# Chapter 6

# Delaunay Triangulations

In Chapter 4 we have discussed triangulations of simple polygons. A triangulation partitions a polygon into triangles, which allows to easily compute the total area, or to derive a small guarding set, for instance. Another typical application is interpolation: Suppose a function $f$ is defined on the vertices of the polygon $P$, and we want to extend it "reasonably" and continuously to the entire $P$. To this end we take a triangulation $\mathcal{T}$. Given any point $p \in P$ we find a triangle $v_1v_2v_3 \in \mathcal{T}$ that contains $p$, and so $p = \sum_{i=1}^{3} \lambda_i v_i$ can be written as a (unique) convex combination of the three vertices. We may use the same coefficients to define an interpolated value $f(p) := \sum_{i=1}^{3} \lambda_i f(v_i)$.

If triangulations are a useful tool when working with polygons, they might also turn out useful for other geometric objects, such as point sets. But what could be a triangulation of a point set? Polygons have a clearly defined interior, which naturally lends itself to be covered by triangles. A point set does not have an interior, unless... we make one. Here the notion of convex hull comes handy. One way to think of a point set is as a convex polygon (its convex hull) potentially with some little holes (those points in the interior of the hull). A triangulation should then partition the convex hull while *respecting* the points in the interior. Figure 6.1b shows an example. In contrast, Figure 6.1c gives a counterexample: although the triangles do partition the convex hull, some points in the interior are not respected as they are swallowed by large triangles.

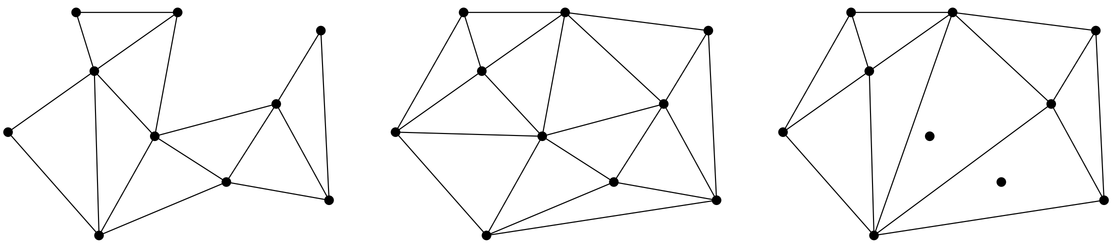

**(a)** Simple polygon triangulation. **(b)** Point set triangulation. **(c)** Not a triangulation.

**Figure 6.1:** *Examples of (non-)triangulations.*

This interpretation directly leads to the following adaption of Definition 4.4.

**Definition 6.1.** *A* **triangulation** *of a finite point set* $P \subset \mathbb{R}^2$ *is a collection* $\mathcal{T}$ *of triangles, such that*

*(1)* $\text{conv}(P) = \bigcup_{T \in \mathcal{T}} T$*;*

*(2)* $P = \bigcup_{T \in \mathcal{T}} V(T)$*; and*

*(3) for every distinct pair* $T, T' \in \mathcal{T}$*, the intersection* $T \cap T'$ *is either a common vertex, or a common edge, or empty.*

Just as for polygons, triangulations are universally available for point sets, meaning that (almost) every point set admits at least one.

**Proposition 6.2.** *Every set* $P \subseteq \mathbb{R}^2$ *of* $n \geqslant 3$ *points has a triangulation, unless all points in* $P$ *are collinear.*

*Proof.* In order to construct a triangulation for $P$, consider the lexicographically sorted sequence $p_1, \ldots, p_n$ of points in $P$. Let $m$ be minimal such that $p_1, \ldots, p_m$ are not collinear. We triangulate $p_1, \ldots, p_m$ by connecting $p_m$ to all of $p_1, \ldots, p_{m-1}$ (which *are* on a common line), see Figure 6.2a.

**(a)** Getting started. **(b)** Adding a point.

**Figure 6.2:** *Constructing the scan triangulation of* $P$.

Then we add $p_{m+1}, \ldots, p_n$ one by one. Let us inductively assume that we had built a triangulation of $P_{i-1} := \{p_1, \ldots, p_{i-1}\}$, and we are about to add $p_i$. Note that $p_i$ is not contained in $C_{i-1} := \text{conv}(P_{i-1})$ because of the lexicographic order. We connect it with all "visible" vertices of $C_{i-1}$; that is, every vertex $v$ of $C_{i-1}$ for which $\overline{p_i v} \cap C_{i-1} = \{v\}$. Among these vertices are two tangent points from $p_i$ to $C_{i-1}$, and the vertices in between are exactly the visible ones. After adding these connections, we have covered $C_i \setminus C_{i-1}$ by several new disjoint triangles, so overall we obtain a triangulation of $P_i$. □

The triangulation constructed in Proposition 6.2 is called a *scan triangulation*. Figure 6.3a shows a larger example. It is usually "ugly", though, as the lexicographic order tends to produce many long and skinny triangles. This is not only an aesthetic deficit, but also a practical concern in the context of interpolation, for example, since long and skinny triangles imply a less local interpolation. In contrast, the *Delaunay triangulation* of the same point set (Figure 6.3b) looks much nicer, and we will discuss in the next section how to get this triangulation.

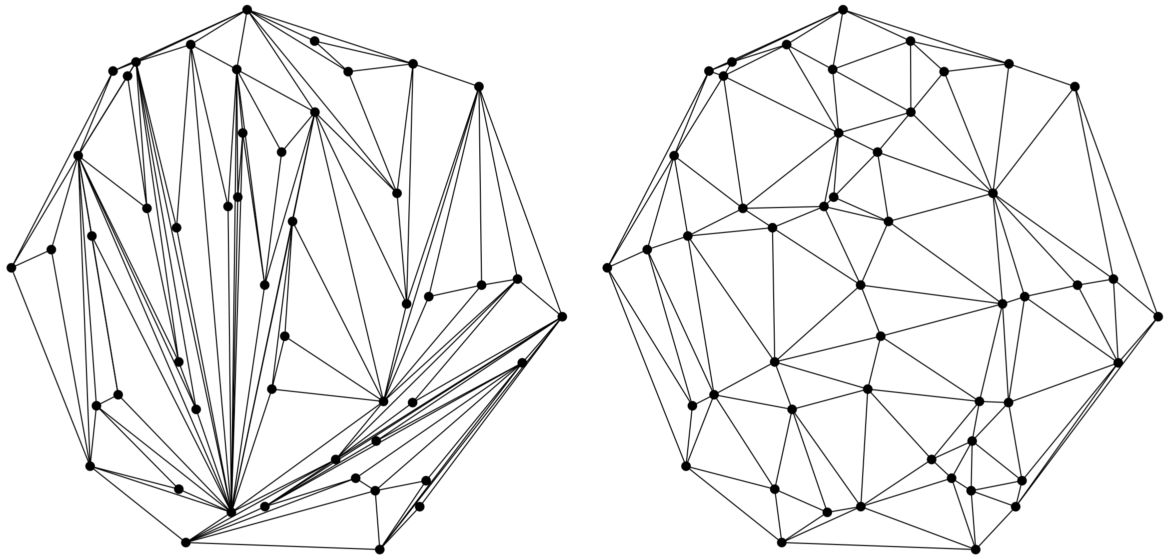

(a) Scan triangulation. (b) Delaunay triangulation.

