CHAPTER 1

Delaunay triangulations in the plane

1.1 The Voronoi Diagram

Definition: Consider a finite set $S = \{p_1, \ldots, p_n\} \subseteq \mathbb{R}^2$ of n distinct points in the plane. The *Voronoi cell* V_i of $p_i \in S$ is the set of points x that are closer to p_i than to any other points of the set:

$$V_i = \left\{x \in \mathbb{R}^2 \mid \left\|x - p_i\right\| < \left\|x - p_j\right\|, \forall 1 \leq i \leq n, i \neq j\right\}$$

where $\left\|x - y\right\|$ is the euclidian distance between x and y.

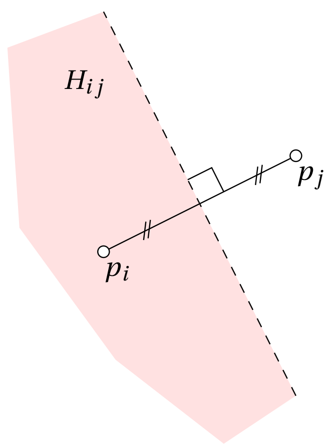

Figure 1.1: Points p_i and p_j, their perpendicular bissector (in dashed lines) and half-plane H_{ij}.

Consider first the case where $S = \{p_i, p_j\}$. The *p*erpendicular bisector of the line segment $p_i p_j$ is a line perpendicular to $p_i p_j$ and passing through its midpoint. The perpendicular bissector of $p_i p_j$ divides $\mathbb{R}^2$ into two halfplanes H_{ij} and H_{ji}:

$$H_{ij} = \{x \in \mathbb{R}^2 \mid \|x - p_i\| < \|x - p_j\|\}.$$

Here, we have clearly $V_i = H_{ij}$. The perpendicular bissector of line segment $p_i p_j$ is the intersection of the closures of the two half planes: $\overline{H_{ij}} \cap \overline{H_{ji}}$.

Let's make the problem a little more complicated and consider a set $S = \{p_i, p_j, p_k\}$ of 3 points. The Voronoi cell associated to p_i is the intersection of half planes H_{ij} and H_{ik}: $V_i = H_{ij} \cap H_{ik}$ (see Figure 1.2).

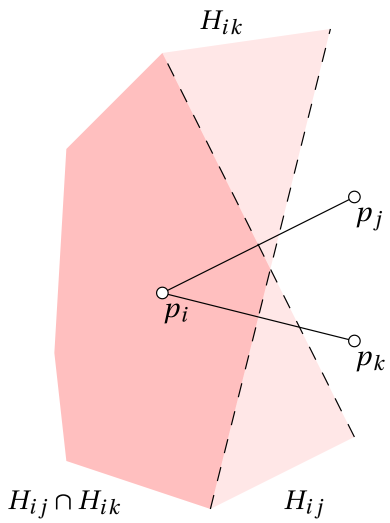

Figure 1.2: Points p_i, p_j and p_k and their perpendicular bissectors.

In the general case, the Voronoi cell relative to p_i is the intersection of all half planes:

$$V_i = \bigcap_{1 \le j \le n, j \ne i} H_{ij}. \tag{1.1}$$

By definition (1.1), each Voronoi cell V_i is the intersection of open half planes containing vertex p_i. The intersection of two convex polygon being itself a convex polygon, V_i is therefore *a convex polygon.*

Definition: *The Voronoi diagram* $V(S)$ is the unique subdivision of the plane into n cells is the union of all Voronoi cells V_p:

$$V = \bigcup_{1 \le i \le n} V_i. \tag{1.2}$$

Each point $x \in \mathbb{R}^2$ having at least one closest point in S, the Voronoi diagram covers the entire plane. Different Voronoi regions are disjoint. Therefore, the Voronoi diagram is unique.

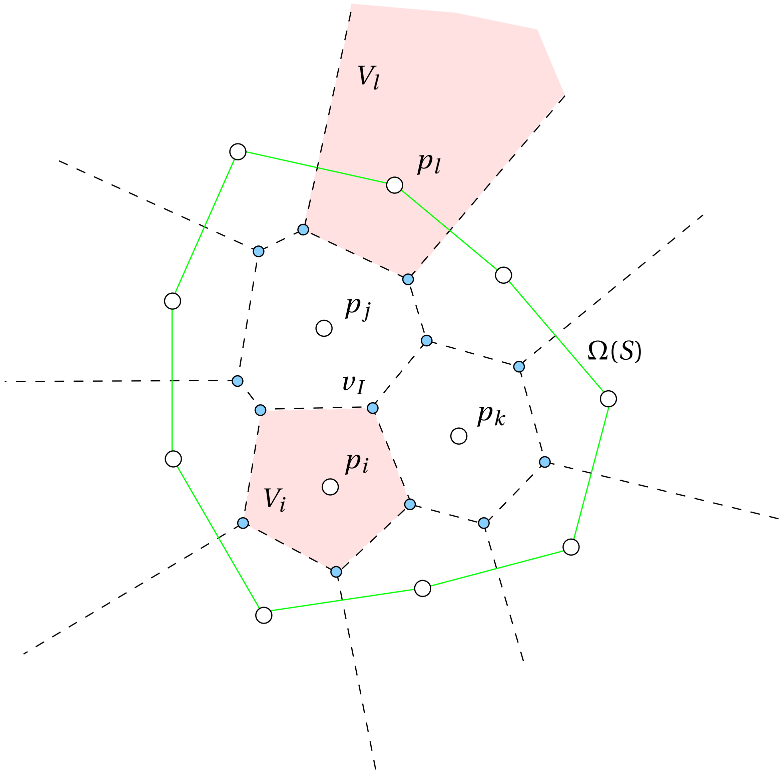

Figure 1.3: Voronoi Diagram. Voronoi cell V_i is closed because it correspond to point p_i that is not on $\Omega(S)$. Voronoi cell V_l is open because $p_l \in \Omega(S)$.

