Planar triangulations

Žiga Virk

October 7, 2021

In the previous lecture we learned about metric spaces along with the homeomorphic and homotopic type. However, the descriptions we used are not of combinatorial nature, and one would have difficulties using them for computations. In this lecture we will introduce one of the simplest combinatorial descriptions of planar spaces: triangulations in the plane. Essentially, we would like to describe a planar region as a “nice” union of triangles. The triangles are used primerely because they are easy to describe: we only have to provide the three points. In later sections we will use these triangulations to compute various invariants of the space: components, homology, etc.

It turns out that not every planar subset can be triangulated. However, finite triangulations (i.e., triangulations with finitely many triangles) can be obtained for most planar subsets of interest.

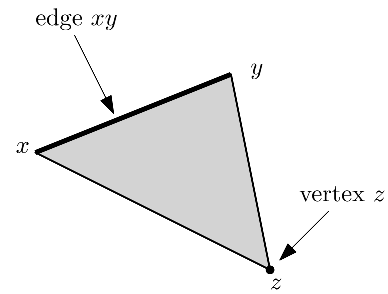

Figure 1: Triangle xyz.

1 Definition of planar triangulations

A triangle in the plane has three edges and three vertices.

Definition 1.1. *A* ***triangulation*** *of a closed region $D \subset \mathbb{R}^2$ is a decomposition of D into triangles, so that:*

1. *no triangle is degenerate (i.e., a point or just a line segment),*
2. *interiors of triangles are disjoint, and*
3. *intersection of any pair of triangles is either a common edge, a common vertex, or empty.*

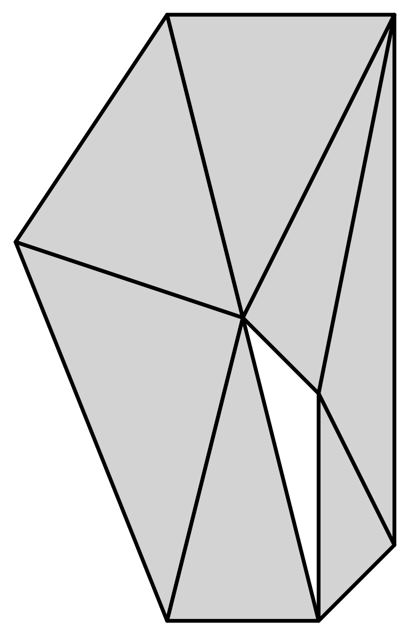

Figure 2: A planar triangulation.

Geometric description of Definition 1.1 is provided by Figure 3.

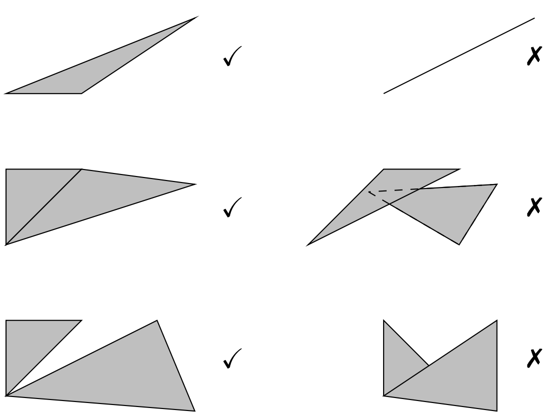

Figure 3: Conditions of Definition 1.1.

The idea of a triangulation may be generalized in various ways. One could use different shapes of pieces to decompose a planar region or the entire plane. Such decompositions are called tessellations. Generalizing by dimension, one could use "higher dimensional triangles", such as tetrahedra, to decompose a higher dimensional space. This idea will be formalized as simplicial complex in the next lecture.

☞ A planar region admitting a triangulation is called a polygonal region.

Modifications of triangulations

Occasionally we will want to modify a triangulation. Here are some of the most used modifications:

- add a triangle;
- remove a triangle;
- flip a common edge;
- refine using the barycentric subdivision: for each edge and each triangle consider its geometric center (centroid) as a new vertex in our triangulation, and then decompose each triangle as demonstrated by Figure 4. This modification is convenient when we want to refine a triangulation, i.e., systematically decompose the triangles into smaller triangles.

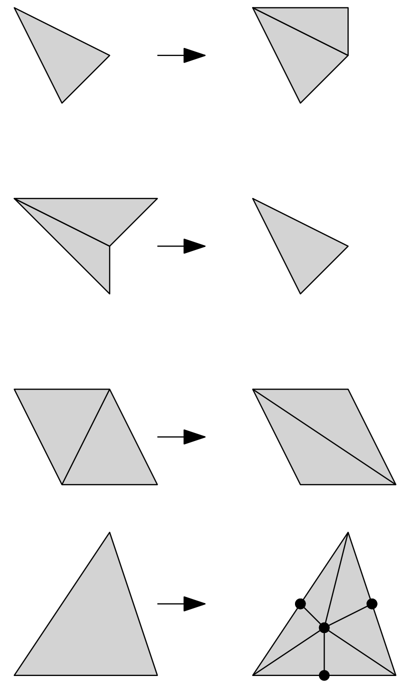

Figure 4: Modifications of tirangulations.