**Figure 6.3:** *Two triangulations of the same set of* 50 *points.*

**Exercise 6.3.** *Describe how to implement the scan triangulation in* $O(n \log n)$ *time for a set of* $n$ *points in* $\mathbb{R}^2$.

On another note, if you look closely into the Graham Scan algorithm for planar convex hulls in Chapter 5, then you will realize that we also could have used it to prove Proposition 6.2. Whenever a point $q$ is discarded during Graham Scan due to a right turn $p \to q \to r$, we add the triangle $pqr$ to fill the space. Eventually this leads to a triangulation of the point set.

Every triangulation of $P$ induces a plane straight-line graph $G = (P, E)$, where the edges are the sides of the triangles. As shown by the lemma below (cf. Corollary 2.5), the counts of edges and triangles are determined by $P$.

**Lemma 6.4.** *Any triangulation of a set* $P \subset \mathbb{R}^2$ *of* $n$ *points has exactly* $3n-h-3$ *edges and* $2n-h-1$ *faces in its induced graph, where* $h := |P \cap \partial\text{conv}(P)|$ *is the number of points on the outer cycle.*

*Proof.* Consider the graph induced by an arbitrary triangulation of $P$. Denote by $E$ the set of edges and by $F$ the set of faces. We count the number of edge-face incidences in two ways. Denote $X = \{(e, f) \in E \times F : e \text{ bounds } f\}$.

On the one hand, every edge is incident to exactly two faces and therefore $|X| = 2|E|$. On the other hand, every inner face is a triangle and the outer face is bounded by $h$ edges, therefore $|X| = 3(|F| - 1) + h$. Together we obtain $3|F| = 2|E| - h + 3$. Combining with Euler's formula $n - |E| + |F| = 2$ we can solve for $|E| = 3n - h - 3$ and $|F| = 2n - h - 1$. □

In graph theory, the term "triangulation" is sometimes used as a synonym for "maximal planar graph". But geometric triangulations are somewhat weaker: They are not maximal in the sense that no abstract edge can be added; rather, only in the sense that no *straight-line* edge can be added without sacrificing planarity.

**Corollary 6.5.** *A triangulation of a set* $P \subset \mathbb{R}^2$ *of* $n$ *points is maximal planar, if and only if* $\text{conv}(P)$ *is a triangle.*

*Proof.* Combine Corollary 2.5 and Lemma 6.4. □

**Exercise 6.6.** *Find for every* $n \geqslant 3$ *a simple polygon* $P$ *with* $n$ *vertices such that* $P$ *has exactly one triangulation.* $P$ *should be in general position, meaning that no three vertices are collinear.*

**Exercise 6.7.** *Show that every set of* $n \geqslant 5$ *points in general position (no three points are collinear) has at least two different triangulations.*
*Hint: Show first that every set of five points in general position contains a convex 4-hole, that is, a subset of four points that span a convex quadrilateral that does not contain the fifth point.*

## 6.1 The Empty Circle Property

We will now move on to study the ominous and supposedly nice Delaunay triangulations mentioned above. They are defined in terms of an "empty circumcircle" property. The *circumcircle* of a triangle is the unique circle passing through the three vertices of the triangle, see Figure 6.4. Observe that long and skinny triangles usually have unproportionally large circumcircles, which tend to (though not always) enclose other points inside. A Delaunay triangulation forbids such enclosure, in hope of avoiding skinny triangles as much as possible.

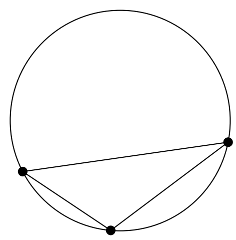

**Figure 6.4:** *Circumcircle of a triangle.*

**Definition 6.8.** *A triangulation* $\mathcal{T}$ *of a finite point set* $P \subset \mathbb{R}^2$ *is a* **Delaunay triangulation***, if the circumcircle of every triangle* $T \in \mathcal{T}$ *is* **empty***, that is, the circle does not enclose any point from* $P$ *strictly inside.*

Consider the example depicted in Figure 6.5. It shows a Delaunay triangulation of a set of six points: The circumcircles of all five triangles are empty (we also say that the

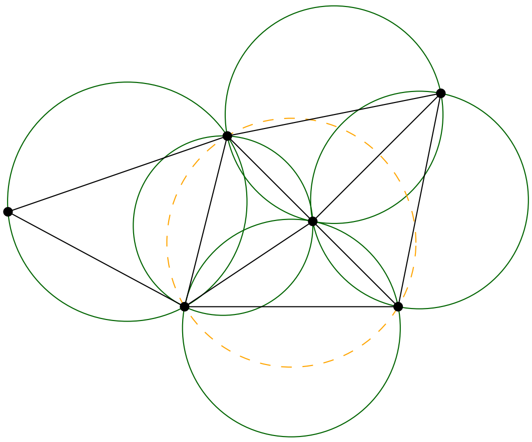

**Figure 6.5:** *All triangles satisfy the empty circle property.*

triangles satisfy the empty circle property). The dashed circle is not empty, but that is fine since it is not a circumcircle of any triangle.

It is instructive to look at the toy example where four points are arranged in convex position. Obviously, there are two possible triangulations. If the four points happen to lie on the same cycle $C$, the circumcircle of any three points is exactly $C$, which is empty, so both triangulations shall be Delaunay (see Figure 6.6a). But in general position, i.e. when the four points are not cocircular, only one triangulation is Delaunay (see Figures 6.6b and 6.6c). This case is formalized in the proposition below, whose proof technique will show up frequently in this chapter.