Definition: The *convex hull* $\Omega(S)$ of a finite point set S is the smallest convex polygon that contains S.

Voronoi cells are either closed or open. They can only be open for points (like p_l in see Figure 1.3) that are located on the convex hull $\Omega(S)$ of the point set.

1.2 Triangulations

Definition: A triangulation $T(S)$ of S is a set of non overlapping triangles that exactly covers the convex hull $\Omega(S)$ with all points of S being among the vertices of the triangulation.

Different triangulations of the same point set S may exist (e.g. Figure 1.4), but we are going to show that they all have the same number of edges and of triangles.

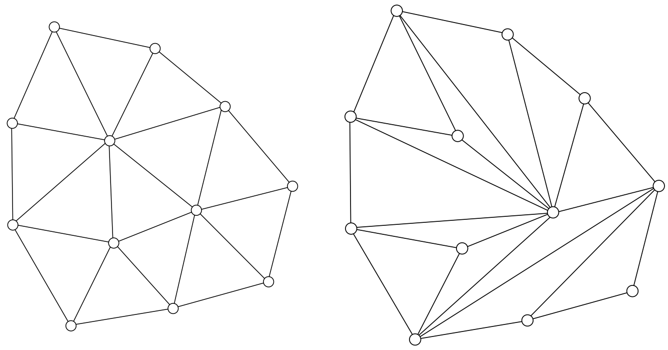

Figure 1.4: Two triangulations of the same point set, both containing $n_f = 13$ triangles defined by a total of $n = 12$ points with $n_h = 9$ points that lie on $\Omega(S)$.

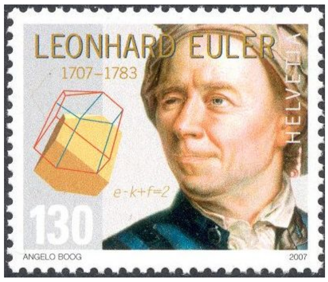

Figure 1.5: Commemorative stamp with Euler and its formula.

Property 1.2.1 *Every triangulation $T(S)$ contains exacltly $n_f = 2(n-1) - n_h$ triangles and $n_e = 3(n-1) - n_h$ edges.*

Proof The proof uses a very well know result of Euler that he proved in 1758. Here is what Euler had to say: Consider any polyhedron and let n be the number of its vertices, n_f the number of its faces, and n_e be the number of its edges. Then

$$n + n_f - n_e = 2. \tag{1.3}$$

A commemorative stamp put out by the Swiss Post shows Euler together with that very famous formula (Figure 1.5). David Eppstein gives 20 different proofs of Euler's formula in [**?**]. Here is one that is quick and elegant. The skeleton of any convex polyhedron is a planar graph. This is geometrically easy to see: in order tu build such a planar graph, dispose the polyhedron on one on a plane and dig a hole on one of its face (an upper face). Then, enlarge this hole in order to unfold the polyhedron up to the point it is completely flattened. The upper face then becomes "infinite" and can be seen as the outer face of the graph. Figure 1.6 shows an example of such a flattening procedure: a cube is shown with its corresponding skelton that is actually

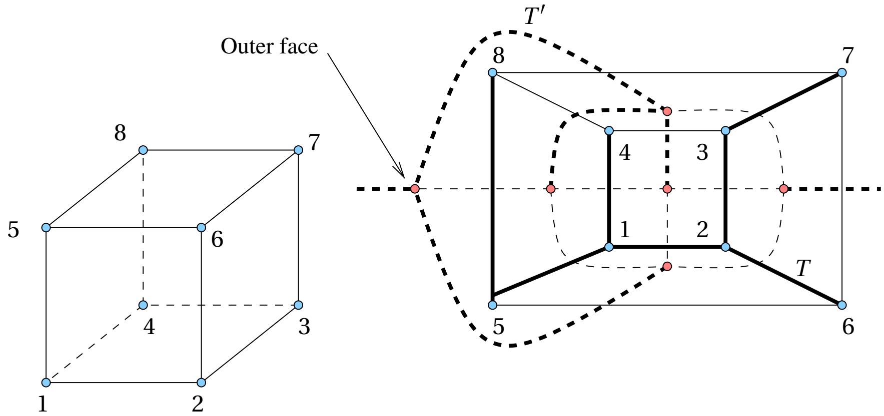

Figure 1.6: A cube (left) and its corresponding skelton Γ (right, plain lines) and the dual Γ' of Γ (right, dashed lines). Edges in bold correspond to a spanning tree T of Γ and edges in bold and dashed correspond to a spanning tree of Γ'.

a planar graph Γ. Let Γ' be the dual of Γ i.e. a graph with its 6 nodes that correspond to the faces of the cubes and its 12 edges that correspond to the edges of the cube. Edges of both graphs have a one-to-one correspondance.