At this time we will focus on triangulations of convex polyhedra, i.e., convex hulls of finitely many points in the plane, as defined below. Given a finite $S \subset \mathbb{R}^2$ we say a **triangulation on** S is any triangulation of the convex hull of S, whose vertex set is S.

2 Recap on convexity

Given points $x, y \in \mathbb{R}^n$, the line segment between them is parameterized as

$$\gamma(t) = tx + (1-t)y, \quad t \in [0,1].$$

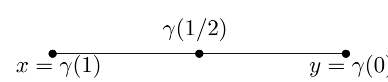

Figure 5: Line segment.

Note that $\gamma(0) = y, \gamma(1) = x$, and $\gamma(1/2)$ corresponds to the midpoint of the line segment.

Definition 2.1. *A subset $A \subset \mathbb{R}^n$ is **convex**, if for each $a, b \in A$ the entire line segment between a and b lies in A, i.e., if $\forall t \in [0,1]$ we have $ta + (1-t)b \in A$.*

*Given a subset $B \subset \mathbb{R}^n$, its **convex hull** $\mathrm{Conv}(B)$ is the smallest convex set containing B.*

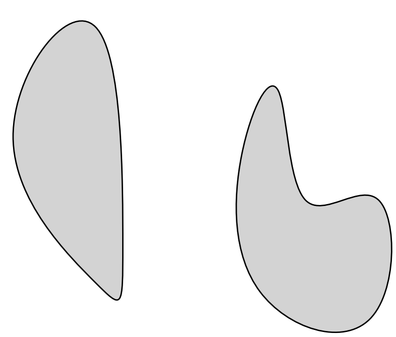

Figure 6: A convex (left) and a non-convex (right) subset of the plane.

The closed region on Figure 2 is not convex, while the ones on Figure 4 are convex. A triangle is the convex hull of the set of its vertices, which provides a convenient description of a triangle: the triangle with affinely independent vertices $x, y, z \in \mathbb{R}^2$ can be parameterized by all possible convex combinations of these vertices:

$$\{\alpha_1 x + \alpha_2 y + \alpha_3 z, \quad | \quad \forall i \in \{1,2,3\} : \alpha_i \in [0,1], \quad \sum_{i=1}^{3} \alpha_i = 1\}.$$

The term "convex combination" (as opposed to the linear combination) refers to the fact that the coefficients α_i are from $[0,1]$ and add up to 1. These coefficients are called **barycentric coordinates** in a triangle. The point with $\alpha_1 = \alpha_2 = \alpha_3 = 1/3$ is the centroid of the triangle, while points with two barycentric coordinates[1] $1/2$ are the midpoints of the corresponding edges; all these points are used in the barycentric subdivision shown in Figure 3.

[1] ...and the third coordinate equal to 0.

Convex hull can be constructed by iteratively adding all feasible line segments. It is important to note that for $B \subset \mathbb{R}^n$ the set obtained by adding all line segments

$$\{\alpha_1 x + \alpha_2 y \quad | \quad x, y \in B, \quad \forall i : \alpha_i \in [0,1], \quad \sum_{i=1}^{2} \alpha_i = 1\}$$

is typically not the convex hull. For example, starting with three vertices and adding the line segments between all three pairs we would obtain the set consisting of the edges but not the interior of the triangle (which constituted the convex hull of three points). Instead, we have to apply the procedure of adding all possible line segments again and again, or alternatively, add all convex combinations in one step:

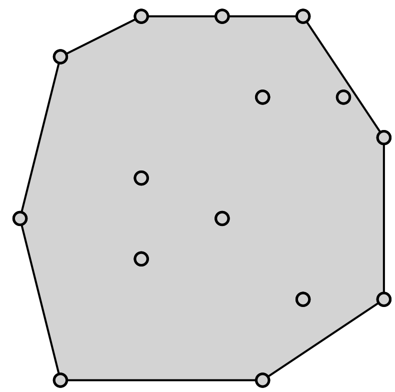

Figure 7: A collection of points and its convex hull.

$$\mathrm{Conv}\, B = \bigcup_{m \in \mathbb{N}} \{\sum_{i=1}^{m} \alpha_i x_i \quad | \forall i : \quad x_i \in B, \quad \alpha_i \in [0,1], \quad \sum_{i=1}^{m} \alpha_i = 1\}.$$

By the Carathéodory Theorem we can bound the number of summands by the dimension of the ambient space plus one:

$$\mathrm{Conv}\, B = \{\sum_{i=1}^{n+1} \alpha_i x_i \quad | \forall i : \quad x_i \in B, \quad \alpha_i \in [0,1], \quad \sum_{i=1}^{n+1} \alpha_i = 1\}.$$

In particular: for a finite subset $F \subset \mathbb{R}^2$, each point of $\mathrm{Conv}(F)$ is contained in the convex hull of some triple of points from F.

3 Euler characteristic

Along with the number of components of a space, the Euler characteristic is one of the first real topological invariants we come across. In particular, while there are many triangulations of $\mathrm{Conv}(S)$ on a finite subset $S \subset \mathbb{R}^2$, the Euler characteristic is the same for all of them.