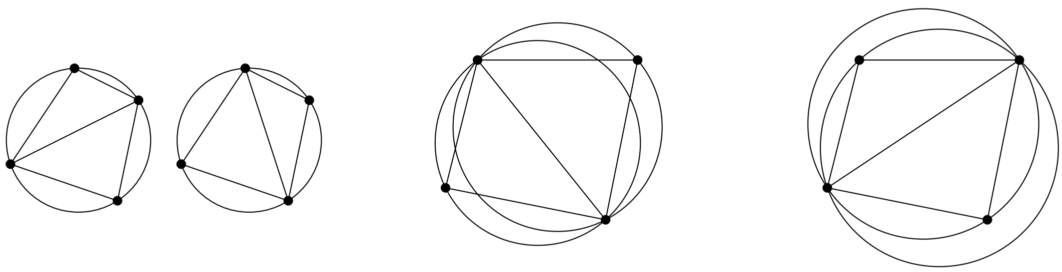

**(a)** Two Delaunay triangulations. **(b)** Delaunay triangulation. **(c)** Non-Delaunay triangulation.

**Figure 6.6:** *Triangulations of four points in convex position.*

**Proposition 6.9.** *Given a set $P \subset \mathbb{R}^2$ of four points that are in convex position but not cocircular. Then $P$ has exactly one Delaunay triangulation.*

*Proof.* Consider a set of four points $P = \{p, q, r, s\}$ arranged counterclockwise in convex position. There are only two possible triangulations: $\mathcal{T}_1 := \{prq, prs\}$ and $\mathcal{T}_2 := \{qsp, qsr\}$.

Let $C_1$ be the circumcircle of triangle $prq \in \mathcal{T}_1$, and $C_1'$ be the circumcircle of the other triangle $prs \in \mathcal{T}_1$. Since the four points are not cocircular, we have only two cases:

**$s$ is strictly outside $C_1$.** First we argue that $q$ must be strictly outside $C_1'$. Imagine the process of continuously moving from $C_1$ to $C_1'$ while keeping $p, r$ on the cycle (Figure 6.7a). More precisely, we move the center towards $s$ along the perpendicular bisector of $pr$. At some point the cycle hits $s$ and becomes $C_1'$ and the point $q$ must be "left behind". Thus $q$ is strictly outside $C_1'$, indeed.

As both $C_1$ and $C_1'$ are empty, $\mathcal{T}_1$ is a Delaunay triangulation. Next we argue that $\mathcal{T}_2$ is not Delaunay. Consider the continuous motion from $C_1$ to $C_2$, the circumcircle of $qsp \in \mathcal{T}_2$, while keeping $p, q$ intact (Figure 6.7b). The point $r$ is on $C_1$ and remains within the circle all the way up to $C_2$. This means $C_2$ is not empty, thus $\mathcal{T}_2$ is not Delaunay.

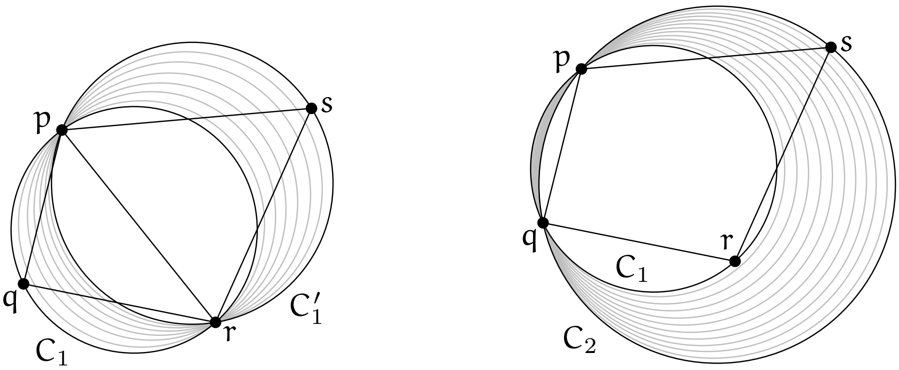


**(a)** Going from $C_1$ to $C_1'$ in $\mathcal{C}_1$. **(b)** Going from $C_1$ to $C_2$ in $\mathcal{C}_2$.

**Figure 6.7:** *Circumcircles and containment for triangulations of four points.*

**$s$ is strictly inside $C_1$.** The case is symmetric: just shift the roles of $pqrs$ to $qrsp$.

□

**Exercise 6.10.** *Prove or disprove that every minimum weight triangulation (that is, a triangulation for which the sum of edge lengths is minimum) is a Delaunay triangulation.*

## 6.2 The Lawson Flip algorithm

It is not clear yet that every point set $P$ of $n$ points actually has a Delaunay triangulation (given that not all points are collinear). In this and the next two sections, we will prove that this is the case, via the *Lawson flip algorithm*:

1. Compute some triangulation of $P$ (for example, the scan triangulation).

2. While there exist two adjacent triangles $\Delta, \Delta'$ such that the circumcircle of $\Delta$ encloses a vertex of $\Delta'$ (see Figure 6.6c; observe that the four vertices must be in

convex position), replace them by the other pair of adjacent triangles (Figure 6.6b). In other words, we flip the diagonal of the convex quadrilateral.

We call the replacement operation in the second step a *(Lawson) flip*.

**Theorem 6.11.** *Let $P \subseteq \mathbb{R}^2$ be a set of $n$ points, equipped with some triangulation $\mathcal{T}$. The Lawson flip algorithm terminates after at most $\binom{n}{2} = O(n^2)$ flips, and the resulting triangulation $\mathcal{D}$ is a Delaunay triangulation of $P$.*

We will prove Theorem 6.11 in two steps: In Section 6.3 we show that the program described above always terminates and, therefore, is an algorithm indeed. (If you think about it a little, it is not obvious whether the algorithm would loop indefinitely.) Then in Section 6.4 we show that the algorithm does produce a Delaunay triangulation upon termination.

## 6.3 Termination of the Lawson Flip Algorithm

For the termination proof, we make use of the (parabolic) *lifting map* $\ell$:

$$p = (x, y) \in \mathbb{R}^2 \quad \mapsto \quad \ell(p) = (x, y, x^2 + y^2) \in \mathbb{R}^3.$$

Geometrically, $\ell$ "lifts" the point vertically up until hitting the *unit paraboloid* $\{(x, y, z) \mid z = x^2 + y^2\} \subseteq \mathbb{R}^3$, see Figure 6.8a.

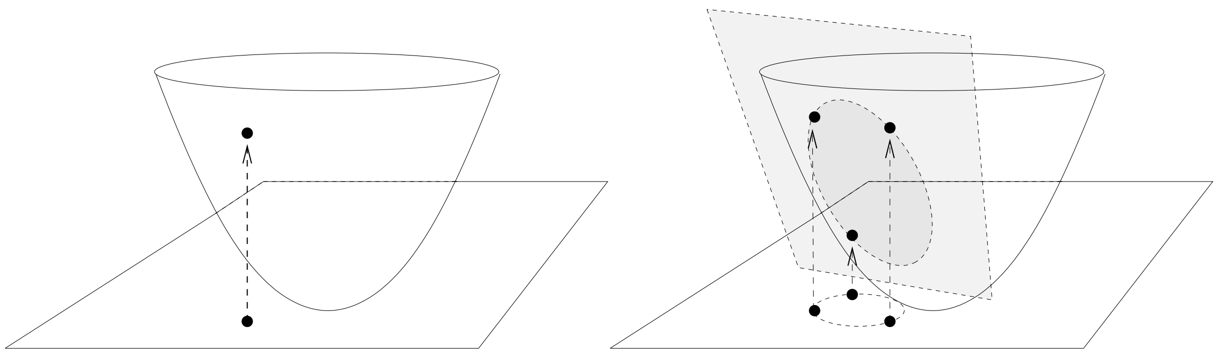

**(a)** The lifting map.

**(b)** Points on/inside/outside a circle are lifted to points on/below/above a plane.

**Figure 6.8:** *The lifting map: circles map to planes.*

Recall the following important property of the lifting map that we proved in Exercise 5.31. It is illustrated in Figure 6.8b.

