Chapter 2

Two-dimensional Delaunay triangulations

The Delaunay triangulation is a geometric structure that engineers have used for meshes since mesh generation was in its infancy. In two dimensions, it has a striking advantage: among all possible triangulations of a fixed set of points, the Delaunay triangulation maximizes the minimum angle. It also optimizes several other geometric criteria related to interpolation accuracy. If it is our goal to create a triangulation without small angles, it seems almost silly to consider a triangulation that is not Delaunay. Delaunay triangulations have been studied thoroughly, and excellent algorithms are available for constructing and updating them.

A constrained triangulation is a triangulation that enforces the presence of specified edges—for example, the boundary of a nonconvex object. A constrained Delaunay triangulation relaxes the Delaunay property just enough to recover those edges, while enjoying optimality properties similar to those of a Delaunay triangulation. Constrained Delaunay triangulations are nearly as popular as their unconstrained ancestors.

This chapter surveys two-dimensional Delaunay triangulations, constrained Delaunay triangulations, weighted Delaunay triangulations, and their geometric properties.

2.1 Triangulations of a planar point set

The word *triangulation* usually refers to a simplicial complex, but it has multiple meanings when we discuss a *triangulation of* some geometric entity that is being triangulated. There are triangulations of point sets, polygons, polyhedra, and many other structures. Consider points in the plane (or in any Euclidean space).

Definition 2.1 (triangulation of a point set)**.** Let S be a finite set of points in the plane. A *triangulation of* S is a simplicial complex $\mathcal{T}$ such that S is the set of vertices in $\mathcal{T}$, and the union of all the simplices in $\mathcal{T}$ is the convex hull of S—that is, $|\mathcal{T}| = \operatorname{conv} S$.

Does every point set have a triangulation? Yes. Consider the *lexicographic triangulation* illustrated in Figure 2.1. To construct one, sort the points *lexicographically* (that is, by x-coordinate, ordering points with the same x-coordinate according to their y-coordinates),

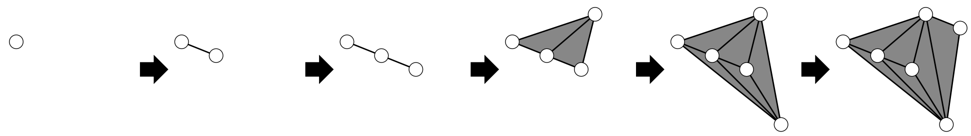

Figure 2.1: Incremental construction of a lexicographic triangulation.

yielding a sorted sequence $v_1, v_2, \ldots, v_n$ of points. Define the lexicographic triangulation $\mathcal{T}_i$ of the first i points by induction as follows. The first triangulation is $\mathcal{T}_1 = \{v_1\}$. Each subsequent triangulation is $\mathcal{T}_i = \mathcal{T}_{i-1} \cup \{v_i\} \cup \{\text{conv}(\{v_i\} \cup \sigma) : \sigma \in \mathcal{T}_{i-1}$ and the relative interior of $\text{conv}(\{v_i\} \cup \sigma)$ intersects no simplex in $\mathcal{T}_{i-1}\}$.

Even if the points in S are all collinear, there is a triangulation of S: $\mathcal{T}_n$ contains n vertices, $n-1$ collinear edges connecting them, and no triangles.

A triangulation of n points in the plane has at most $2n-5$ triangles and $3n-6$ edges as a consequence of Euler's formula. With no change, Definition 2.1 defines triangulations of point sets in higher-dimensional Euclidean spaces as well.

2.2 The Delaunay triangulation

The *Delaunay triangulation* of a point set S, introduced by Boris Nikolaevich Delaunay in 1934, is characterized by the *empty circumdisk property*: no point in S lies in the interior of any triangle's circumscribing disk; recall Definition 1.17.

Definition 2.2 (Delaunay)**.** In the context of a finite point set S, a triangle is *Delaunay* if its vertices are in S and its open circumdisk is *empty*— i.e. contains no point in S. Note that any number of points in S can lie on a Delaunay triangle's circumcircle. An edge is *Delaunay* if its vertices are in S and it has at least one empty open circumdisk. A *Delaunay triangulation* of S, denoted Del S, is a triangulation of S in which every triangle is Delaunay, as illustrated in Figure 2.2.

One might wonder whether every point set has a Delaunay triangulation, and how many Delaunay triangulations a point set can have. The answer to the first question is "yes." Section 2.3 gives some intuition for why this is true, and Section 2.5 gives a proof.

The Delaunay triangulation of S is unique if and only if no four points in S lie on a common empty circle, a fact proved in Section 2.7. Otherwise, there are Delaunay triangles

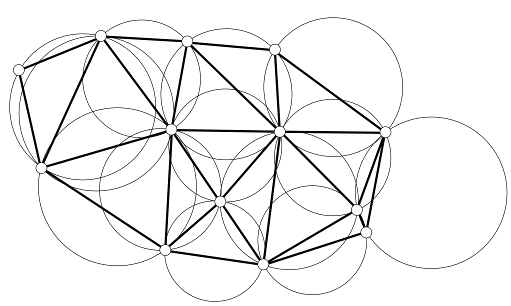

Figure 2.2: Every triangle in a Delaunay triangulation has an empty open circumdisk.

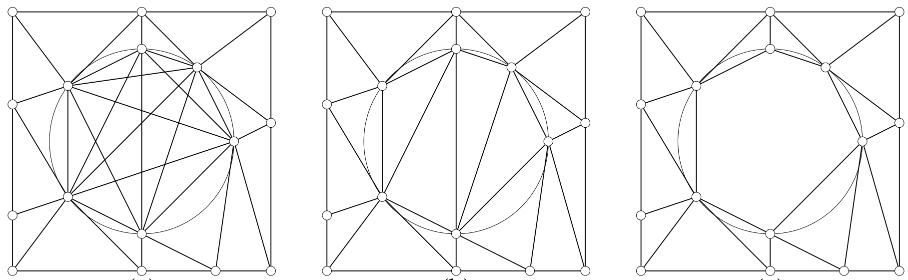

Figure 2.3: Three ways to define a Delaunay structure in the presence of cocircular vertices. (a) Include all the Delaunay simplices. (b) Choose a subset of Delaunay simplices that constitutes a triangulation. (c) Exclude all crossing Delaunay edges, and fuse overlapping Delaunay triangles into Delaunay polygons.

and edges whose interiors intersect, as illustrated in Figure 2.3(a). Most applications omit some of these triangles and edges so that the survivors form a simplicial complex, as in Figure 2.3(b). Depending on which Delaunay simplices one keeps and which one discards, one obtains different Delaunay triangulations.

It is sometimes useful to unite the intersecting triangles into a single polygon, depicted in Figure 2.3(c). The *Delaunay subdivision* obtained this way is a polyhedral complex, rather than a simplicial complex. It has the advantage of being the geometric dual of the famous *Voronoi diagram*, discussed in Section 7.1.

Clearly, a simplex's being Delaunay does not guarantee that it is in *every* Delaunay triangulation of a point set. But a slightly stronger property does provide that guarantee.

Definition 2.3 (strongly Delaunay)**.** In the context of a finite point set S, a triangle τ is *strongly Delaunay* if its vertices are in S and its *closed* circumdisk contains no point in S except the vertices of τ. An edge e is *strongly Delaunay* if its vertices are in S and it has at least one closed circumdisk that contains no point in S except the vertices of e. Every point in S is a strongly Delaunay vertex.

Every Delaunay triangulation of S contains every strongly Delaunay simplex, a fact proved in Section 2.7. The Delaunay subdivision contains the strongly Delaunay edges and triangles, and no others.