For a given triangulation let:

- V be the number of its vertices,
- E be the number of its edges,
- F be the number of its triangles.

Definition 3.1. *The* ***Euler characteristic*** χ *of a triangulation is defined as*
$$\chi = F - E + V.$$

Theorem 3.2 (A simple version of the Euler-Poincaré formula)**.** *Let* $S \subset \mathbb{R}^2$ *be finite. For each triangulation on* S *we have* $\chi = 1$.

Proof. Let us assume our triangulation has no vertical edge: if necessary this can be achieved by a small rotation. Assign to each triangle value $+1$ and to each edge value -1. Slide each of these values towards the unique rightmost vertex of the corresponding triangle/edge as suggested by Figure 9. Assign to each vertex value $+1$. The total sum of all assigned values is χ.

For each single vertex add: the value at the vertex and all the values of the triangles and edges, that gathered at that vertex. We can see that for each vertex the total sum is zero (arising from a sequence edge-triangle-edge-...-edge-triangle-edge on the left from the vertex plus the vertex itself) except for the leftmost vertex, where the value equals one.

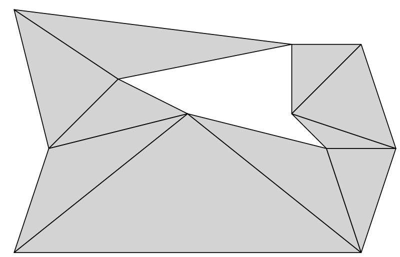

$\chi = 0$

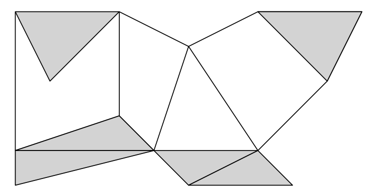

$\chi = -3$

Figure 8: A few planar triangulations along with their Euler characteristics.

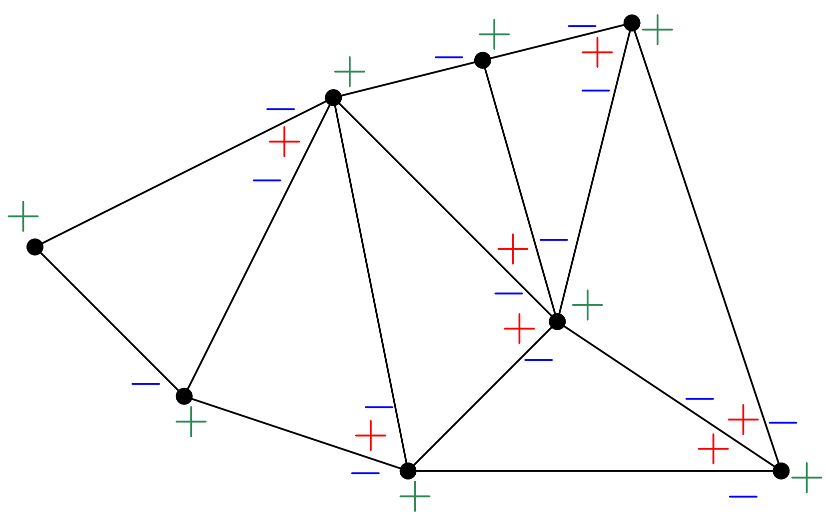

Figure 9: The assignment of values on triangles (red $+$), edges (blue $-$) and vertices (green $+$) from the proof of Theorem 3.2. Vertices also hold additional $+1$ value. The triangles are present but not shaded.

$\square$

Remark 3.3. *It turns out that $\chi = 1$ for each triangulation of a contractible set*[2] *in $\mathbb{R}^2$. In fact, for a triangulation of $D \subset \mathbb{R}^2$, χ equals the number of components minus the number of holes of D. More technical details on this fact will be provided in later sections. Let us just mention that the number of holes of a bounded set $D \subset \mathbb{R}^2$ equals the number of components of $\mathbb{R}^2 \setminus D$ minus one (see Figure 10).*

[2] Even more: χ is homotopy invariant.

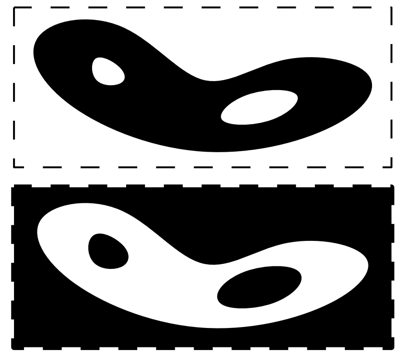

Figure 10: The top set has **2** holes. Equivalently, its complement on the bottom has $\mathbf{2}+1$ components.

4 Constructing planar triangulations with line sweep

Let $S \subset \mathbb{R}^2$ be finite. Perhaps the simplest way to construct a triangulation on S is using a **line sweep**, which we now describe. Assume no two points of S have the same x-coordinate (this can be achieved by a small rotation if necessary). Now imagine a vertical line sweeping $\mathsf{Conv}(S)$ from left to right. Each time the line reaches a point of S (a vertex in our triangulation), add all possible edges towards left without creating intersections. Furthermore, for each new bounded region add the corresponding triangle. As the line sweeps S we thus obtain a triangulation on S.