Let T be a spanning tree of Γ i.e. a subgraph $T \subset \Gamma$ that includes all of the vertices of Γ and that is a tree i.e. that contains no cycles. T does not contain any cycles, so it does not disconnect the plane. The co-tree T^* of T is the set of edges of the graph that are not in T. Consider now the set of edges $T' \subset \Gamma'$ that correspond to T^*. Set T' contains no cycles: if one cycle exists in T', then the corresponding edges of Γ would create some isolated vertices in T, which is impossible because T is a spanning tree and it contains all vertices of Γ. T' contains all vertices (the faces of the polyhedron) of Γ' because T does not contain any cycles. Then, T' is a spanning tree of Γ'.

The number of edges on a spanning tree can be computed in a general fashion. Let's construct a spanning tree in the following way: start with one random edge e of Γ and add it to T. This first edge e connects 2 vertices that are inserted in a stack. While this stack is not empty, we take the vertex v at the top of the stack and look for all edges $e_i(v, v_i)$ that are incident to v. We add e_i to T if neither v or v_i is not yet in T. So, each edge of T correspond to one vertex of Γ, except the first one that correspond to two. Then, a spanning tree has exactly $n-1$ vertices.

So T has n vertices and $k \equiv n-1$ edges. Similarly, T' has n_f vertices and $k' \equiv n_f - 1$ edges. Since $k + k' = n_e$, we have $n - 1 + n_f - 1 = n_e$ and formula 1.3 follows.

Euler's formula applies to polyhedron i.e. meshes that are topologically equivalent to a sphere. Euler generalized its formula to general orientable manifolds as

$$n - n_e + n_f = \chi. \tag{1.4}$$

Here, χ is the Euler characteristic is a topological invariant: it is a number that de-

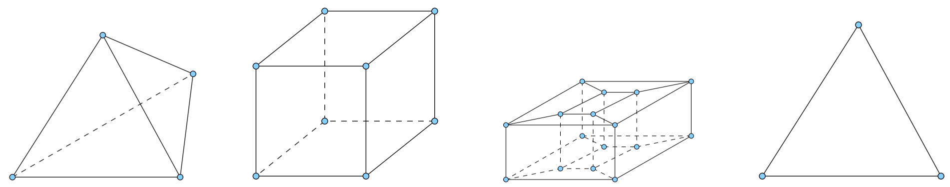

Figure 1.7: Computation of Euler's characteristic χ for different manifold meshes.

scribes the topological structure of the domain. For a sphere, we have obviously $\chi = 2$. In this section, we are essentially concerned by domains that are topologically equivalent to a disk i.e. when $\chi = 1$. Figure 1.7 shows manifold meshes of different objects together with their Euler characteristic.

Now let's specialize Euler's formula to triangulations. In the case of a triangulation of a closed manifold (a sphere or a torus), every edge is is connected to 2 triangles and each triangle has 3 incident edges. We have then

$$2n_e = 3n_f.$$

This last result combined with Eulers's formula gives, for a closed manifold

$$n_f = 2(n - \chi) \text{ and } n_e = 3(n - \chi).$$

For planar triangulations, we consider domains that are topologically equivalent to the a disk i.e. where $\chi = 1$. Those domains have one boundary and every edge of the boundary of the domain of is connected to one single triangle. Assume that n_h is the number of edges (or of points) of the boundary. In this case, $n_e - n_h$ edges are internal with two adjacent triangles and n_h edges and have only one adjacent triangle. Every triangle being always incident to 3 edges, we get the following result

$$3n_f = 2(n_e - n_h) + n_h. \tag{1.5}$$

Combining Euler's formula (1.4) with (1.5) gives the result $n_f = 2(n-1) - n_h$ and $n_e = 3(n-1) - n_h$. ∎

1.3 The Delaunay triangulation

The Delaunay triangulation DT(S) is the geometric dual of the Voroinoi diagram (see Figure 1.8). The Voronoi diagram V is made of n Voronoi cells V_i that correspond to the points p_i, $1 \le i \le n$ of S. The line segments that form the boundaries of Voronoi cells and are the Voronoi edges. Voronoi edges are orthogonal bissectors of neighboring points in the diagram. The endpoints of the Voronoi edges are called Voronoi vertices $v_I, 1 \le I \le N$, N being the number of Voronoi vertices. Voronoi vertices v_I are those points that are equidistant to three or more vertices.

Definition: Points of S are said to be *in general position* if there exist no quadruplet of points of S that are co-circular.

When the points of S are in general position, Voronoi vertices are *triple points* i.e. they are equidistant of three points of S. Consider a Voronoi Vertex v_I that is equidistant to points p_i, p_j and $p_k \in S$ (see Figure 1.3). Voronoi point $v_I = H_{ij} \cap H_{jk} \cap H_{ki}$ is the circumcenter of a triangle $\Delta_I = p_i p_j p_k$.

Definition: The Delaunay triangulation DT(S) is the triangulation of S that consist in the union of the N triangles $\Delta_I, 1 \leq I \leq N$ that correspond to the triple points of the Voronoi diagram (see Figure 1.8).