Consider two examples of strongly Delaunay edges. First, every edge on the boundary of a triangulation of S is strongly Delaunay. Figure 2.4 shows why. Second, the edge connecting a point $v \in S$ to its nearest neighbor $w \in S$ is strongly Delaunay, because the smallest closed disk containing v and w does not contain any other point in S. Therefore, every Delaunay triangulation connects every vertex to its nearest neighbor.

2.3 The parabolic lifting map

Given a finite point set S, the *parabolic lifting map* transforms the Delaunay subdivision of S into faces of a convex polyhedron in three dimensions, as illustrated in Figure 2.5.

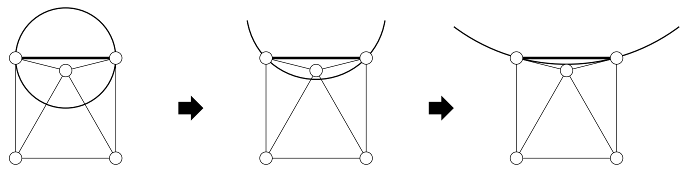

Figure 2.4: Every edge on the boundary of a convex triangulation is strongly Delaunay, because it is always possible to find an empty disk that contains its endpoints and no other vertex.

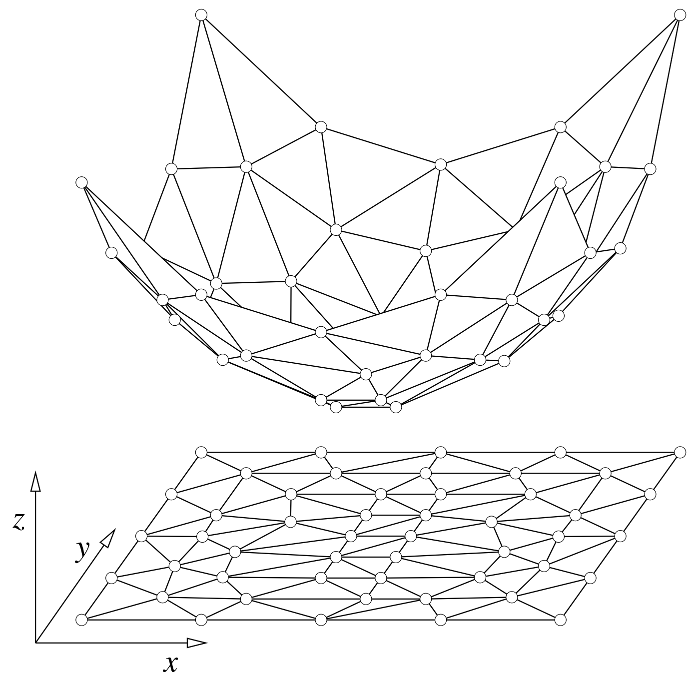

Figure 2.5: The parabolic lifting map.

This relationship between Delaunay triangulations and convex hulls has two consequences. First, it makes many properties of the Delaunay triangulation intuitive. For example, from the fact that every finite point set has a polyhedral convex hull, it follows that every finite point set has a Delaunay triangulation. Second, it brings to mesh generation the power of a huge literature on polytope theory and algorithms. For example, every convex hull algorithm is a Delaunay triangulation algorithm!

The parabolic lifting map sends each point $p = (x, y) \in \mathbb{R}^2$ to a point $p^+ = (x, y, x^2 + y^2) \in \mathbb{R}^3$. Call p^+ the *lifted companion* of p.

Consider the convex hull $\operatorname{conv} S^+$ of the lifted points $S^+ = \{v^+ : v \in S\}$. Figure 2.5 illustrates its downward-facing faces. Formally, a face f of $\operatorname{conv} S^+$ is *downward-facing* if no point in $\operatorname{conv} S^+$ is directly below any point in f, with respect to the z-axis. Call the collection of downward-facing faces the *underside* of $\operatorname{conv} S^+$. Projecting the underside of $\operatorname{conv} S^+$ to the x-y plane (by discarding every point's z-coordinate) yields the Delaunay subdivision of S. If S has more than one Delaunay triangulation, this Delaunay subdivision has non-triangular polygons, like the hexagon in Figure 2.3(c). Triangulating these polygonal faces yields a Delaunay triangulation.

For a simplex σ in the plane, its *lifted companion* σ^+ is the simplex embedded in $\mathbb{R}^3$ whose vertices are the lifted companions of the vertices of σ. Note that σ^+ is flat and does not curve to hug the paraboloid. The following lemmas show that every Delaunay simplex's lifted companion is included in a downward-facing face of conv S^+.

Lemma 2.1 (Lifting Lemma)**.** *Let C be a circle in the plane. Let $C^+ = \{p^+ : p \in C\}$ be the ellipse obtained by lifting C to the paraboloid. Then the points of C^+ lie on a plane h, which is not parallel to the z-axis. Furthermore, every point p inside C lifts to a point p^+ below h, and every point p outside C lifts to a point p^+ above h. Therefore, testing whether a point p is inside, on, or outside C is equivalent to testing whether the lifted point p^+ is below, on, or above h.*

Proof. Let o and r be the center and radius of C, respectively. Let p be a point in the plane. The z-coordinate of p^+ is $\|p\|^2$. By expanding $d(o, p)^2$, we have the identity $\|p\|^2 = 2\langle o, p\rangle - \|o\|^2 + d(o, p)^2$. With o and r fixed and $p \in \mathbb{R}^2$ varying, the equation $z = 2\langle o, p\rangle - \|o\|^2 + r^2$ defines a plane h in $\mathbb{R}^3$, not parallel to the z-axis. For every point $p \in C$, $d(o, p) = r$, so $C^+ \subset h$. For every point $p \notin C$, if $d(o, p) < r$, then the lifted point p^+ lies below h, and if $d(o, p) > r$, then p^+ lies above h. □

Proposition 2.2. *Let σ be a simplex whose vertices are in S, and let σ^+ be its lifted companion. Then σ is Delaunay if and only if σ^+ is included in some downward-facing face of* conv S^+*. The simplex σ is strongly Delaunay if and only if σ^+ is a downward-facing face of* conv S^+.

Proof. If σ is Delaunay, σ has a circumcircle C that encloses no point in S. Let h be the unique plane in $\mathbb{R}^3$ that includes C^+. By the Lifting Lemma (Lemma 2.1), no point in S^+ lies below h. Because the vertices of σ^+ are in C^+, $h \supset \sigma^+$. Therefore, σ^+ is included in a downward-facing face of the convex hull of S^+. If σ is strongly Delaunay, every point in S^+ lies above h except the vertices of σ^+. Therefore, σ^+ is a downward-facing face of the convex hull of S^+. The converse implications follow by reversing the argument. □

The parabolic lifting map works equally well for Delaunay triangulations in three or more dimensions; the Lifting Lemma (Lemma 2.1) and Proposition 2.2 generalize to higher dimensions without any new ideas. Proposition 2.2 implies that any algorithm for constructing the convex hull of a point set in $\mathbb{R}^{d+1}$ can construct the Delaunay triangulation of a point set in $\mathbb{R}^d$.

2.4 The Delaunay Lemma

Perhaps the most important result concerning Delaunay triangulations is the *Delaunay Lemma*, proved by Boris Delaunay himself. It provides an alternative characterization of the Delaunay triangulation: a triangulation whose edges are *locally Delaunay*.

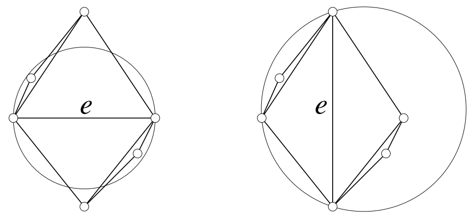

Figure 2.6: At left, e is locally Delaunay. At right, e is not.