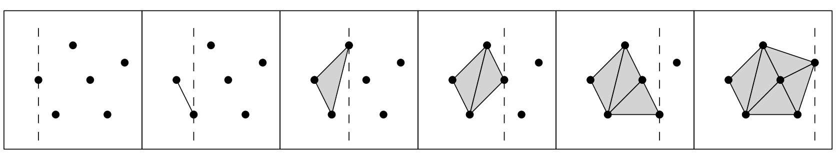

Figure 11: A line sweep using the vertical dashed line. Each time the vertical line reaches a point, we add all possible edges from that point to a point with smaller x-coordinate.

The condition that no two points have the same x-coordinate was added for reasons of simplicity only. If more points, say $a_1, a_1, \ldots, a_k$ have the same x-coordinate then, instead of adding all edges for all points a_i at once, proceed point by point: add all possible edges for a_1, then for a_2, etc. Depending on the order of points a_i we typically get a different triangulation.

It should also be obvious that the line sweep does not need to proceed from left to right, but can proceed along any direction by sweeping a line perpendicular to that direction.

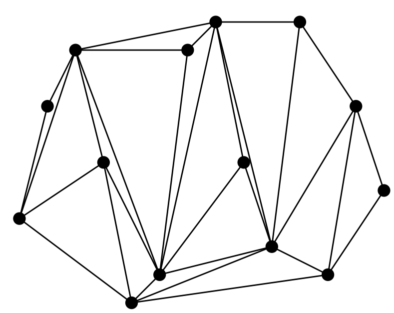

Figure 12: A line sweep triangulation resulting in thin triangles.

While the line sweep is conceptually simple, it does tend to construct triangulations with very thin triangles, which may be undesirable in applications. The triangulation that avoids thin triangles as much as possible is the Delaunay triangulation.

5 Voronoi diagram and Delaunay triangulation

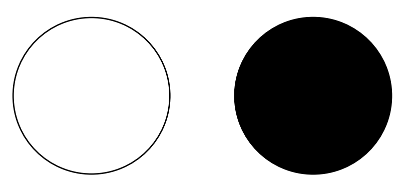

Figure 13: Circle on left, and a ball on right. The boundary of the ball is the circle.

Throughout this section let $S \subset \mathbb{R}^2$ be a finite subset satisfying a general position property: no four points of S lie on the same circle. We will first present the Voronoi diagram of S, which is a decomposition of the plane into specific regions.

For each $s \in S$ define the **Voronoi region** of s:

$$V_s = \{x \in \mathbb{R}^2 \mid \forall u \in S \setminus \{s\} : d(x,s) \leq d(x,u)\}.$$

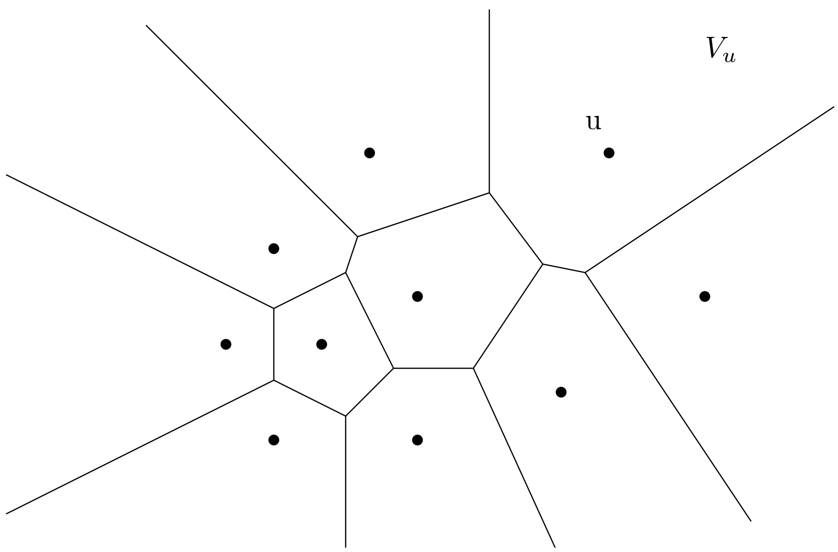

Figure 14: An example of a Voronoi decomposition.

If a pair of Voronoi regions V_{s_1}, V_{s_2} has a non-empty intersection, then (due to the general position condition above) this intersection is a bounded or unbounded line segment called a **Voronoi edge** and lies on the bisector between s_1 and s_2.

If a triple of Voronoi regions $V_{s_1}, V_{s_2}, V_{s_3}$ has a non-empty intersection, then this intersection is a point called a **Voronoi vertex**. As this point lies on all three pairwise bisectors, it is the center of the circle containing s_1, s_2 and s_3.

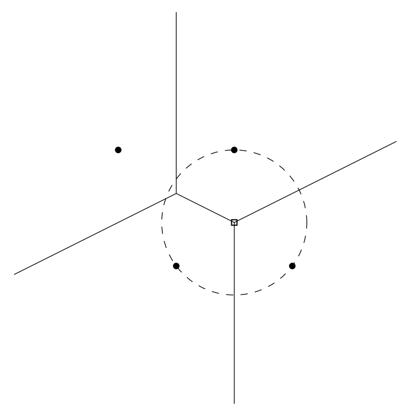

Figure 15: A Voronoi vertex $\square$ is the center of the circle containing the corresponding points $\bullet$ of S. Voronoi edges lie on the bisector lines between the corresponding points of S.