**Lemma 6.12.** *Let $C \subseteq \mathbb{R}^2$ be a circle of positive radius. The "lifted circle" $\ell(C) := \{\ell(p) \mid p \in C\}$ is contained in a unique plane $h_C \subseteq \mathbb{R}^3$. Moreover, a point $p \in \mathbb{R}^2$ is strictly inside (respectively outside) $C$ if and only if the lifted point $\ell(p)$ is strictly below (respectively above) $h_C$.*

Using the lifting map, we can interpret triangulations in the 3D space. For each triangle $\Delta = pqr$, we define its "lifted version" as $\ell(\Delta) := \text{conv}\{\ell(p), \ell(q), \ell(r)\}$, which is a triangle hanging in the space with $\Delta$ being its "shadow". This way, the triangulation is lifted to a piecewise linear surface in the space consisting of patches of triangles.

Consider a Lawson flip on adjacent triangles $\Delta = pqr$ and $\Delta' = pqs$, where $p, q, r, s$ are in convex position. Let $C$ and $C'$ be their respective circumcircles. By the condition of a flip, $C$ encloses $s$, and similarly $C'$ encloses $r$. In the lifted picture, Lemma 6.12 states that $\ell(s)$ is strictly below the plane that contains $\ell(\Delta)$, and similarly $\ell(r)$ is strictly below the plane that contains $\ell(\Delta')$. In other words, the triangles $\ell(\Delta)$ and $\ell(\Delta)$ form a mountain that protrudes upward; see Figure 6.9a.

After the flip, the two triangles are replaced by $prs$ and $qrs$. In the lifted picture, triangles form a valley that protrudes downward by a similar reasoning; see Figure 6.9b.

More pictorially, imagine an opaque tetrahedron $\text{conv}\{\ell(p), \ell(q), \ell(r), \ell(s)\}$ in the space. When you look at it from the top, you see two faces corresponding to the two triangles before the flip; and when you look from the bottom, you see the other two faces corresponding to the two triangles after the flip. (You cannot see three faces from either direction, since $p, q, r, s$ are in convex position.) Hence a Lawson flip can be interpreted as replacing the two top faces by the two bottom faces of the tetrahedron.

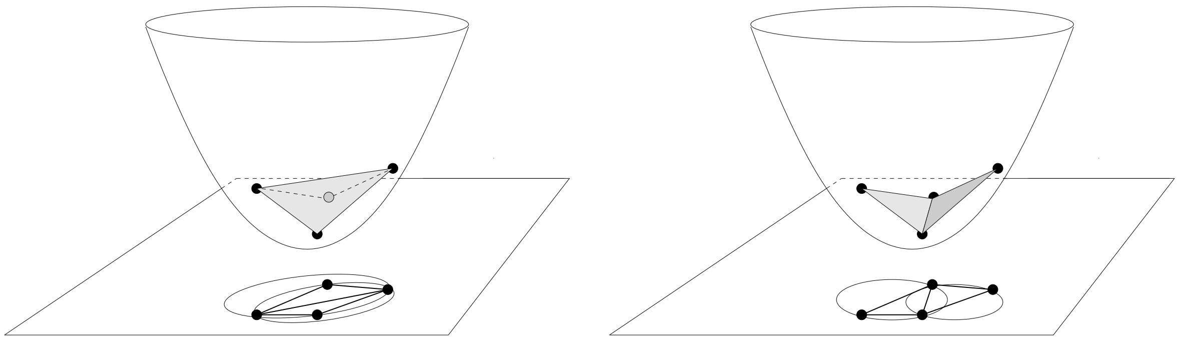

**(a)** Before the flip: the top two triangles of the tetrahedron and the corresponding non-Delaunay triangulation in the plane.

**(b)** After the flip: the bottom two triangles of the tetrahedron and the corresponding Delaunay triangulation in the plane.

**Figure 6.9:** *Lawson flip: the height of the surface of lifted triangles decreases.*

It follows that the 3D surface can only grow strictly downward pointwise. In particular, once an edge $pq$ has been flipped, it becomes strictly above the surface thereafter and thus can never show up again. Since $n$ points can span at most $\binom{n}{2}$ edges, the bound on the number of flips follows.

## 6.4 Correctness of the Lawson Flip Algorithm

The triangulation of $P$ that we get upon termination of the Lawson flip algorithm is "locally Delaunay": it checks the empty circle property for adjacent triangles only. But in fact it is "globally Delaunay", too.

**Proposition 6.13.** *The triangulation $\mathcal{D}$ that results from the Lawson flip algorithm is a Delaunay triangulation.*

*Proof.* Suppose for contradiction that there is some triangle $\Delta \in \mathcal{D}$ and some point $p \in P$ strictly inside the circumcircle $C$ of $\Delta$. Among all such pairs $(\Delta, p)$, we choose one that minimizes the distance from $p$ to $\Delta$. Note that this distance is positive by definition of a triangulation. We assume for now that the point on $\Delta$ closest to $p$ lies on the relative interior of some edge $e$ of $\Delta$; we will deal with the other case later. The situation is as depicted in Figure 6.10a.