Definition 2.4 (locally Delaunay)**.** Let e be an edge in a triangulation $\mathcal{T}$ in the plane. If e is an edge of fewer than two triangles in $\mathcal{T}$, then e is said to be *locally Delaunay*. If e is an edge of exactly two triangles τ_1 and τ_2 in $\mathcal{T}$, then e is said to be *locally Delaunay* if it has an open circumdisk containing no vertex of τ_1 nor τ_2. Equivalently, the open circumdisk of τ_1 contains no vertex of τ_2. Equivalently, the open circumdisk of τ_2 contains no vertex of τ_1.

Figure 2.6 shows two different triangulations of six vertices. In the triangulation at left, the edge e is locally Delaunay, because the depicted circumdisk of e does not contain either vertex opposite e. Nevertheless, e is not Delaunay, thanks to other vertices in e's circumdisk. In the triangulation at right, e is not locally Delaunay; every open circumdisk of e contains at least one of the two vertices opposite e.

The Delaunay Lemma has several uses. First, it provides a linear-time algorithm to determine whether a triangulation of a point set is Delaunay: simply test whether every edge is locally Delaunay. Second, it implies a simple algorithm for producing a Delaunay triangulation called the *flip algorithm* (Section 2.5). The flip algorithm helps to prove that Delaunay triangulations have useful optimality properties. Third, the Delaunay Lemma helps to prove the correctness of other algorithms for constructing Delaunay triangulations.

As with many properties of Delaunay triangulations, the lifting map provides intuition for the Delaunay Lemma. On the lifting map, the Delaunay Lemma is essentially the observation that a simple polyhedron is convex if and only if its has no *reflex edge*. A reflex edge is an edge where the polyhedron is locally nonconvex; that is, two adjoining triangles meet along that edge at a dihedral angle greater than $180°$, measured through the interior of the polyhedron. If a triangulation has an edge that is not locally Delaunay, that edge's lifted companion is a reflex edge of the lifted triangulation (by the Lifting Lemma, Lemma 2.1).

Lemma 2.3 (Delaunay Lemma)**.** *Let $\mathcal{T}$ be a triangulation of a point set S. The following three statements are equivalent.*

(i) *Every triangle in $\mathcal{T}$ is Delaunay (i.e. $\mathcal{T}$ is Delaunay).*

(ii) *Every edge in $\mathcal{T}$ is Delaunay.*

(iii) *Every edge in $\mathcal{T}$ is locally Delaunay.*

Proof. If the points in S are all collinear, S has only one triangulation, which trivially satisfies all three properties.

Otherwise, let e be an edge in $\mathcal{T}$; e is an edge of at least one triangle $\tau \in \mathcal{T}$. If τ is Delaunay, τ's open circumdisk is empty, and because τ's circumdisk is also a circumdisk

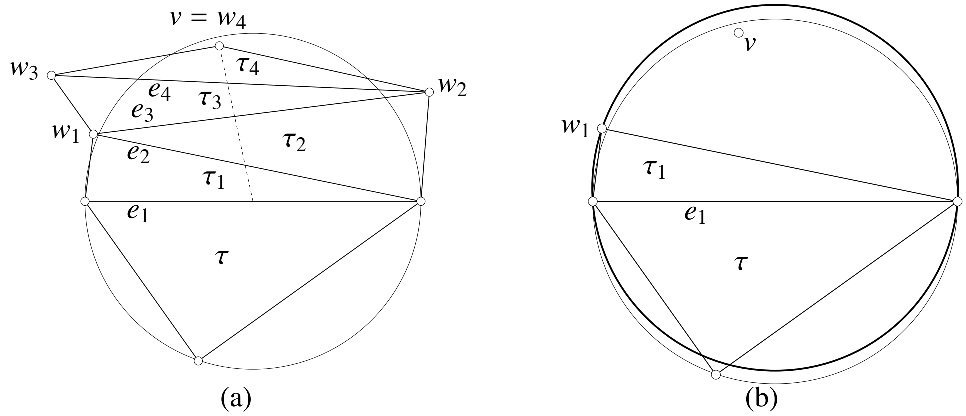

Figure 2.7: (a) Because τ's open circumdisk contains v, some edge between v and τ is not locally Delaunay. (b) Because v lies above e_1 and in τ's open circumdisk, and because w_1 lies outside τ's open circumdisk, v must lie in τ_1's open circumdisk.

of e, e is Delaunay. Therefore, Property (i) implies Property (ii). If an edge is Delaunay, it is clearly locally Delaunay too, so Property (ii) implies Property (iii). The proof is complete if Property (iii) implies Property (i). Of course, this is the hard part.

Suppose that every edge in $\mathcal{T}$ is locally Delaunay. Suppose for the sake of contradiction that Property (i) does not hold. Then some triangle $\tau \in \mathcal{T}$ is not Delaunay, and some vertex $v \in S$ is inside τ's open circumdisk. Let e_1 be the edge of τ that separates v from the interior of τ, as illustrated in Figure 2.7(a). Without loss of generality, assume that e_1 is oriented horizontally, with τ below e_1.

Draw a line segment ℓ from the midpoint of e_1 to v—see the dashed line in Figure 2.7(a). If the line segment intersects some vertex other than v, replace v with the lowest such vertex and shorten ℓ accordingly. Let $e_1, e_2, e_3, \ldots, e_m$ be the sequence of triangulation edges (from bottom to top) whose relative interiors intersect ℓ. Because $\mathcal{T}$ is a triangulation of S, every point on the line segment lies either in a single triangle or on an edge. Let w_i be the vertex above e_i that forms a triangle τ_i in conjunction with e_i. Observe that $w_m = v$.

By assumption, e_1 is locally Delaunay, so w_1 lies outside the open circumdisk of τ. As Figure 2.7(b) shows, it follows that the open circumdisk of τ_1 includes the portion of τ's open circumdisk above e_1 and, hence, contains v. Repeating this argument inductively, we find that the open circumdisks of $\tau_2, \ldots, \tau_m$ contain v. But $w_m = v$ is a vertex of τ_m, which contradicts the claim that v is in the open circumdisk of τ_m. □

2.5 The flip algorithm

The *flip algorithm* has at least three uses: it is a simple algorithm for computing a Delaunay triangulation, it is the core of a constructive proof that every finite set of points in the plane has a Delaunay triangulation, and it is the core of a proof that the Delaunay triangulation optimizes several geometric criteria when compared with all other triangulations of the same point set.

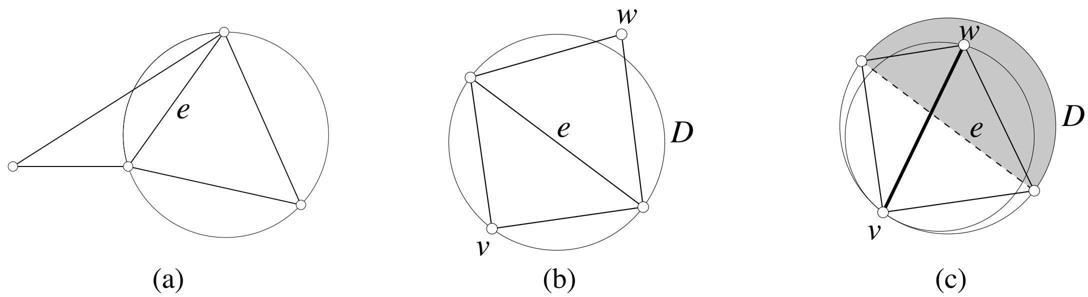

Figure 2.8: (a) In this nonconvex quadrilateral, e cannot be flipped, and e is locally Delaunay. (b) The edge e is locally Delaunay. (c) The edge e is not locally Delaunay. The edge created by a flip of e is locally Delaunay.