Definition 5.1. *The* ***Voronoi diagram*** *(or decomposition) of S is the collection of Voronoi regions, edges and vertices.*

A Voronoi region V_s consists of points, whose closest point of S is s. If for some point w there are two such closest points in S, then w is on the corresponding edge. If for some point w there are three such closest points in S, then w is a Voronoi vertex. The general position criterion above states that for each point in the plane, there can be no four closest points in S.

Voronoi diagram can be thought of as a result of a uniform expansion from the points of S. Suppose that in the initial stage we start with a finite set of locations S. Then, as time goes by, each point of S is being expanded into a region by growing at the same speed in all directions. At the beginning all these regions are balls centered at the points of S. As the growing regions collide, the growth towards the regions (edges) of contact stops, but continues along all other directions. The Voronoi decomposition is the final result of such growth, with each Voronoi region V_s containing the points that were first reached from s.

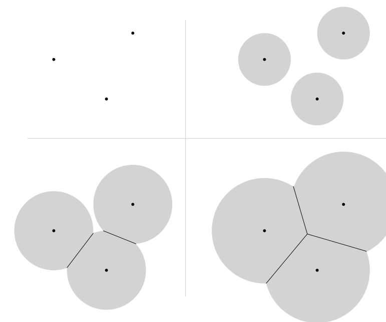

Figure 16: Voronoi diagram arising from expansion around points.

Definition 5.2. *The **Delaunay triangulation** on S, denoted by $\mathcal{D}(S)$, is the triangulation on S, such that:*

- *its vertices are all points of S,*
- *xy is an edge iff $V_x \cap V_y \neq \emptyset$, and*
- *xyz is a triangle iff $V_x \cap V_y \cap V_z \neq \emptyset$.*

It turns out that Definition 5.2 indeed defines a planar triangulation on S. As a curiosity we mention that an edge xy of a Delaunay triangulation may partially lie outside of the union $V_x \cup V_y$.

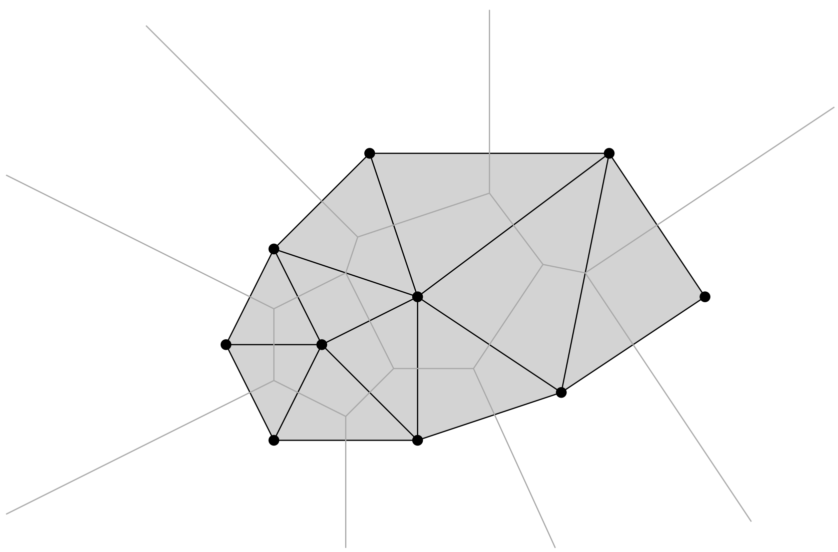

Figure 17: An example of a Delaunay triangulation with its underlying Voronoi decomposition.

Note that the edge xy of a Delaunay triangulation is a boundary edge (meaning it is contained in precisely one triangle) iff $V_x \cap V_y$ is unbounded. Similarly, x is a boundary vertex of a Delaunay triangulation (meaning it is an endpoint of some boundary edge) iff V_x is unbounded.

Locally Delaunay condition

For a triple of non-colinear points $x, y, z \in \mathbb{R}^2$ in the plane define $\mathcal{C}(x,y,z)$ to be the circle containing x, y, z, and let $\mathcal{B}(x,y,z)$ be the ball whose boundary is $\mathcal{C}(x, y, z)$.

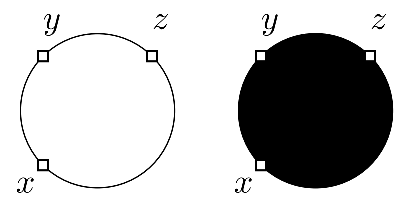

Figure 18: $\mathcal{C}(x, y, z)$ on the left and $\mathcal{B}(x, y, z)$ on the right.