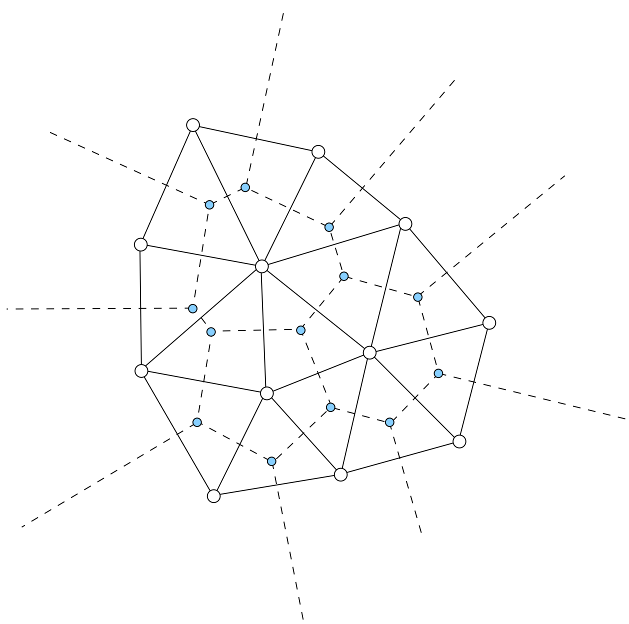

Figure 1.8: Voronoi Diagram (in dashed lines) and Delaunay triangulation. White points are points of S and blue points are Voronoi vertices that are the circumcenters of the triangles.

We should now show that the set of triangles in question is a triangulation in the sense of Definition 1.2. If this is the case, then Property 1.2.1 applies and $N = 2(n-1) - n_h$. The fact that DT is a triangulation will be the consequense of the following properties of the Delaunay triangles. The fact that DT is a triangulation will be the consequence of the two following properties.

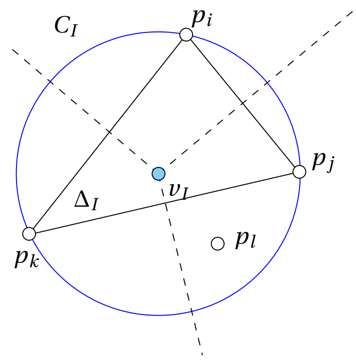

Figure 1.9: Illustration of why the empty circle property is true.

1.3.1 The empty circumcircle property

We will first demonstrate the following remarkable result that is called the empty circumcircle property.

Property 1.3.1 *The empty circumcircle of any triangle in the Delaunay triangulation is empty i.e. it contains no point of S.*

Proof Consider the Delaunay triangle $\Delta_I = p_i p_j p_k$ (see Figure 1.9). Assume now that point $p_l \in C_I$ where C_I is the circumcircle of Δ_I. By definition, the triple point v_I is at equal distance to p_i, p_j and p_k and no other points of S are closer to v_I than those three points. Then, if a point like p_l exist in S, v_I is not a triple point and triangle Δ_I cannot be a Delaunay triangle.

1.3.2 Delaunay Edges

It is useful at that point to look at some geometrical properties of circle bundles that share two points p_i and p_j. The centers of such circles lie on the perpendicular bissectors of line segment $p_i p_j$ (see Figure 1.10). Edge $p_i p_j$ divides disk C_1 into two disk sectors and one of the two sectors completely lies inside C_2. On the Figure, the pink sector of C_1 is inside C_2 and the yellow sector of C_2 lies inside C_1.

Definition: An edge $p_i p_j$ of a triangulation is a *Delaunay edge* if there exist a circle that contains p_i and p_j and that is empty i.e. that contain no point of S.

Property 1.3.2 *A mesh is a Delaunay Triangulation if and only if all its edges are Delaunay edges.*

Proof Let us first show that a Delaunay triangulation has only Delaunay edges. Assume a Delaunay triangulation $T(S)$ and an edge $p_i p_j$ that is not Delaunay. This means that there exist no circle passing through p_i and p_j that is empty. Consider

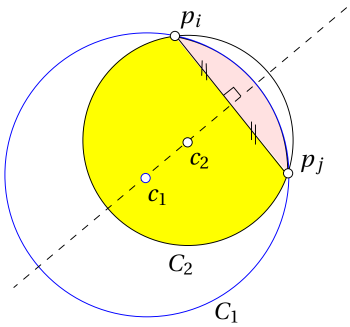

Figure 1.10: Two circles C_1 and C_2 sharing an edge $p_i p_j$. The centers of the circles c_1 and c_2 lie on the perpendicular bissector of segment $p_i p_j$ (in dashed lines).

Delaunay triangle $\Delta_I = p_i p_j p_k$ that contains edge $p_i p_j$. Its circumcircle is empty by definition because T is a Delaunay triangulation. This is in contradiction with the hypothesis that there exist no circle passing through p_i and p_j and that is not empty.