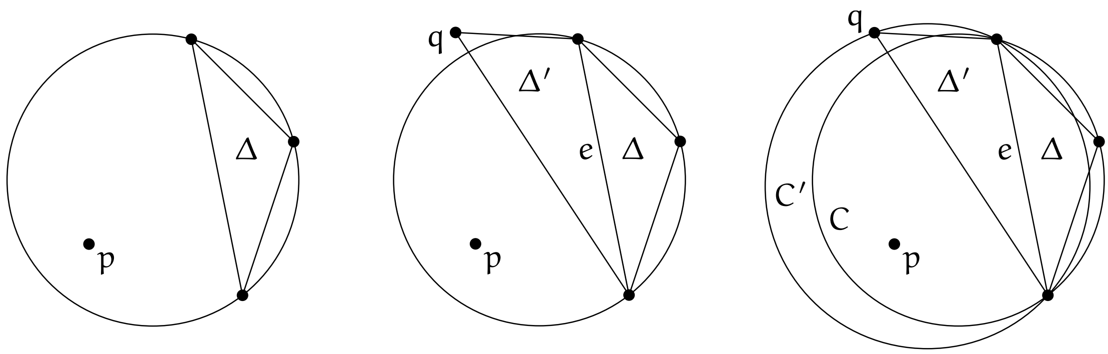


**(a)** A point $p$ inside the circumcircle $C$ of a triangle $\Delta$.

**(b)** The edge $e$ of $\Delta$ closest to $p$ and the second triangle $\Delta'$ incident to $e$.

**(c)** The circumcircle $C'$ of $\Delta'$ also contains $p$, and $p$ is closer to $\Delta'$ than to $\Delta$.

**Figure 6.10:** *Correctness of the Lawson flip algorithm.*

There must be another triangle $\Delta'$ in $\mathcal{D}$ that is incident to the edge $e$. By the local Delaunay property of $\mathcal{D}$, the third vertex $q$ of $\Delta'$ is on or outside of $C$, see Figure 6.10b. But then the circumcircle $C'$ of $\Delta'$ contains the whole portion of $C$ on $p$'s side of $e$, hence it also contains $p$; moreover, $p$ is closer to $\Delta'$ than to $\Delta$ (Figure 6.10c). But this is a contradiction to our choice of $\Delta$ and $p$. Hence there was no $(\Delta, p)$, and $\mathcal{D}$ is a Delaunay triangulation.

Consider now the special case where the point on $\Delta$ closest to $p$ happens to be a vertex $v$ of the triangle $\Delta$. In this case, we need some additional care when choosing $\Delta$. Among all triangles that use $v$ as a vertex and that have $p$ in their circumcircle, we choose our actual $\Delta$ as the one for which the edge $e$ (as before, this is the edge that faces $p$ in the circumcircle) and the segment $\overline{pv}$ form an angle closest 90 degrees.

From here, the proof proceeds as in the first case. We construct a new triangle $\Delta'$ that also uses the edge $e$ and that also contains the point $p$ in its circumcircle. The difference is that we do not necessarily get the same type of contradiction because the point on $\Delta'$ closest to $p$ might still be $v$. If that is the case, however, the angle between the edge $e'$ (this is the edge that faces $p$ in the circumcircle of $\Delta'$) and the segment $\overline{pv}$

has will be closer to $90^\circ$ compared to $e$. This now stands in contradiction to our more careful choice of the triangle $\Delta$, which finishes the proof. □

**Exercise 6.14.** *The Euclidean minimum spanning tree (EMST) of a finite point set* $P \subset \mathbb{R}^2$ *is a spanning tree for which the sum of the edge lengths is minimum (among all spanning trees of* $P$*). Show:*

*(a) Every EMST of* $P$ *is a plane graph.*

*(b) Every EMST of* $P$ *contains a closest pair, that is, an edge between two points* $p, q \in P$ *that have minimum distance to each other among all point pairs in* $\binom{P}{2}$*.*

*(c) Every Delaunay Triangulation of* $P$ *contains an EMST of* $P$*.*

**Exercise 6.15.** *(a) Show that for any two triangulations* $T_1$ *and* $T_2$ *on a point set* $P$*, it is possible to transform* $T_1$ *into* $T_2$ *using* $O(n^2)$ *edge flips.*

*(b) Let* $D = P \cup Q$ *where* $P = \{p_1, p_2, \ldots, p_n\}$ *and* $Q = \{q_1, q_2, \ldots, q_n\}$ *each forms a slightly bent arc, facing against each other. For any line* $q_i q_j$ *the set* $P$ *is on its left; and symmetrically, for any line* $p_i p_j$ *the set* $Q$ *is on its right. Show that there are two triangulations* $T_1$ *and* $T_2$ *on* $D$ *such that at least* $\Omega(n^2)$ *edge flips are needed to transform* $T_1$ *into* $T_2$*.*

*(c) Show that* $D$ *can be constructed in such a way that one of the triangulations from (b), say,* $T_1$ *is a Delaunay triangulation.*

## 6.5 The Delaunay Graph

Despite the fact that a point set may have more than one Delaunay triangulation, there are certain edges that are present in every Delaunay triangulation, for instance, the edges of the convex hull.

**Definition 6.16.** *The* **Delaunay graph** *of* $P \subseteq \mathbb{R}^2$ *consists of all line segments* $\overline{pq}$*, for* $p, q \in P$*, that are contained in every Delaunay triangulation of* $P$*.*

The following characterizes the edges of the Delaunay graph.

**Lemma 6.17.** *The segment* $\overline{pq}$*, for* $p, q \in P$*, is in the Delaunay graph of* $P$ *if and only if there exists a circle through* $p$ *and* $q$ *for which all other points of* $P$ *are strictly outside.*

*Proof.* "⇒": Let $pq$ be an edge in the Delaunay graph of $P$, and let $\mathcal{D}$ be a Delaunay triangulation of $P$. Then there exists a triangle $\Delta = pqr$ in $\mathcal{D}$, whose circumcircle $C$ does not enclose any point from $P$ strictly inside.

If there is a point $s$ on $C$ such that $\overline{rs}$ intersects $\overline{pq}$, then let $\Delta' = pqt \neq \Delta$ denote the other triangle in $\mathcal{D}$ that is incident to $pq$ (Figure 6.11a). Note that $t$ must be on $C$, for

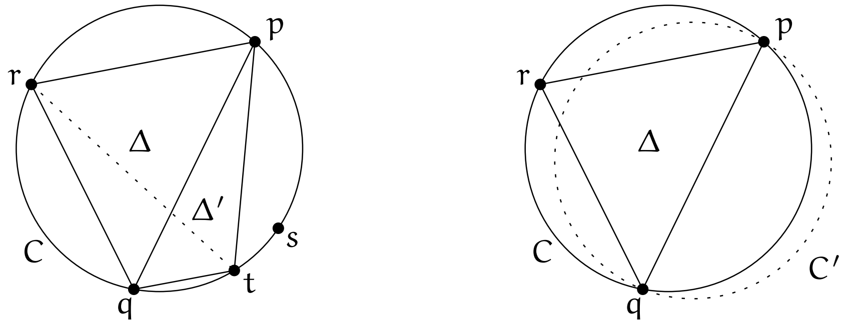


**(a)** Another point $s \in \partial C$. **(b)** Moving C away from s.

**Figure 6.11:** *Characterization of edges in the Delaunay graph (I).*

otherwise the circumcircle of $\Delta'$ would enclose s. Now flipping the edge pq to rt yields another Delaunay triangulation that does not contain the edge pq, in contradiction to pq being an edge in the Delaunay graph. Therefore, there is no such point s.