Let S be a point set to be triangulated. The flip algorithm begins with any triangulation $\mathcal{T}$ of S; for instance, the lexicographic triangulation described in Section 2.1. The Delaunay Lemma tells us that $\mathcal{T}$ is Delaunay if and only if every edge in $\mathcal{T}$ is locally Delaunay. The flip algorithm repeatedly chooses any edge that is not locally Delaunay and *flips* it.

The union of two triangles that share an edge is a quadrilateral, and the shared edge is a diagonal of the quadrilateral. To flip an edge is to replace it with the quadrilateral's other diagonal, as illustrated in Figure 2.6. An edge flip is legal only if the two diagonals cross each other—equivalently, if the quadrilateral is strictly convex. Fortunately, unflippable edges are always locally Delaunay, as Figure 2.8(a) shows.

Proposition 2.4. *Let e be an edge in a triangulation of S. Either e is locally Delaunay, or e is flippable and the edge created by flipping e is locally Delaunay.*

Proof. Let v and w be the vertices opposite e. Consider the quadrilateral defined by e, v, and w, illustrated in Figure 2.8. Let D be the open disk whose boundary passes through v and the vertices of e.

If w is outside D, as in Figure 2.8(b), then the empty circumdisk D demonstrates that e is locally Delaunay.

Otherwise, w is in the section of D bounded by e and opposite v. This section is shaded in Figure 2.8(c). The quadrilateral is thus strictly convex, so e is flippable. Furthermore, the open disk that is tangent to D at v and has w on its boundary does not contain the vertices of e, because D includes it, as Figure 2.8(c) demonstrates. Therefore, the edge vw is locally Delaunay. □

Proposition 2.4 shows that the flip algorithm can flip any edge that is not locally Delaunay, thereby creating an edge that is. Unfortunately, the outer four edges of the quadrilateral might discover that they are no longer locally Delaunay, even if they were locally Delaunay before the flip. If the flip algorithm repeatedly flips edges that are not locally Delaunay, will it go on forever? The following proposition says that it won't.

Proposition 2.5. *Given a triangulation of n points, the flip algorithm terminates after $O(n^2)$ edge flips, yielding a Delaunay triangulation.*

Proof. Let $\mathcal{T}$ be the initial triangulation provided as input to the flip algorithm. Let $\mathcal{T}^+ = \{\sigma^+ : \sigma \in \mathcal{T}\}$ be the initial triangulation lifted to the parabolic lifting map; $\mathcal{T}^+$ is a simplicial complex embedded in $\mathbb{R}^3$. If $\mathcal{T}$ is Delaunay, then $\mathcal{T}^+$ triangulates the underside of conv S^+; otherwise, by the Lifting Lemma (Lemma 2.1), the edges of $\mathcal{T}$ that are not locally Delaunay lift to reflex edges of $\mathcal{T}^+$.

By Proposition 2.4, an edge flip replaces an edge that is not locally Delaunay with one that is. In the lifted triangulation $\mathcal{T}^+$, a flip replaces a reflex edge with a convex edge. Let Q be the set containing the four vertices of the two triangles that share the flipped edge. Then conv Q^+ is a tetrahedron whose upper faces are the pre-flip simplices and whose lower faces are the post-flip simplices. Imagine the edge flip as the act of gluing the tetrahedron conv Q^+ to the underside of $\mathcal{T}^+$.

Each edge flip monotonically lowers the lifted triangulation, so once flipped, an edge can never reappear. The flip algorithm can perform no more than $n(n-1)/2$ flips—the number of edges that can be defined on n vertices—so it must terminate. But the flip algorithm terminates only when every edge is locally Delaunay. By the Delaunay Lemma, the final triangulation is Delaunay. □

The fact that the flip algorithm terminates helps to prove that point sets have Delaunay triangulations.

Proposition 2.6. *Every finite set of points in the plane has a Delaunay triangulation.*

Proof. Section 2.1 demonstrates that every finite point set has at least one triangulation. By Proposition 2.5, the application of the flip algorithm to that triangulation produces a Delaunay triangulation. □

An efficient implementation of the flip algorithm requires one extra ingredient. How quickly can one find an edge that is not locally Delaunay? To repeatedly test every edge in the triangulation would be slow. Instead, the flip algorithm maintains a list of edges that might not be locally Delaunay. The list initially contains every edge in the triangulation. Thereafter, the flip algorithm iterates the following procedure until the list is empty, whereupon the algorithm halts.

- Remove an edge from the list.
- Check whether the edge is still in the triangulation, and if so, whether it is locally Delaunay.
- If the edge is present but not locally Delaunay, flip it, and add the four edges of the flipped quadrilateral to the list.

The list may contain multiple copies of the same edge, but they do no harm.

Implemented this way, the flip algorithm runs in $O(n + k)$ time, where n is the number of vertices (or triangles) of the triangulation and k is the number of flips performed. In the worst case, $k = \Theta(n^2)$, giving $O(n^2)$ running time. But there are circumstances where the flip algorithm is fast in practice. For instance, if the vertices of a Delaunay mesh are perturbed by small displacements during a physical simulation, it might take only a small

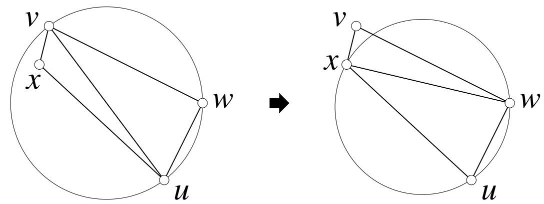

Figure 2.9: A Delaunay flip increases the angle opposite edge uw and, if $\angle wxu$ is acute, reduces the size of the circumdisk of the triangle adjoining that edge.

number of flips to restore the Delaunay property. In this circumstance, the flip algorithm probably outperforms any algorithm that reconstructs the triangulation from scratch.

2.6 The optimality of the Delaunay triangulation

Delaunay triangulations are valuable in part because they optimize several geometric criteria: the smallest angle, the largest circumdisk, and the largest min-containment disk. Recall from Definition 1.20 that the *min-containment disk* of a triangle is the smallest closed disk that includes it. For a triangle with no obtuse angle, the circumdisk and the min-containment disk are the same, but for an obtuse triangle, the min-containment disk is smaller.

Proposition 2.7. *Flipping an edge that is not locally Delaunay increases the minimum angle and reduces the largest circumdisk among the triangles changed by the flip.*

Proof. Let uv be the flipped edge, and let wvu and xuv be the triangles deleted by the flip, so wxu and xwv are the triangles created by the flip.

The angle opposite the edge uw is $\angle wvu$ before the flip, and $\angle wxu$ after the flip. As Figure 2.9 illustrates, because the open circumdisk of wvu contains x, the latter angle is greater than the former angle by the Inscribed Angle Theorem, a standard fact about circle geometry that was known to Euclid. Likewise, the flip increases the angles opposite wv, vx, and xu.

Each of the other two angles of the new triangles, $\angle xuw$ and $\angle wvx$, is a sum of two pre-flip angles that merge when uv is deleted. It follows that all six angles of the two post-flip triangles exceed the smallest of the four angles in which uv participates before the flip.

Suppose without loss of generality that the circumdisk of wxu is at least as large as the circumdisk of xwv, and that $\angle wxu \leq \angle uwx$, implying that $\angle wxu$ is acute. Because the open circumdisk of wvu contains x, it is larger than the circumdisk of wxu, as illustrated in Figure 2.9. It follows that the largest pre-flip circumdisk is larger than the largest post-flip circumdisk. □

In Section 4.3, we show that a Delaunay flip never increases the largest min-containment disk among the triangles changed by the flip. These local results imply a global optimality result.