Now let's proof that if every edge of a triangulation is Delaunay, then every triangle is Delaunay as well. Assume that triangle $\Delta_I = p_i p_j p_k$ is not Delaunay, but all its 3 edges $p_i p_j$, $p_i p_k$ and $p_j p_k$ are Delaunay. Figure 1.11 shows a configuration whith a non Delaunay triangle $\Delta_I = p_i p_j p_k$ which circumcircle contains p_l. Because we deal with triangulations as defined in Definition 1.2, p_l cannot be inside triangle Δ_I. It is then situated inside one of the three circular sectors delimited by p_i, p_j and p_k. Assume that p_l and p_j are on opposite sides of $p_i p_k$ like in Figure 1.11. By hypothesis, there exist a circle passing through p_i and p_k and that is empty. The center of such a circle lies on the orthogonal bissector of $p_i p_k$. Any circle like C_1 with its center c_1 that is below c_I contains p_j any circle C_2 that is above c_I contains p_l, which is in contradiction with the hypothesis that there exist a circle passing through $p_i p_k$ and that is empty.

1.3.3 Local Delaunayhood

Definition: Given a triangulation $T(S)$ and an edge $p_i p_j$ in the triangulation that is adjacent to two triangles $\Delta_I = p_i p_j p_k$ and $\Delta_J = p_i p_l p_j$. We call edge $p_i p_j$ *locally Delaunay* if p_l lies on or outside the circumcircle of Δ_I.

Figure 1.12 gives an illustration of an edge $p_i p_j$ that is not locally Delaunay: point p_l lies inside circle C_I. It is easy to see that this condition is symmetric: if point p_l lies inside circle C_I, then point p_k lies inside circle C_J. We'll prove that below.

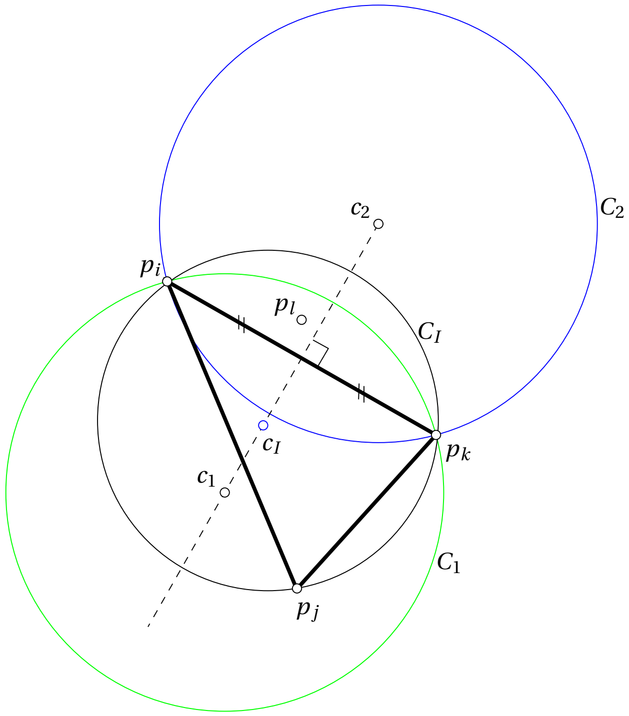

Figure 1.11: Two circles C_1 and C_2 sharing an edge $p_i p_j$. The centers of the circles c_1 and c_2 lie on the perpendicular bissector of segment $p_i p_j$.

1.3.4 Edge Flip

Consider again the situation of two triangles adjacent to edge $p_i p_j$ as depicted in Figure 1.12. Flipping edge $p_i p_j$ consist in replacing triangles $p_i p_j p_k$ and $p_j p_i p_l$ by triangles $p_l p_k p_i$ and $p_k p_l p_j$. Edge $p_i p_j$ has been flipped and replaced by edge $p_k p_l$.

The edge flip operator can only be applied to a pair of triangles that form a convex quadrilateral. If it is concave, then flipping the edge leads to an invalid configuration with two overlapping triangles (see Figure 1.13).

Property 1.3.3 *An edge that is not locally Delaunay is flippable and the new edge resulting of the flip operation is locally Delaunay.*

Proof Let us first shwow that any edge that is not locally Delaunay is flippable. Consider Figure 1.12. Edge $p_i p_j$ is not locally Delaunay because $p_k \in C_J$ and $p_l \in C_I$. A simple way of checking wether edges $p_i p_j$ and $p_k p_l$ can be flipped is to verify that they actually intersect. Consider triangle $p_j p_k p_i$ on Figure 1.12. The fact that p_l is on the opposite side of $p_i p_j$ than p_k and that it lies inside C_I ensures that $p_k p_l$ in-

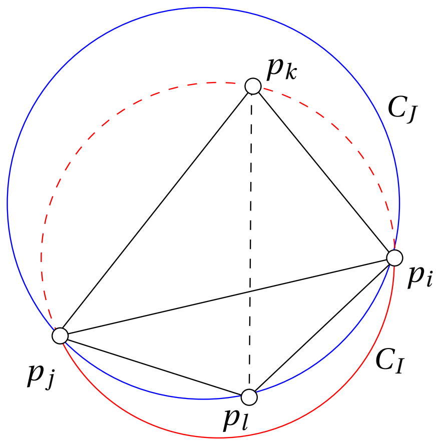

Figure 1.12: An edge $p_i p_j$ that is not locally Delaunay.