Otherwise we can slightly change the circle C by moving away from r while keeping p and q on the circle. As P is a finite point set, we can do such a modification without catching another point from P with the circle. In this way we obtain a circle $C'$ through p and q such that all other points from P are strictly outside $C'$ (Figure 6.11b).

"$\Leftarrow$": Let $\mathcal{D}$ be a Delaunay triangulation of P. If $\overline{pq}$ is not an edge of $\mathcal{D}$, there must be another edge of $\mathcal{D}$ that crosses $\overline{pq}$ (otherwise, we could add $\overline{pq}$ to $\mathcal{D}$ and still have a plane graph, a contradiction to $\mathcal{D}$ being a triangulation of P). Let rs denote the first edge of $\mathcal{D}$ that the directed line segment $\overrightarrow{pq}$ intersects.

Consider the triangle $\Delta$ of $\mathcal{D}$ that is incident to rs on the side that faces p (given that $\overline{rs}$ intersects $\overline{pq}$ this is a well defined direction). By the choice of rs neither of the other two edges of $\Delta$ intersects $\overline{pq}$, and $p \notin \Delta^\circ$ because $\Delta$ is part of a triangulation of P. The only remaining option is that p is a vertex of $\Delta = prs$. As $\Delta$ is part of a Delaunay triangulation, its circumcircle $C_\Delta$ needs to be empty.

Consider now a circle C through p and q for which all other points are strictly outside. Fixing p and q, we expand C towards r to eventually obtain the circle $C'$ through $p, q, r$ (Figure 6.12a). Recall that r and s are on different sides of the line through p and q. Therefore, s lies strictly outside $C'$. Next fix p and r and expand $C'$ towards s to eventually obtain the circle $C_\Delta$ through $p, r, s$ (Figure 6.12b). Recall that s and q are on the same side of the line through p and r. Therefore, $q \in C_\Delta$, which is in contradiction to $C_\Delta$ being empty. It follows that there is no Delaunay triangulation of P that does not contain the edge pq. □

The Delaunay graph is useful to prove uniqueness of the Delaunay triangulation in case of general position.

**Corollary 6.18.** *Let $P \subset \mathbb{R}^2$ be a finite set of points in general position (no four points of P are cocircular). Then P has a unique Delaunay triangulation.*

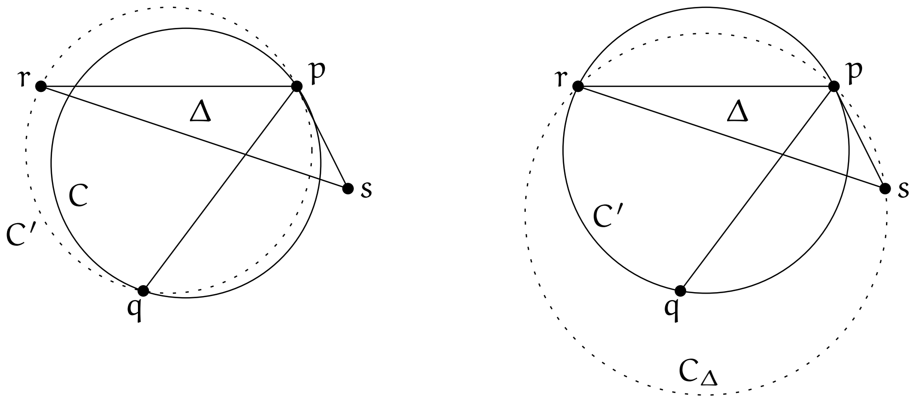


**(a)** Expanding $C$ towards $r$. **(b)** Expanding $C'$ towards $s$.

**Figure 6.12:** *Characterization of edges in the Delaunay graph (II).*

**Exercise 6.19.** *Prove Corollary 6.18.*

## 6.6 Every Delaunay Triangulation Maximizes the Smallest Angle

Why are we interested in Delaunay triangulations? It turns out that Delaunay triangulations satisfy a number of interesting properties. Here we give a scientific explanation for their nice looks.

Recall that when we compared a scan triangulation with a Delaunay triangulation of the same point set in Figure 6.3, we claimed that the scan triangulation is "ugly" because it contains many long and skinny triangles. The triangles of the Delaunay triangulation, at least in this example, look much nicer, that is, much closer to an equilateral triangle. One way to quantify this "niceness" is to look at the angles that appear in a triangulation: If all angles are large, then all triangles are reasonably close to an equilateral triangle. Indeed, we will show that Delaunay triangulations maximize the smallest angle among all triangulations of a given point set. This is not saying that there are no long and skinny triangles in a Delaunay triangulation. But if there is one, then the small angle is inherent: there would exist at least as skinny triangle in *every* triangulation of the point set.

Every triangulation $\mathcal{T}$ of $P$ induces an *angle sequence* $A(\mathcal{T}) = (\theta_1, \theta_2, \ldots, \theta_{3m})$ which lists the measures of interior angles of all $T \in \mathcal{T}$, sorted increasingly so that $0 < \theta_1 \leqslant \theta_2 \leqslant \cdots \leqslant \theta_{3m} < \pi$. Here $m$ is the number of triangles, which is a constant determined by $P$; see Lemma 6.4. Let $\mathcal{T}, \mathcal{T}'$ be two triangulations of $P$. We say that $A(\mathcal{T}) < A(\mathcal{T}')$ if there is some $i$ for which $\theta_i < \theta'_i$ and $\theta_j = \theta'_j$, for all $j < i$. (This is nothing but the lexicographic order on angle sequences.) We write $A(\mathcal{T}) \leqslant A(\mathcal{T}')$ if $A(\mathcal{T}) < A(\mathcal{T}')$ or $A(\mathcal{T}) = A(\mathcal{T}')$.

**Theorem 6.20.** *Let $P \subseteq \mathbb{R}^2$ be a finite set of points in general position (not all collinear and no four cocircular). Let $\mathcal{D}^*$ be the unique Delaunay triangulation of $P$, and let*

*$\mathcal{T}$ be any triangulation of $P$. Then $A(\mathcal{T}) \leqslant A(\mathcal{D}^*)$.*

In particular, $\mathcal{D}^*$ maximizes the smallest angle among all triangulations of $P$.