Definition 5.3. *Suppose an edge xy is shared by two different triangles xyz and xyw of a triangulation. The edge xy is* ***locally Delaunay*** *[abbreviation: LD], if $w \notin \mathcal{B}(x, y, z)$.*

Proposition 5.4. *Suppose edge xy is shared by two different triangles xyz and xyw from a triangulation.*

1. *Definition 5.3 is symmetric, i.e., $w \notin \mathcal{B}(x, y, z)$ iff $z \notin \mathcal{B}(x, y, w)$.*

2. *Each edge in a Delaunay triangulation is LD.*

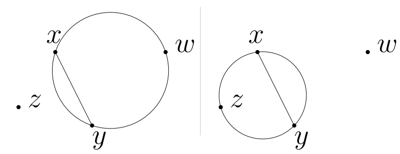

Figure 19: Proof of Proposition 5.4 (1).

Proof. Part (1) is apparent from Figure 19.

(2): Since abc is a triangle in $\mathcal{D}(S)$, there exists the corresponding Voronoi vertex $q = V_x \cap V_y \cap V_z$, which is the center of $\mathcal{C}(x, y, z)$. As $q \notin V_w$ (recall that no four Voronoi regions have a nonempty interesection by the general position property), $d(q, x) = d(q, y) = d(q, z) < d(q, w)$ by the definition of Voronoi regions, hence $w \notin \mathcal{B}(x, y, z)$. □

The property of being LD is a local property, shared by all edges of a Delaunay triangulation. It turns out that the converse of Proposition 5.4(2) is also true.

Theorem 5.5. *Suppose K is a triangulation on S. Then K is the Delaunay triangulation iff each edge is locally Delaunay.*

Construction of $\mathcal{D}(S)$

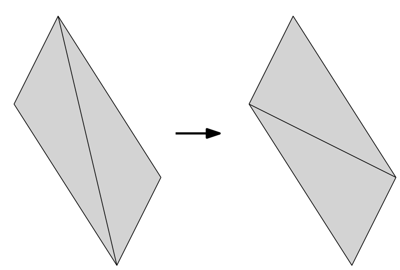
Figure 20: Edge flip.

Theorem 5.5 motivates the edge-flipping construction of Delaunay triangulations: starting with any triangulation on S (say, one obtained by a line sweep), keep flipping the non-LD edges. In order to algorithmically implement this construction we have to clarify two issues:

1. How do we verify the LD condition?

2. Does the procedure stop?

We address 1. first. It turns out it is not hard to verify the condition of LD using the incircle test.

Proposition 5.6. *[Incircle test] Suppose $x = (x_1, x_2), y = (y_1, y_2), z = (z_1, z_2)$ and $w = (w_1, w_2)$ are four points in $\mathbb{R}^2$. Assume x, y, z are* *not collinear and form a positively oriented triple, i.e.:*

$$\begin{vmatrix} 1 & x_1 & x_2 \\ 1 & y_1 & y_2 \\ 1 & z_1 & y_2 \end{vmatrix} > 0$$

Then $w \notin \mathcal{B}(x,y,z)$ *iff*

$$\begin{vmatrix} 1 & x_1 & x_2 & x_1^2 + x_2^2 \\ 1 & y_1 & y_2 & y_1^2 + y_2^2 \\ 1 & z_1 & z_2 & z_1^2 + z_2^2 \\ 1 & w_1 & w_2 & w_1^2 + w_2^2 \end{vmatrix} > 0.$$

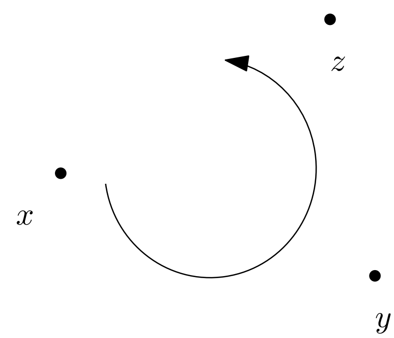

Figure 21: Positively oriented triple (x, y, z).

A proof and technical details of Proposition 5.6 are provided in the Appendix. While Proposition 5.6 provides a convenient way to verify LD property (and answer 1.), it does not suggest whether the edge flip algorithm actually stops (2.). In order to address this question we provide a couple more equivalent conditions to LD.

Suppose edge xy is shared by two different triangles xyz and xyw from a triangulation K on S. We say that edge xy is a **MaxMin** edge, if the minimal angle appearing in triangles xyz and xyw is larger than the minimal angle appearing in triangles xzw and yzw.

Proposition 5.7. *Suppose edge* xy *is shared by two different triangles* xyz *and* xyw *from a triangulation* K *on* S. *Then the following conditions are equivalent:*

(i) xy *is LD.*

(ii) xy *is a MaxMin edge.*

(iii) $\measuredangle xzy + \measuredangle xwy < \pi$.

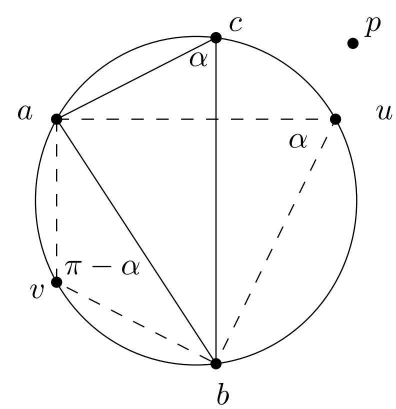

Figure 22: Inscribed angle theorem. Suppose $u, v \in \mathcal{C}(a,b,c)$, as the figure demonstrates. Then $\measuredangle acb = \measuredangle aub = \pi - \measuredangle avb$. This obviously implies $\measuredangle acb > \measuredangle apb$.