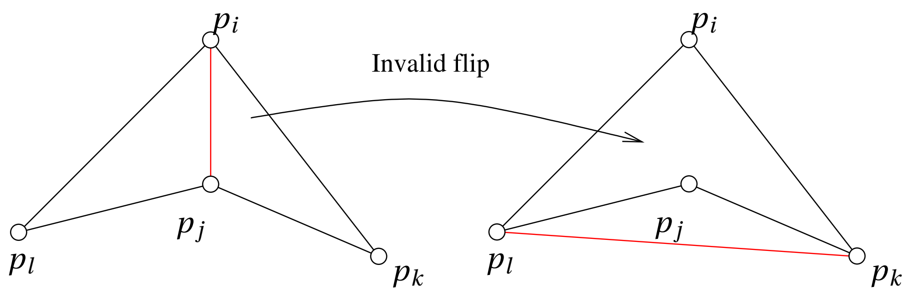

Figure 1.13: Invalid edge flip configurations.

tersects $p_i p_j$ which proves that edge $p_i p_j$ is flippable if $p_i p_j$ is not locally Delaunay. Move now to Figure 1.14. and prove that, if $p_i p_j$ is not locally Delaunay, then $p_k p_l$ is locally Delaunay. In other words, we'd like to prove that, provided that p_l is inside C, then p_i is outside C'.

Circles C and C' share edge $p_k p_j$ and points p_i and p_l are on the same sides of edge $p_k p_j$. Edge $p_i p_j$ is not Delaunay by hypothesis. Then point p_l is inside C, as well as the whole arc $\widehat{p_k p_l p_j}$ (in dashed line on Figure 1.14) of C'. Point p_i belongs to C and is on the same side of $p_k p_j$ as p_l, it is then outside C' and edge $p_k p_l$ is locally delaunay. ∎

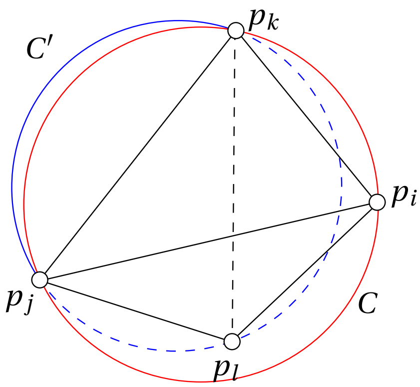

Figure 1.14: If $p_i p_j$ is not locally Delaunay, then $p_k p_l$ is locally Delaunay.

1.3.5 Locally Delaunay vs. Globally Delaunay

Property 1.3.4 *If all edges of triangulation $T(S)$ are locally Delaunay, then T is the Delaunay triangulation $DT(S)$.*

The fact that a specific edge is locally Delaunay does not imply that both its two adjacent triangles are Delaunay triangles. Yet, if all edges are locally Delaunay, then the resulting triangulation is Delaunay.

Proof We prove property 1.3.4 by contradiction. Assume all edges of a triangulation to be locally Delaunay. Assume that triangle $\Delta_I = p_i p_j p_k$ has its circumcircle C_I that contains point $p_l \in S$. The situation is summarized on Figure 1.15. Assume that point p_l and p_i are on opposite sides of $p_j p_k$. Edge $p_j p_k$ is locally Delaunay but triangle $p_i p_j p_k$ is not Delaunay because its circumcircle is not empty (it actually contains point p_l). Consider triangle $p_k p_j p_m$. Points p_i and p_m are on opposite sides of $p_j p_k$ and and p_m is outside C_I. This implies that C_J contains p_l as well. We can continue that and show that C_K and C_L both contain p_l as well. Yet, edge $p_0 p_n$ is supposed to be locally Delaunay which means that p_l should be outside C_L. This is indeed a contradiction.

1.3.6 The Flip Algorithm

Result 1.3.4 is of high importance. Combined with the flip algorithm, we can forsee a simple algorithm that would start with any triangulation $T(S)$ and would produce the Delaunay triangulation DT(S) using edge successive flips. The algorithm could be summarized as follows

- Insert all the internal edges of $T(S)$ in a `stack`.
- Do while the `stack` is not empty

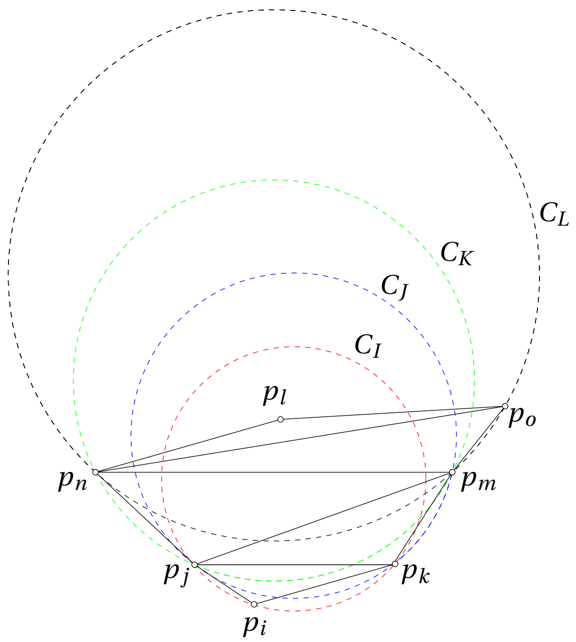

Figure 1.15: An edge $p_i p_j$ that is locally Delaunay (point p_m is outside C_I) but with triangle $p_i p_j p_k$ that is not Delaunay.