Theorem 2.8. *Among all the triangulations of a point set, there is a Delaunay triangulation that maximizes the minimum angle in the triangulation, a Delaunay triangulation that* *minimizes the largest circumdisk, and a Delaunay triangulation that minimizes the largest min-containment disk.*

Proof. Each of these properties is locally improved when an edge that is not locally Delaunay is flipped—or at least not worsened, in the case of min-containment disks. There is at least one optimal triangulation $\mathcal{T}$. If $\mathcal{T}$ has an edge that is not locally Delaunay, flipping that edge produces another optimal triangulation. When the flip algorithm runs with $\mathcal{T}$ as its input, every triangulation it iterates through is optimal by induction, and by Proposition 2.5, that includes a Delaunay triangulation. □

Theorem 2.8 is not the strongest statement we can make, but it is easy to prove. With more work, one can show that *every* Delaunay triangulation of a point set optimizes these criteria. See Exercise 3 for details.

Unfortunately, the only optimality property of Theorem 2.8 that generalizes to Delaunay triangulations in dimensions higher than two is the property of minimizing the largest min-containment ball. However, the list of optimality properties in Theorem 2.8 is not complete. In the plane, the Delaunay triangulation maximizes the mean inradius of its triangles and minimizes a property called the *roughness* of a piecewise linearly interpolated function, which is the integral over the triangulation of the square of the gradient. Section 4.3 discusses criteria related to interpolation error for which Delaunay triangulations of any dimension are optimal.

Another advantage of the Delaunay triangulation arises in numerical discretizations of the Laplacian operator ∇^2. Solutions to Dirichlet boundary value problems associated with Laplace's equation $\nabla^2\varphi = 0$ satisfy a *maximum principle*: the maximum value of φ always occurs on the domain boundary. Ideally, an approximate solution found by a numerical method should satisfy a discrete maximum principle, both for physical realism and because it helps to prove strong convergence properties for the numerical method and to bound its error. A piecewise linear finite element discretization of Laplace's equation over a Delaunay triangulation in the plane satisfies a discrete maximum principle. Moreover, the stiffness matrix is what is called a *Stieltjes matrix* or an *M-matrix*, which implies that it can be particularly stable in numerical methods such as incomplete Cholesky factorization. These properties extend to three-dimensional Delaunay triangulations for some finite volume methods but, unfortunately, not for the finite element method.

2.7 The uniqueness of the Delaunay triangulation

The strength of a strongly Delaunay simplex is that it appears in *every* Delaunay triangulation of a point set. If a point set has multiple Delaunay triangulations, they differ only in their choices of simplices that are merely Delaunay. Hence, if a point set is *generic*—if it has no four cocircular points—it has only one Delaunay triangulation.

Let us prove these facts. Loosely speaking, the following proposition says that strongly Delaunay simplices intersect nicely.

Proposition 2.9. *Let σ be a strongly Delaunay simplex, and let τ be a Delaunay simplex. Then $\sigma \cap \tau$ is either empty or a shared face of both σ and τ.*

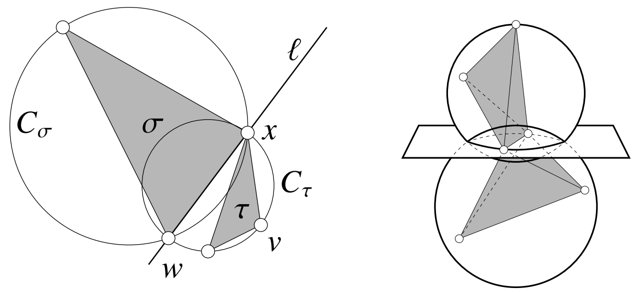

Figure 2.10: A strongly Delaunay simplex σ intersects a Delaunay simplex τ at a shared face of both. The illustration at right foreshadows the fact that this result holds in higher dimensions too.

Proof. If τ is a face of σ, the proposition follows immediately. Otherwise, τ has a vertex v that σ does not have. Because τ is Delaunay, it has an empty circumcircle C_τ. Because σ is strongly Delaunay, it has an empty circumcircle C_σ that does not pass through v, illustrated in Figure 2.10. But v lies on C_τ, so $C_\sigma \neq C_\tau$.

The intersection of circumcircles $C_\sigma \cap C_\tau$ contains zero, one, or two points. In the first two cases, the proposition follows easily, so suppose it is two points w and x, and let ℓ be the unique line through w and x. On one side of ℓ, an arc of C_σ encloses an arc of C_τ, and because C_σ is empty, no vertex of τ lies on this side of ℓ. Symmetrically, no vertex of σ lies on the other side of ℓ. Therefore, $\sigma \cap \tau \subset \ell$. It follows that $\sigma \cap \ell$ is either $\emptyset$, $\{w\}$, $\{x\}$, or the edge wx. The same is true of $\tau \cap \ell$ and, therefore, of $\sigma \cap \tau$. □

Proposition 2.9 leads us to see that if a point set has several Delaunay triangulations, they differ only by the simplices that are not strongly Delaunay.

Proposition 2.10. *Every Delaunay triangulation of a point set contains every strongly Delaunay simplex.*

Proof. Let $\mathcal{T}$ be any Delaunay triangulation of a point set S. Let σ be any strongly Delaunay simplex. Let p be a point in the relative interior of σ.

Some Delaunay simplex τ in $\mathcal{T}$ contains the point p. By Proposition 2.9, $\sigma \cap \tau$ is a shared face of σ and τ. But $\sigma \cap \tau$ contains p, which is in the relative interior of σ, so $\sigma \cap \tau = \sigma$. Therefore, σ is a face of τ, so $\sigma \in \mathcal{T}$. □

An immediate consequence of this proposition is that "most" point sets—at least, most point sets with randomly perturbed real coordinates—have just one Delaunay triangulation.

Theorem 2.11. *Let S be a point set. Suppose no four points in S lie on a common empty circle. Then S has one unique Delaunay triangulation.*

Proof. By Proposition 2.6, S has at least one Delaunay triangulation. Because no four points lie on a common empty circle, every Delaunay simplex is strongly Delaunay. By Proposition 2.10, every Delaunay triangulation of S contains every Delaunay simplex. By definition, no Delaunay triangulation contains a triangle that is not Delaunay. Hence, the

Delaunay triangulation is uniquely defined as the set of all Delaunay triangles and their faces. □

Theorem 2.11 does not preclude the possibility that all the vertices might be collinear. In that case, the vertices have a unique triangulation that has edges but no triangles and is vacuously Delaunay.

2.8 The weighted Delaunay triangulation

The parabolic lifting map connects Delaunay triangulations with convex hulls. It also suggests a generalization of Delaunay triangulations in which lifted vertices are not required to lie on the paraboloid. This observation is exploited by several mesh generation algorithms.

The simplest version of this idea begins with a planar point set S and assigns each point an arbitrary *height* to which it is lifted in $\mathbb{R}^3$. Imagine taking the convex hull of the points lifted to $\mathbb{R}^3$ and projecting its underside down to the plane, yielding a convex subdivision called the *weighted Delaunay subdivision*. If some of the faces of this subdivision are not triangular, they may be triangulated arbitrarily, and S has more than one *weighted Delaunay triangulation*.

In recognition of the special properties of the parabolic lifting map, it is customary to endow each point $v \in S$ with a scalar *weight* ω_v that represents how far its height deviates from the paraboloid. Specifically, v's height is $\|v\|^2 - \omega_v$, and its lifted companion is

$$v^+ = (v_x, v_y, v_x^2 + v_y^2 - \omega_v).$$

A positive weight thus implies that the point's lifted companion is below the paraboloid; a negative implies above. The reason weight is defined in opposition to height is so that increasing a vertex's weight will tend to increase its influence on the underside of the convex hull $\operatorname{conv} S^+$.

