation. □

In this connected graph, given two triangulations, how far apart are they? What is the smallest number of flips I need to go from one to the other?

Theorem 3.4.4. *Any pair of triangulations of a planar point set are at distance no more than the number of edge pair crossings.*

Here is another curious property of the graph of triangulations of a planar point set.

Theorem 3.4.5. *Every triangulation of an* n *point configuration in the plane has at least* $n-3$ *geometric bistellar flips. In other words, the graph of triangulations of a planar point set with* n *points has degree at least* $n-3$ *in all of its vertices.*

Proof. Let T be a triangulation of an n point configuration A in $\mathbb{R}^2$. If there is a flip that inserts a point P, then T can be considered as a triangulation of $A \setminus \{P\}$ and induction on n shows that it has at least other $n-4$ flips. Hence we assume that the triangulation uses all the points of A.

We say that an edge of T is flippable if it is interior (not contained in the boundary of A) and the two triangles incident to it form either a strictly convex quadrilateral or a quadrilateral with two consecutive edges whose union is a straight line segment contained in the boundary of the convex hull of A. In the first case there is a flip of type (2,2) which removes the flippable edge and inserts the other diagonal of the quadrilateral, and in the second case there is a flip of type $(2,1)$ which removes the interior edge and joins the two consecutive collinear edges into one (corresponding to the upper half of Figure 2.16(a)).

Let e_b be the number of boundary edges (note that e_b also equals the number of boundary points of A). Denote by e_i the number of interior edges and by f the number of triangles. Euler's formula for the disk gives $n - e_i - e_b + f = 1$ and a counting argument shows that $3f = 2e_i + e_b$. With these two equalities we obtain:

$$e_i = 3n - 3 - 2e_b$$

For an interior non-flippable edge a, the union of the two triangles sharing a is a quadrangle with a concave or flat vertex which we will call the vertex associated to a. If a vertex p is associated to four interior edges, then the four edges form two pairs of collinear edges with p as a common end and there are two flips of type $(2,1)$ which make p disappear. If p is associated to three interior edges, then the star of p looks like either part (a) or (b) of Figure 2.16, and there is one flip (of type $(2,1)$ or $(3,1)$) which makes the point p disappear.

Hence, the number of interior non-flippable edges is no greater than twice the number of interior points plus the number of flips which make a point disappear. In other words, the total number e_i of interior edges is no greater than the total number of flips plus twice the number

$n - e_b$ of interior points. Thus the number of flips is at least $e_i - 2(n - e_b) = n - 3$, as desired. □