Proof. Let us prove the equivalence (i) $\Leftrightarrow$ (iii) first using the Inscribed angle theorem (see Figure 22).

xy is LD $\overset{\text{definition}}{\Longleftrightarrow} w \notin \mathcal{B}(x,y,z) \overset{\text{inscribed angle}}{\Longleftrightarrow} \pi - \measuredangle xzy > \measuredangle xwy \Leftrightarrow \measuredangle xzy + \measuredangle xwy < \pi$.

We now turn our attention to (i) $\Leftrightarrow$ (ii). Let α be the minimal angle appearing in triangles xyz, xyw, xzw and yzw. It is easy to see that α has to lie either along xy or zw as all the other angles get dissected (and hence decreased) by either xy or zw in one of the configurations.

Assume xy is LD. According to the Inscribed angle theorem, $\measuredangle xyz > \measuredangle xwz$, hence $\measuredangle xyz$ does not equal α. In the same way we can prove that no angle along xy equals α, hence xy is the MaxMin edge.

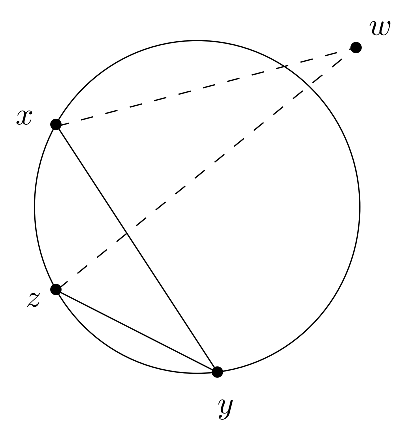

Figure 23: Proof of Proposition 5.7, (i) $\Leftrightarrow$ (ii).

Assume now that xy is not LD. Using the identical argument as in the previous paragraph we can prove that each angle along zw is larger than the corresponding angle along xy. Hence α has to lie along xy and therefore xy is not a MaxMin edge. □

Proposition 5.7 implies that each edge flip, which makes an edge in a triangulation LD, increases the minimal angles in the triangulation. Let us explain this statement in more detail. For each triangle T_i in a triangulation K on S let t_i denote the size of its minimal angle. Construct a lexicographically ordered list of these minimal angles, i.e., $t_{i_0} \leq t_{i_1} \leq \ldots \leq t_{i_m}$. Proposition 5.7 implies that every time we execute an edge flip making an edge in a triangulation LD, the new lexicographically ordered list of the minimal angles $t'_{i_0} \leq t'_{i_1} \leq \ldots \leq t'_{i_m}$ is lexicographical larger than the previous list, i.e., $t_{i_j} \leq t'_{i_j}, \forall j$ with strict inequality holding for at least one index j. Hence by making the required edge flips that keep turning edges into LD edges we can't return to the initial or any already visited triangulation. Since there are only finitely many triangulations on S, and therefore finitely many possible ordered lists of minimal angles, the edge flipping algorithm terminates, answering 2. above affirmatively.

Conclusion: the edge flipping algorithm terminates with $\mathcal{D}(S)$.

A triangulation, for which the lexicographically ordered list of the minimal angles is maximal in the lexicographical order, is called a MaxMin triangulation.

Theorem 5.8. *A MaxMin triangulation on S coincides with $\mathcal{D}(S)$.*

In particular, there exists only one MaxMin triangulation on S.

6 Concluding remarks

Recap (highlights) of this chapter

- Planar triangulations;
- Convexity;
- Euler characteristic;
- Line sweep;
- Voronoi diagram and Delaunay triangulation;
- Constructing the Delaunay triangulation using the locally Delaunay condition and the incircle test.

Background and applications

The Euler characteristic has been introduced by Leonhard Euler in the 18th century. The line sweep algorithm, Voronoi diagram and Delaunay triangulation are basic notions studied especially in computational geometry. The edge flip algorithm we mentioned requires $O(n^2)$ edge flips, where n is the number of vertices of S. There are known algorithms to construct the Delaunay triangulation in $O(n \log n)$.

The above mentioned properties of the Delaunay triangulation make it one of the favorite choices for a triangulation on a finite planar set S. For example, assume you are given a collection of points modelling a geographic profile in a small region. The points consist of coordinates and elevations at these coordinates. The task is to model the surface modelling the geographic profile. A standard solution would be to construct the Delaunay triangulation on the set of coordinate points, and then lift these points and triangles according to the given elevations. Triangles lifted this way provide a good approximation of the geographic profile on the sampled region.