Given a weight assignment $\omega : S \rightarrow \mathbb{R}$, we denote the *weighted point set* $S[\omega]$. If the weight of a vertex v is so small that its lifted companion v^+ is not on the underside of $\operatorname{conv} S[\omega]^+$, as illustrated in Figure 2.11, it does not appear in the weighted Delaunay subdivision at all, and v is said to be *submerged* or *redundant*. Submerged vertices create some confusion of terminology, because a weighted Delaunay triangulation of S is not necessarily a triangulation of S—it might omit some of the vertices in S.

If every vertex in S has a weight of zero, every weighted Delaunay triangulation of S is a Delaunay triangulation of S. No vertex is submerged because every point on the paraboloid is on the underside of the convex hull of the paraboloid.

The weighted analog of a Delaunay simplex is called a *weighted Delaunay* simplex, and the analog of a circumdisk is a *witness* plane.

Definition 2.5 (weighted Delaunay triangulation; witness)**.** Let $S[\omega]$ be a weighted point set in $\mathbb{R}^3$. A simplex σ whose vertices are in S is *weighted Delaunay* if σ^+ is included in a downward-facing face of $\operatorname{conv} S[\omega]^+$. In other words, there exists a non-vertical plane $h_\sigma \subset \mathbb{R}^3$ such that $h_\sigma \supset \sigma^+$ and no vertex in $S[\omega]^+$ lies below h_σ. The plane h_σ is called a *witness* to the weighted Delaunay property of σ. A *weighted Delaunay triangulation* of

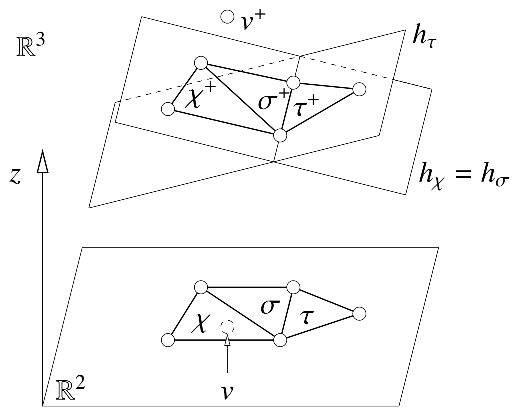

Figure 2.11: The triangles χ, σ, and τ are all weighted Delaunay, but only τ is strongly weighted Delaunay. Triangles χ and σ have the same witness plane $h_\chi = h_\sigma$, and τ has a different witness h_τ. The vertex v is submerged.

$S[\omega]$, denoted $\mathrm{Del}\, S[\omega]$, is a triangulation of a subset of S such that $|\mathrm{Del}\, S[\omega]| = \mathrm{conv}\, S$ and every simplex in $\mathrm{Del}\, S[\omega]$ is weighted Delaunay with respect to $S[\omega]$.

Figure 2.11 illustrates three weighted Delaunay triangles and their witnesses. All their edges and vertices are weighted Delaunay as well, but the submerged vertex v is not weighted Delaunay. A triangle has a unique witness, but an edge or vertex can have an infinite number of witnesses.

The weighted analog of a strongly Delaunay simplex is a *strongly weighted Delaunay* simplex.

Definition 2.6 (strongly weighted Delaunay)**.** A simplex σ is *strongly weighted Delaunay* if σ^+ is a downward-facing face of $\mathrm{conv}\, S[\omega]^+$ and no vertex in $S[\omega]^+$ lies on σ^+ except the vertices of σ^+. In other words, there exists a non-vertical plane $h_\sigma \subset \mathbb{R}^3$ such that $h_\sigma \supset \sigma^+$ and in $S[\omega]^+$ lies above h_σ, except the vertices of σ^+. The plane h_σ is a *witness* to the strongly weighted Delaunay property of σ.

Of the three triangles in Figure 2.11, only τ is strongly weighted Delaunay. All the edges are strongly weighted Delaunay except the edge shared by χ and σ. All the vertices are strongly weighted Delaunay except v.

Proposition 2.2 shows that if all the weights are zero, "weighted Delaunay" is equivalent to "Delaunay" and "strongly weighted Delaunay" is equivalent to "strongly Delaunay." If a simplex σ is weighted Delaunay, it appears in at least one weighted Delaunay triangulation of S. If σ is strongly weighted Delaunay, it appears in *every* weighted Delaunay triangulation of S (by a generalization of Proposition 2.10).

Definition 2.7 (generic)**.** A weighted point set $S[\omega]$ in $\mathbb{R}^2$ is *generic* if no four points in $S[\omega]^+$ lie on a common non-vertical plane in $\mathbb{R}^3$.

If a weighted point set is generic, then every weighted Delaunay simplex is strongly weighted Delaunay, and the point set has exactly one weighted Delaunay triangulation. For points with weight zero, this definition is equivalent to the statement that no four points are cocircular.

2.9 Symbolic weight perturbations

Some algorithms for constructing Delaunay triangulations, like the gift-wrapping algorithm described in Section 3.11, have difficulties triangulating point sets that have multiple Delaunay triangulations. These problems can be particularly acute in three or more dimensions. One way to make the points generic is to perturb them in space so that their Delaunay triangulation is unique. A better way is to assign the points infinitesimal weights such that they have one unique weighted Delaunay triangulation. Because the weights are infinitesimal, that triangulation is also an ordinary Delaunay triangulation.

To put this idea on firm mathematical ground, replace the infinitesimals with tiny, finite weights that are *symbolic*—their magnitudes are not explicitly specified. Given a point set $S = \{v_1, v_2, \ldots, v_n\}$, assign vertex v_i a weight of ϵ^i, where $\epsilon > 0$ is presumed to be so small that making it smaller would not change the weighted Delaunay triangulation. There is no need to compute an ϵ that satisfies this presumption; it is enough to know that such an ϵ exists.

An intuitive way to understand these weights is to imagine the result of a procedure that perturbs the vertex weights one at a time. Initially every vertex has a weight of zero, and the Delaunay subdivision may have some polygons that are not triangular. Perturb the weight of each vertex $v_1, v_2, \ldots$ in turn to subdivide the non-triangular polygons adjoining the perturbed vertex. The perturbation of vertex v_i's weight is chosen so that v_i^+ is not coplanar with any three other lifted vertices, and it is chosen sufficiently small that if v_i^+ was above the affine hull of three other lifted vertices, it remains above that affine hull after the perturbation. Therefore, every polygon adjoining v_i in the Delaunay subdivision after the perturbation must be a triangle, and no face not adjoining v_i is changed by the perturbation. Both these goals are achieved by choosing each weight perturbation to be infinitesimally smaller than all the previous weight perturbations—for instance, a weight of ϵ^i in the limit as $\epsilon > 0$ approaches zero.

Proposition 2.12. *Let S be a set of points in the plane. Let $\omega : S \to \mathbb{R}$ be the weight assignment described above. The following statements hold.*

- *If a simplex is strongly Delaunay with respect to S, it is strongly weighted Delaunay with respect to $S[\omega]$.*
- *If a simplex is strongly weighted Delaunay with respect to $S[\omega]$, it is Delaunay with respect to S.*
- *There is exactly one weighted Delaunay triangulation of $S[\omega]$, which is a Delaunay triangulation of S.*

Proof. Let σ be a simplex that is strongly Delaunay with respect to S. Then σ is strongly weighted Delaunay with respect to $S[0]$, and some plane $h \supset \sigma^+$ is a witness to that fact. If σ is a triangle, vertical perturbations of the vertices of σ^+ induce a unique perturbation of h. If σ is an edge or a vertex, choose one or two arbitrary points in h that are affinely independent of σ^+ and fix them so that a perturbation of σ^+ uniquely perturbs h.