- Take edge $p_i p_j$ at the top of the `stack`. This edge is adjacent to triangles $p_i p_j p_k$ and $p_j p_i p_l$. If $p_i p_j$ is not locally Delaunay, then flip it and add edges $p_i p_k$,$p_k p_j$, $p_j p_l$ and $p_l p_i$ in the stack. If one of those edges was already present in the stack, update its neighbors.
- Remove $p_i p_j$ from the `stack`.

Two questions should be asked at that point: (i) does this algorithm produce the Delaunay triangulation of S and (ii) if it achieves to create DT(S), what is its complexity?

Proposition 1.3.1 *The edge flip algorithm converges to DT(S) in at most $\mathcal{O}(n^2)$ flips.*

Proof Consider an edge $p_i p_j$ that is not Delaunay (Figure 1.16) with its two adjacent triangles $p_i p_j p_k$ and $p_j p_i p_l$ and their respective circumcircles C_I and C_J, with $p_l \in C'_J$ and $p_k \in C'_I$. Edge flip will produce triangles $p_j p_k p_l$ and $p_i p_l p_k$ and their respective circumcircles C'_I and C'_J. Edge $p_k p_l$ is locally Delaunay i.e. $p_i \notin C'_I$ and $p_j \notin C'_J$.

Consider now the set of all possible point-triangle relations in a mesh T and a function $F(T)$ that counts how many of those relations violate the Delaunay empty circle property. There is at most $\mathcal{O}(n^2)$ point-triangle pairs in a mesh (see property 1.2.1). So, F's magnitude is not bigger than $\mathcal{O}(n^2)$. Assume now that edge $p_i p_j$ is

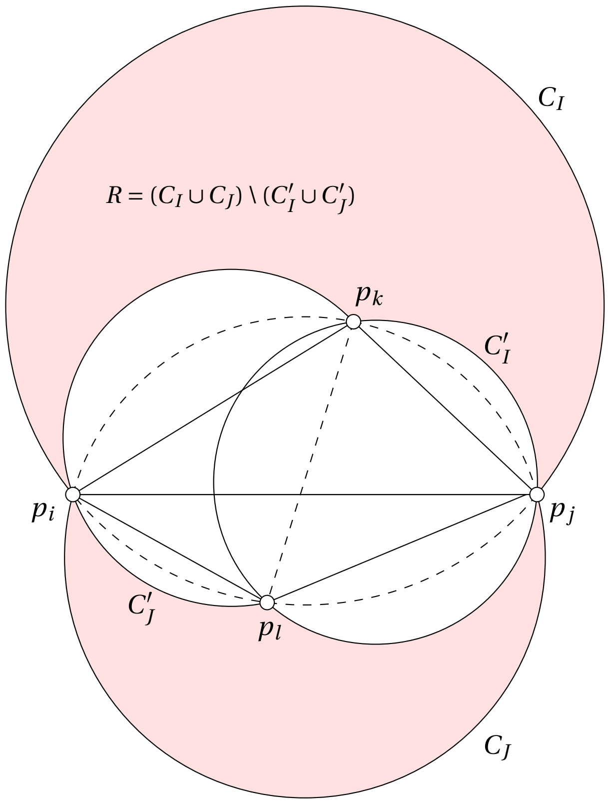

Figure 1.16: Edge flip: $C_I \cup C_J \subset C'_I \cup C'_J$.

flipped, leading to a new triangulation T'. Flipping an edge always leads to $F(T') < F(T)$. Figure 1.16 shows visually that

$$C_I \cup C_J \subset C'_I \cup C'_J,$$

the colored zone in the Figure representing

$$R = (C_I \cup C_J) \setminus (C'_I \cup C'_J).$$

If some points of S were inside circumcenters of triangles $p_i p_j p_k$ and $p_j p_i p_l$ in T, then edge flip will not increase that number because those points will not be anymore invalid. If R contains no points of S, then $F(T') = F(T) - 2$ because the two point-triangle relations associated to points p_l and p_k and triangles $p_i p_j p_k$ and $p_j p_i p_l$ disappear from F. In conclusion, we have

$$F(T') \leq F(T) + 2$$

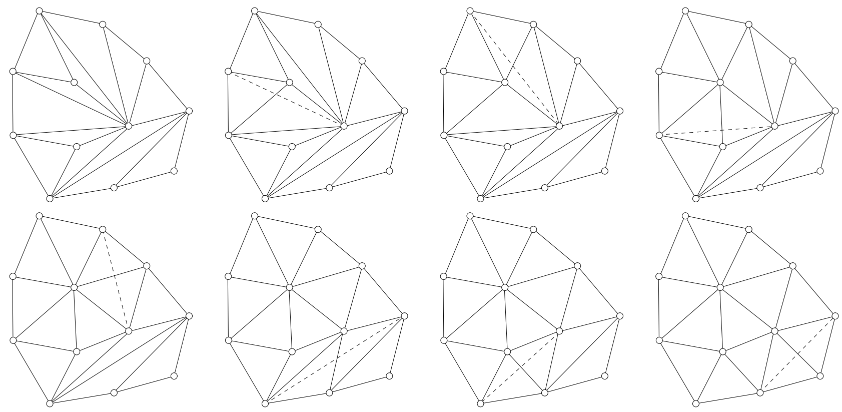

Figure 1.17: Building an angle-optimal triangulation using swaps.

which means that F decreases at each edge flip.