Appendix: Proof of Proposition 5.6

Proof. Let us explain the positively oriented criterion first, see Figure 24. Points x, y, z form a positively oriented triple iff vectors $\overrightarrow{xy}, \overrightarrow{yz}$ are positively oriented, meaning that the third component of the cross product $(y_1 - x_1, y_2 - x_2, 0) \times (z_1 - x_1, z_2 - x_2)$ is positive. This third component equals the 3×3 determinant

$$\begin{vmatrix} 1 & x_1 & x_2 \\ 1 & y_1 & y_2 \\ 1 & z_1 & y_2 \end{vmatrix}.$$

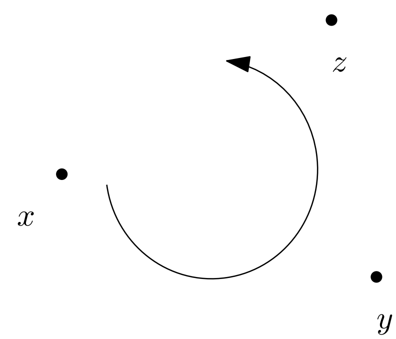

Figure 24: Positively oriented triple x, y, z.

We now turn our attention to the proof of the proposition. Surprisingly enough, we need to use the three-dimensional geometry, see Figure 25 throughout the proof. Embed $\mathbb{R}^2$ (and points x, y, z, w) into $\mathbb{R}^3$ by assigning the third coordinate to be 0, i.e., $x = (x_1, x_2, 0)$, etc. Consider the graph of the function $f(x, y) = x^2 + y^2$. Lift points x, y, z, w to the graph of f and let denote x', y', z', w' the lifted points, i.e., $x' = (x_1, x_2, x_1^2 + x_2^2)$, etc. Let Π denote the plane containing x', y', z'.

Let C be the intersection of the graph of f and Π. Note that the vertical projection of C onto $\mathbb{R}^2 \times \{0\}$ is a circle: substituting z in $z = x^2 + y^2$ by an equation of a plane $z = ax + by + c$ we obtain an equation of a circle in the plane of the form $x^2 + y^2 - ax - by - c = 0$. As this circle contains x, y, z, it coincides with $\mathcal{C}(x, y, z)$. It is geometrically apparent that $w \notin \mathcal{B}(x, y, z)$ iff w' lies above Π (the region where the graph of f is below Π is $\mathcal{B}(x, y, z)$). Since x, y, z are

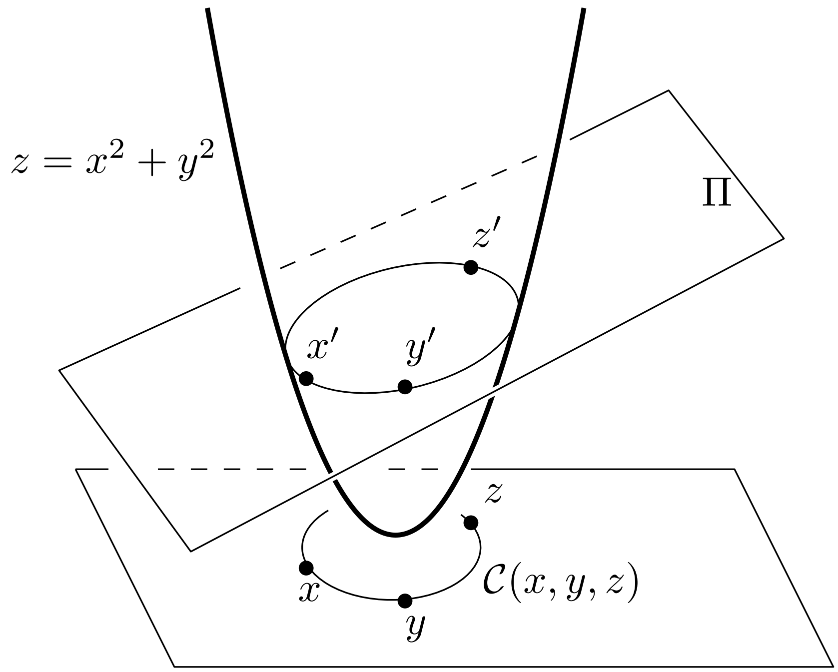

Figure 25: Proof of Proposition 5.6.

positively oriented, a normal of Π with a positive third component is $\vec{n} = \overrightarrow{x'y'} \times \overrightarrow{x'z'}$. Point w' lies above Π iff $\vec{n} \cdot \overrightarrow{x'w'}$ is positive. It is elementary to verify that $\vec{n} \cdot \overrightarrow{x'w'}$ equals

$$\begin{vmatrix} 1 & x_1 & x_2 & x_1^2 + x_2^2 \\ 0 & y_1 - x_1 & y_2 - x_2 & y_1^2 + y_2^2 - (x_1^2 + x_2^2) \\ 0 & z_1 - x_1 & z_2 - x_2 & z_1^2 + z_2^2 - (x_1^2 + x_2^2) \\ 0 & w_1 - x_1 & w_2 - x_2 & w_1^2 + w_2^2 - (x_1^2 + x_2^2) \end{vmatrix} = \begin{vmatrix} 1 & x_1 & x_2 & x_1^2 + x_2^2 \\ 1 & y_1 & y_2 & y_1^2 + y_2^2 \\ 1 & z_1 & z_2 & z_1^2 + z_2^2 \\ 1 & w_1 & w_2 & w_1^2 + w_2^2 \end{vmatrix}.$$

□