Every vertex in S^+ lies above h except the vertices of σ^+, which lie on h. If ϵ is sufficiently small, the vertically perturbed points $S[\omega]^+$ preserve this property: every vertex in $S[\omega]^+$ lies above the perturbed witness plane for the perturbed σ^+, except the vertices of the perturbed σ^+. Therefore, σ is strongly weighted Delaunay with respect to $S[\omega]$ too, confirming the first statement of the proposition.

If a simplex σ is strongly weighted Delaunay with respect to $S[\omega]$, then every vertex in $S[\omega]^+$ lies above some witness plane $h \supset \sigma^+$ except the vertices of σ^+. If ϵ is sufficiently small, the vertically perturbed points S^+ nearly preserve this property: every vertex in S^+ lies above or on the perturbed witness plane for the perturbed σ^+. (If this were not so, some vertex would have moved from below the affine hull of three other vertices to above their affine hull; this can be prevented by making ϵ smaller.) This confirms the second statement of the proposition.

If ϵ is sufficiently small, no four vertices of $S[\omega]^+$ lie on a common non-vertical plane, so every face of the weighted Delaunay subdivision of $S[\omega]$ is a triangle, and the weighted Delaunay triangulation of $S[\omega]$ is unique. Every simplex of this triangulation is Delaunay with respect to S, so it is a Delaunay triangulation of S. $\square$

The converse of the second statement of Proposition 2.12 is not true: a simplex that is Delaunay with respect to S is not necessarily weighted Delaunay with respect to $S[\omega]$. This is not surprising; the purpose of the weight perturbations is to break coplanarities in S^+ and eliminate some of the Delaunay simplices so that the Delaunay triangulation is unique.

An important advantage of symbolic perturbations is that it is easy to simulate them in software—see Exercise 2 in Chapter 3—and they do not introduce the numerical problems associated with finite, numerical perturbations. Software for constructing Delaunay triangulations can ignore the symbolic perturbations until it encounters four cocircular vertices—if constructing a weighted Delaunay triangulation, four vertices whose lifted companions are coplanar. In that circumstance only, the software must simulate the circumstance where the four lifted vertices are perturbed so they are not coplanar.

2.10 Constrained Delaunay triangulations in the plane

As planar Delaunay triangulations maximize the minimum angle, do they solve the problem of triangular mesh generation? No, for two reasons illustrated in Figure 2.12. First, skinny triangles might appear anyway. Second, the Delaunay triangulation of a domain's vertices might not respect the domain's boundary. Both these problems can be solved by introducing additional vertices, as illustrated.

An alternative solution to the second problem is to use a *constrained Delaunay triangulation* (CDT). A CDT is defined with respect to a set of points and *segments* that demarcate the domain boundary. Every segment is required to become an edge of the CDT. The triangles of a CDT are *not* required to be Delaunay; instead, they must be *constrained Delaunay*, a property that partly relaxes the empty circumdisk property.

One virtue of a CDT is that it can respect arbitrary segments without requiring the insertion of any additional vertices besides the vertices of the segments. Another is that the CDT inherits the Delaunay triangulation's optimality: among all triangulations of a point set *that include all the segments*, the CDT maximizes the minimum angle, minimizes the largest circumdisk, and minimizes the largest min-containment disk.

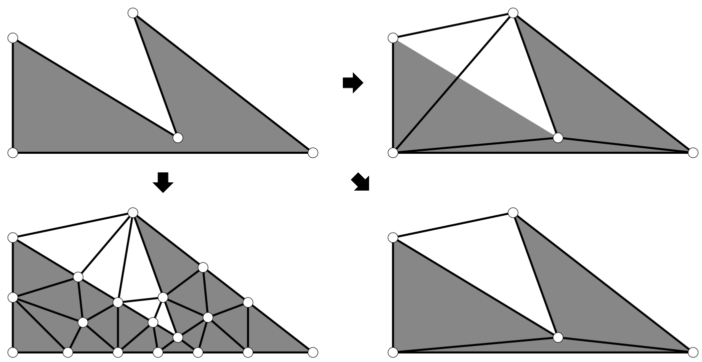

Figure 2.12: The Delaunay triangulation (upper right) may omit domain edges and contain skinny triangles. A Steiner Delaunay triangulation (lower left) can fix these faults by introducing new vertices. A constrained Delaunay triangulation (lower right) fixes the first fault without introducing new vertices.

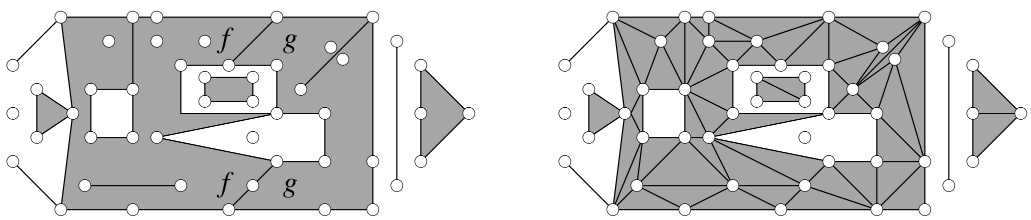

Figure 2.13: A two-dimensional piecewise linear complex and its constrained Delaunay triangulation. Each polygon may have holes, slits, and vertices in its interior.

2.10.1 Piecewise linear complexes and their triangulations

The domain over which a CDT is defined (and the input to a CDT construction algorithm) is not just a set of points; it is a complex composed of points, edges, and polygons, illustrated in Figure 2.13. The purpose of the edges is to dictate that triangulations of the complex must contain those edges. The purpose of the polygons is to specify the region to be triangulated. The polygons are linear 2-cells (recall Definition 1.7), which are not necessarily convex and may have holes.

Definition 2.8 (piecewise linear complex)**.** In the plane, a *piecewise linear complex* (PLC) $\mathcal{P}$ is a finite set of linear cells—vertices, edges, and polygons—that satisfies the following properties.

- The vertices and edges in $\mathcal{P}$ form a simplicial complex. That is, $\mathcal{P}$ contains both vertices of every edge in $\mathcal{P}$, and the relative interior of an edge in $\mathcal{P}$ intersects no vertex in $\mathcal{P}$ nor any other edge in $\mathcal{P}$.
- For each polygon f in $\mathcal{P}$, the boundary of f is a union of edges in $\mathcal{P}$.
- If two polygons in $\mathcal{P}$ intersect, their intersection is a union of edges and vertices in $\mathcal{P}$.

The edges in a PLC $\mathcal{P}$ are called *segments* to distinguish them from other edges in a triangulation of $\mathcal{P}$. The *underlying space* of a PLC $\mathcal{P}$, denoted $|\mathcal{P}|$, is the union of its contents; that is, $|\mathcal{P}| = \bigcup_{f \in \mathcal{P}} f$. Usually, the underlying space is the domain to be triangulated.[1]

Figure 2.13 shows a PLC and a triangulation of it. Observe that the intersection of the linear 2-cells f and g has multiple connected components, including two line segments and one isolated point, which are not collinear. The faces of the complex that represent this intersection are three edges and six vertices.

Every simplicial complex and every polyhedral complex is a PLC. But PLCs are more general, and not just because they permit nonconvex polygons. As Figure 2.13 illustrates, segments and isolated vertices can float in a polygon's interior; they constrain how the polygon can be triangulated. One purpose of these floating constraints is to permit the application of boundary conditions at appropriate locations in a mesh of a PLC.