F is bounded by above by $\mathcal{O}(n^2)$. It is also bounded by below: only the Delaunay triangulation has empty circumcircles, $F(\mathrm{DT}) = 0$. The edge flip algorithm converges to the Delaunay triangulation and its complexity is $\mathcal{O}(n^2)$ in the worst case. ∎

This result is outmost importance. It means that every triangulation $T(S)$ is connected to the Delaunay triangulation $\mathrm{DT}(S)$ by at most $\mathcal{O}(n^2)$ flips. It also means that any two triangulations T and T' are flip connected. Both T and T' being connected to DT, it is therefore possible to go from T to DT using flips and then from DT to T' using "back flipping". The flip-connectness of 2D triangulations allows to generate meshes of arbitrary domains with low complexity. This will be developped in further chapters. Figure 1.17 illustrate the edge flip procedure.

1.3.7 The MaxMin property

Let us first recall a very old geometry theorem from Thales.

Proposition 1.3.2 *Let C_A and C_B be two circumcircles of edge $p_i p_j$ (see Figure 1.18). Let b_1 and b_2 be two points on C_B on the same side of $p_i p_j$. Then, b_1 and b_2 see the edge $p_i p_j$ with the same angle β. Consider now point a on the same side of $p_i p_j$ as b_1 and b_2 but on circle C_A. Assume that b_1, b_2 are inside C_A. Then, $\alpha < \beta$.*

Consider a triangulation $T(S)$ with n_f triangles. This triangulation has $3n_f$ internal angles (3 angles per triangle). Consider the vector of angles $A(T) = (\alpha_1, \ldots, \alpha_{3n_f})$ sorted by increasing values. We can define such a vector for any triangulation. Each triangulation $T(S)$ has the same number of triangles so each vector $A(T)$ has the same length and it is therefore possible to compare them, e.g. lexicographically. We say that one given triangulation T is angle-optimal if $A(T) \leq A(T'), \forall T'$.

Property 1.3.5 *The Delaunay triangulation $\mathrm{DT}(S)$ is angle-optimal: it maximizes the minimum angle among all possible triangulations.*

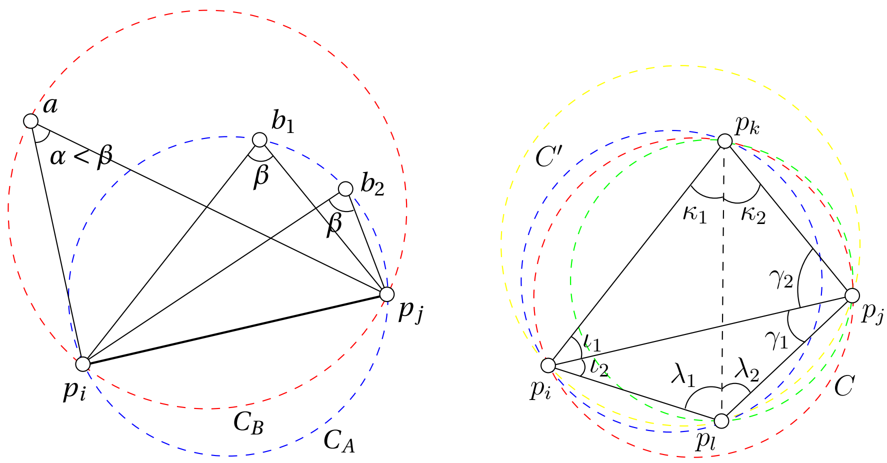

Figure 1.18: Thales theorem (left) and MaxMin property illustrated (right)

Proof Consider two triangulations T and T', where T' differs from T by one edge flip. Let us proove that $A(T') \leq A(T)$. The edge flip procedure consist in replacing triangles $p_i p_j p_k$ and $p_j p_l p_k$ by triangles $p_k p_l p_i$ (see Figure 1.18). The angles of the old configuration are respectively

$$\kappa_1 + \kappa_2, \gamma_2, \iota_1, \iota_2, \gamma_1 \text{ and } \lambda_1 + \lambda_2.$$

The angles of the old configuration are respectively

$$\iota_1 + \iota_2, \kappa_1, \lambda_1, \kappa_2, \lambda_2 \text{ and } \gamma_1 + \gamma_2.$$

Our aim is to bound by above all angles of the old configuration. Two of the 6 relations are obvious: $\gamma_1, \gamma_2 < \gamma_1 + \gamma_2$ and $\iota_1, \iota_1 < \iota_1 + \iota_2$. We use Thales Theorem 1.3.2 for the last four ones. Thales Theorem applied respectively to segments $p_i p_l$ (blue and yellow circles), $p_j p_k$ (red and green circles), $p_i p_k$ (blue and red circles) and $p_l p_j$ (yellow and green circles) gives

$$\gamma_1 < \kappa_1, \;\; \iota_1 < \lambda_2, \;\; \gamma_2 < \lambda_1 \text{ and } \gamma_1 < \kappa_1$$

which are the four relations that were needed. Successive edge flips lead to the Delaunay triangulation and each flip does not increase the minimum angle. The Delaunay triangulation is therefore angle-optimal. ∎