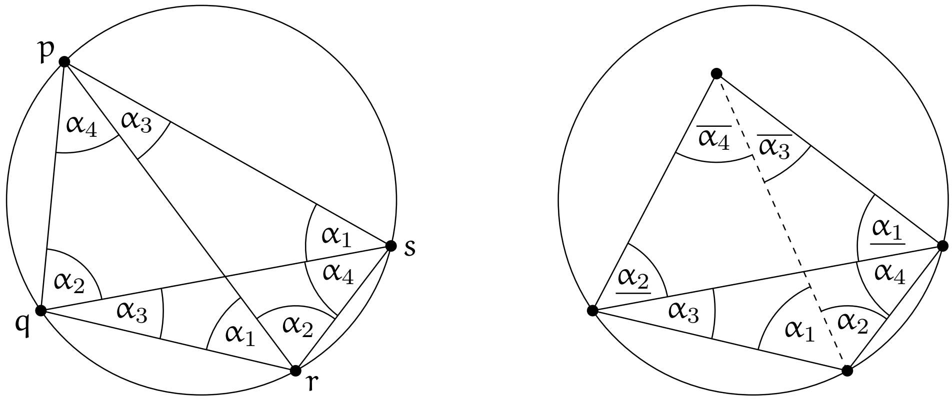


**(a)** Four cocircular points and the induced eight angles.

**(b)** The situation before a flip.

**Figure 6.13:** *Angle-optimality of Delaunay triangulations.*

*Proof.* We know that $\mathcal{T}$ can be transformed into $\mathcal{D}^*$ through the Lawson flip algorithm, and we are done if we can show that each flip lexicographically increases the angle sequence. Recall that a flip involves two triangles and thus effectively expels six angles from the sequence and injects another six. We claim that the minimum of the six new angles is *strictly* larger than the minimum of the six old angles. This claim, once proven, would imply that the sequence increases lexicographically: Before flipping, let $0 < \theta < \pi$ be the minimum of the six old angles and $i \in \{1, \ldots, 3m\}$ be the last position that the value occurs in the sequence; after flipping, all values at positions $j < i$ shall persist whereas the value at position $i$ shall strictly increase.

Next we proceed to show the claim. Let us first look at the situation of four cocircular points; see Figure 6.13a. The *inscribed angle theorem* (a generalization of Thales' Theorem, stated below as Theorem 6.21) tells us that the eight depicted angles come in four equal pairs. For instance, the angles labeled $\alpha_1$ at $s$ and $r$ are angles on the same side of the chord $pq$ of the circle.

In our situation, however, no four points are cocircular. When we perform a Lawson flip, the picture is as in Figure 6.13b where we are about to replace the solid with the dashed diagonal. Here we use under- and over-lines to suggest the relation between angles; angle $\underline{\alpha}$ (repectively $\overline{\alpha}$) is strictly smaller (respectively larger) than $\alpha$. At the flip, the six old angles are

$$\alpha_1 + \alpha_2, \quad \alpha_3, \quad \alpha_4, \quad \underline{\alpha_1}, \quad \underline{\alpha_2}, \quad \overline{\alpha_3} + \overline{\alpha_4},$$

and the six new angles are

$$\alpha_1, \quad \alpha_2, \quad \overline{\alpha_3}, \quad \overline{\alpha_4}, \quad \underline{\alpha_1} + \alpha_4, \quad \underline{\alpha_2} + \alpha_3.$$

Now, *every* new angle is larger than *some* old angle:

$$\begin{aligned} \alpha_1 &> \underline{\alpha_1}, \\ \alpha_2 &> \underline{\alpha_2}, \\ \overline{\alpha_3} &> \alpha_3, \\ \overline{\alpha_4} &> \alpha_4, \\ \underline{\alpha_1} + \alpha_4 &> \alpha_4, \\ \underline{\alpha_2} + \alpha_3 &> \alpha_3. \end{aligned}$$

So the minimum of the new angles is strictly larger than the minimum of the old angles, as claimed. □

**Theorem 6.21** (Inscribed Angle Theorem). *Let* $pq$ *be a chord on a circle* $C$. *Then* $\angle prq$ *stays constant when the point* $r$ *moves along any of the two arcs between* $p$ *and* $q$.

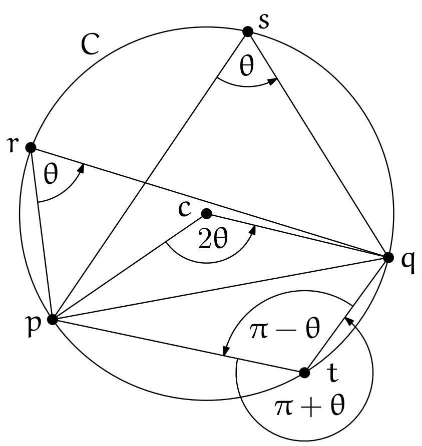


**Figure 6.14:** *The Inscribed Angle Theorem with* $\theta := \angle prq$.

*Proof.* Without loss of generality we may assume that $c$ is located to the left of or on the oriented line $pq$.

Consider first the case that the triangle $\Delta = pqr$ contains $c$. Then $\Delta$ can be partitioned into three triangles: $pcr$, $qcr$, and $cpq$. All three triangles are isosceles, because two sides of each form the radius of $C$. Denote $\alpha = \angle prc$, $\beta = \angle crq$, $\gamma = \angle cpq$, and $\delta = \angle pcq$ (see the figure shown to the right). The angles we are interested in are $\theta = \angle prq = \alpha + \beta$ and $\delta$, and we will show that $\delta = 2\theta$.

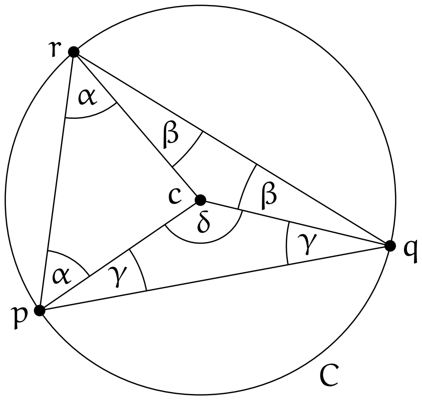


Indeed, the angle sum in $\Delta$ is $\pi = 2(\alpha + \beta + \gamma)$ and the angle sum in the triangle $cpq$ is $\pi = \delta + 2\gamma$. Combining both yields $\delta = 2(\alpha + \beta) = 2\theta$.

Next suppose that pqcr are in convex position and r is to the left of or on the oriented line pq. Without loss of generality let r be to the left of or on the oriented line qc. (The case that r lies to the right of or on the oriented line pc is symmetric.) Define $\alpha$, $\beta$, $\gamma$, $\delta$ as above and observe that $\theta = \alpha - \beta$. Again we show that $\delta = 2\theta$.

The angle sum in the triangle cpq is $\pi = \delta + 2\gamma$ and the angle sum in the triangle rpq is $\pi = (\alpha - \beta) + \alpha + \gamma + (\gamma - \beta) = 2(\alpha + \gamma - \beta)$. Combining both yields $\delta = \pi - 2\gamma = 2(\alpha - \beta) = 2\theta$.

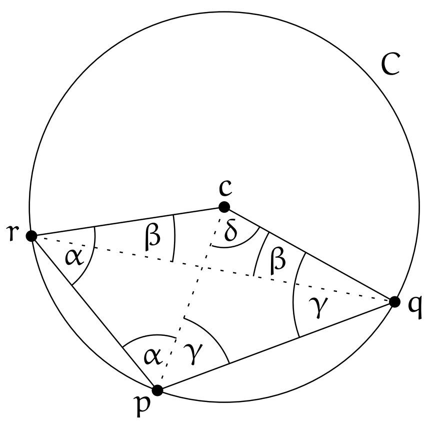


It remains to consider the case that r is to the right of the oriented line pq. Consider the point r′ that is antipodal to r on C, and the quadrilateral Q = prqr′. We are interested in the angle $\phi$ of Q at r. By Thales' Theorem the inner angles of Q at p and q are both $\pi/2$. Hence the angle sum of Q is $2\pi = \theta + \phi + 2\pi/2$ and so $\phi = \pi - \theta$. As shown in the first two cases, $\theta$ is a constant and thus $\phi$ is also a constant.

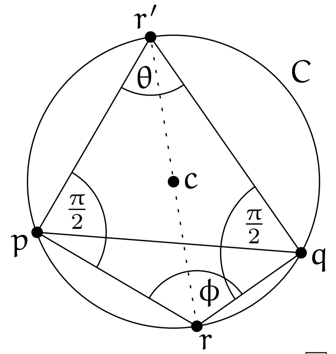


□

What happens in the case where the Delaunay triangulation is not unique? The following still holds.

**Theorem 6.22.** *Let $P \subseteq \mathbb{R}^2$ be a finite set of points, not all on a line. Every Delaunay triangulation $\mathcal{D}$ of $P$ maximizes the smallest angle among all triangulations $\mathcal{T}$ of $P$.*

*Proof.* Let $\mathcal{D}$ be some Delaunay triangulation of P. We infinitesimally perturb the points in P such that no four are on a common circle anymore. Then the Delaunay triangulation becomes unique (Corollary 6.18). Starting from $\mathcal{D}$, we keep applying Lawson flips until we reach the unique Delaunay triangulation $\mathcal{D}^*$ of the perturbed point set. Now we examine this sequence of flips on the original *unperturbed* point set. All these flips must involve four cocircular points (only in the cocircular case, an infinitesimal perturbation can change "good" edges into "bad" edges that still need to be flipped). But as Figure 6.13 (a) easily implies, such a "degenerate" flip does not change the smallest of the six involved angles. It follows that $\mathcal{D}$ and $\mathcal{D}^*$ have the same smallest angle, and since $\mathcal{D}^*$ maximizes the smallest angle among all triangulations $\mathcal{T}$ (Theorem 6.20), so does $\mathcal{D}$. □

## 6.7 Constrained Triangulations

Sometimes one would like to have a Delaunay triangulation, but certain edges are already prescribed. Of course, one cannot expect to be able to get a proper Delaunay triangulation where all triangles satisfy the empty circle property. But it is possible to obtain some triangulation that comes as close as possible to a proper Delaunay triangulation, given that we are forced to include the edges in $E$. Such a triangulation is called a *constrained Delaunay triangulation*, a formal definition of which follows.

Let $P \subseteq \mathbb{R}^2$ be a finite point set and $G = (P, E)$ a geometric graph with vertex set $P$ and straight-line edges $E$. A triangulation $\mathcal{T}$ of $P$ is said to be a *constrained Delaunay triangulation* with respect to $G$ if it contains all edges in $E$ and, for every triangle $\Delta \in \mathcal{T}$,

> The circumcircle of $\Delta$ does not enclose any point $q \in P$ visible from $\Delta^\circ$. A point $q \in P$ is *visible* from $\Delta^\circ$ if there exists a point $p \in \Delta^\circ$ such that the line segment $\overline{pq}$ does not cross any $e \in E$. We can thus imagine the line segments of $E$ as "blocking the view".

For illustration, consider the simple polygon and its constrained Delaunay triangulation shown in Figure 6.15, where the thick edges are prescribed. The circumcircle of the shaded triangle $\Delta$ contains a lot of points in its interior, but that does not matter since the points are blocked by the edge $e$ and are thus invisible from $\Delta^\circ$.

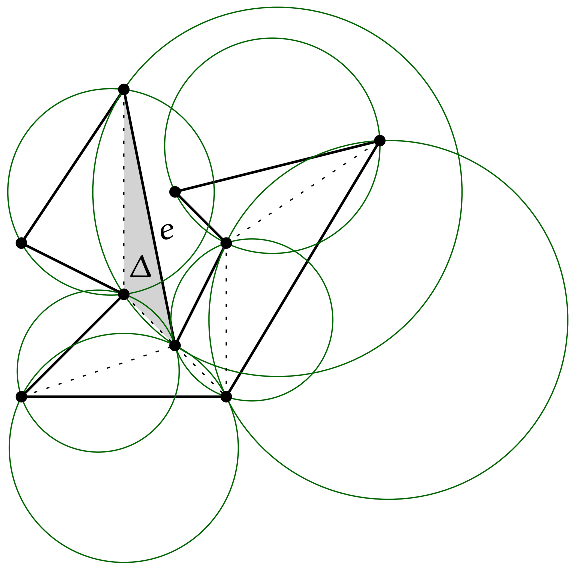


**Figure 6.15:** *Constrained Delaunay triangulation of a simple polygon.*

**Theorem 6.23.** *For every finite point set $P$ and every plane graph $G = (P, E)$, there exists a constrained Delaunay triangulation of $P$ with respect to $G$.*

**Exercise 6.24.** *Prove Theorem 6.23. Also describe a polynomial algorithm to construct such a triangulation.*

## Questions

23. *What is a triangulation?* Provide the definition and prove a basic property: every triangulation with the same set of vertices and the same outer face has the same

number of triangles.

24. *What is a triangulation of a point set?* Give a precise definition.
25. *Does every point set (not all points on a common line) have a triangulation?* You may, for example, argue with the scan triangulation.
26. *What is a Delaunay triangulation of a set of points*? Give a precise definition.
27. *What is the Delaunay graph of a point set?* Give a precise definition and a characterization.
28. *How can you prove that every set of points (not all on a common line) has a Delaunay triangulation?* You can for example sketch the Lawson flip algorithm and the Lifting Map, and use these to show the existence.
29. *When is the Delaunay triangulation of a point set unique?* Show that general position is a sufficient condition. Is it also necessary?
30. *What can you say about the "quality" of a Delaunay triangulation?* Prove that every Delaunay triangulation maximizes the smallest interior angle in the triangulation, among the set of all triangulations of the same point set.