Whereas the faces of a simplex are defined in a way that depends solely on the simplex, and the faces of a convex polyhedron are too, the faces of a polygon are defined in a fundamentally different way that depends on both the polygon and the PLC it is a part of. An edge of a polygon might be a union of several segments in the PLC; these segments and their vertices are faces of the polygon. A PLC may contain segments and edges that lie in the relative interior of a polygon; these are also considered to be faces of the polygon, because they constrain how the polygon can be subdivided into triangles.

Definition 2.9 (face of a linear cell)**.** The *faces* of a linear cell f (polygon, edge, or vertex) in a PLC $\mathcal{P}$ are the linear cells in $\mathcal{P}$ that are subsets of f, including f itself. The *proper faces* of f are all the faces of f except f.

A triangulation of $\mathcal{P}$ must cover every polygon and include every segment.

Definition 2.10 (triangulation of a planar PLC)**.** Let $\mathcal{P}$ be a PLC in the plane. A *triangulation of* $\mathcal{P}$ is a simplicial complex $\mathcal{T}$ such that $\mathcal{P}$ and $\mathcal{T}$ have the same vertices, $\mathcal{T}$ contains every edge in $\mathcal{P}$ (and perhaps additional edges), and $|\mathcal{T}| = |\mathcal{P}|$.

It is not difficult to see that a simplex can appear in a triangulation of $\mathcal{P}$ only if it *respects* $\mathcal{P}$. (See Exercise 4.)

Definition 2.11 (respect)**.** A simplex σ *respects* a PLC $\mathcal{P}$ if $\sigma \subseteq |\mathcal{P}|$ and for every $f \in \mathcal{P}$ that intersects σ, $f \cap \sigma$ is a union of faces of σ.

Proposition 2.13. *Every simple polygon has a triangulation. Every PLC in the plane has a triangulation too.*

Proof. Let P be a simple polygon. If P is a triangle, it clearly has a triangulation. Otherwise, consider the following procedure for triangulating P. Let $\angle uvw$ be a corner of P having an interior angle less than 180°. Two such corners are found by letting v be the lexicographically least or greatest vertex of P.

[1] If one takes the vertices and edges of a planar PLC and discards the polygons, one has a simplicial complex in the plane with no triangles. This complex is called a *planar straight line graph* (PSLG). Most publications about CDTs take a PSLG as the input and assume that the CDT should cover the PSLG's convex hull. PLCs are more expressive, as they can restrict the triangulation to a nonconvex region of the plane.

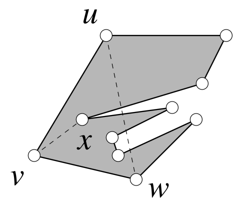

Figure 2.14: The edge vx cuts this simple polygon into two simple polygons.

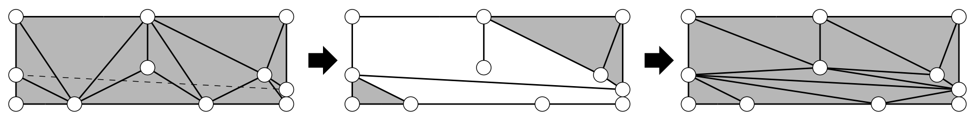

Figure 2.15: Inserting a segment into a triangulation.

If the open edge uw lies strictly in P's interior, then cutting uvw from P yields a polygon having one edge fewer; triangulate it recursively. Otherwise, uvw contains at least one vertex of P besides u, v, and w, as illustrated in Figure 2.14. Among those vertices, let x be the vertex farthest from the line aff uw. The open edge vx must lie strictly in P's interior, because if it intersected an edge of P, that edge would have a vertex further from aff uw. Cutting P at vx produces two simple polygons, each with fewer edges than P; triangulate them recursively. In either case, the procedure produces a triangulation of P.

Let $\mathcal{P}$ be a planar PLC. Consider the following procedure for triangulating $\mathcal{P}$. Begin with an arbitrary triangulation of the vertices in $\mathcal{P}$, such as the lexicographic triangulation described in Section 2.1. Examine each segment in $\mathcal{P}$ to see if it is already an edge of the triangulation. Insert each missing segment into the triangulation by deleting all the edges and triangles that intersect its relative interior, creating the new segment, and retriangulating the two polygonal cavities thus created (one on each side of the segment), as illustrated in Figure 2.15. The cavities might not be simple polygons, because they might have edges dangling in their interiors, as shown. But it is straightforward to verify that the procedure discussed above for triangulating a simple polygon works equally well for a cavity with dangling edges.

The act of inserting a segment never deletes another segment, because two segments in $\mathcal{P}$ cannot cross. Therefore, after every segment is inserted, the triangulation contains all of them. Finally, delete any simplices not included in $|\mathcal{P}|$. □

Definition 2.10 does not permit $\mathcal{T}$ to have vertices absent from $\mathcal{P}$, but mesh generation usually entails adding new vertices to guarantee that the triangles have high quality. This motivates the notion of a Steiner triangulation.

Definition 2.12 (Steiner triangulation of a PLC)**.** Let $\mathcal{P}$ be a PLC. A *Steiner triangulation of* $\mathcal{P}$, also known as a *conforming triangulation of* $\mathcal{P}$ or a *mesh of* $\mathcal{P}$, is a simplicial complex $\mathcal{T}$ such that $\mathcal{T}$ contains every vertex in $\mathcal{P}$ and possibly more, every edge in $\mathcal{P}$ is a union of edges in $\mathcal{T}$, and $|\mathcal{T}| = |\mathcal{P}|$. The new vertices in $\mathcal{T}$, absent from $\mathcal{P}$, are called *Steiner points*. A *Steiner Delaunay triangulation* of $\mathcal{P}$, also known as a *conforming Delaunay triangulation* of $\mathcal{P}$, is a Steiner triangulation of $\mathcal{P}$ in which every simplex is Delaunay.

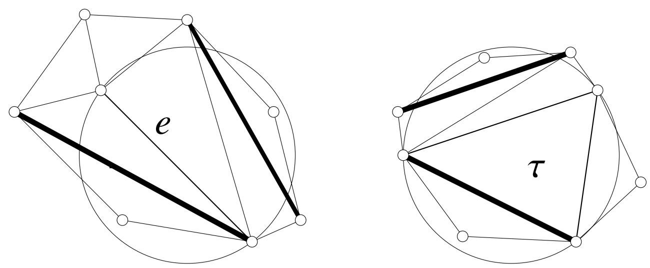

Figure 2.16: The edge e and triangle τ are constrained Delaunay. Bold lines represent segments.

2.10.2 The constrained Delaunay triangulation

Constrained Delaunay triangulations (CDTs) offer a way to force a triangulation to respect the edges in a PLC without introducing new vertices, while maintaining some of the advantages of Delaunay triangulations. However, it is necessary to relax the requirement that all triangles be Delaunay. The terminology can be confusing: whereas every Steiner Delaunay triangulation is a Delaunay triangulation (of some point set), constrained Delaunay triangulations generally are not.

Recall the Delaunay Lemma: a triangulation of a point set is Delaunay if and only if every edge is locally Delaunay. Likewise, there is a Constrained Delaunay Lemma (Section 2.10.3) that offers the simplest definition of a CDT: a triangulation of a PLC is constrained Delaunay if and only if every edge is locally Delaunay *or* a segment. Thus, a CDT differs from a Delaunay triangulation in three ways: it is not necessarily convex, it is required to contain the edges in a PLC, and those edges are exempted from being locally Delaunay.

The defining characteristic of a CDT is that every triangle is constrained Delaunay, as defined below.

Definition 2.13 (visibility)**.** Two points x and y are *visible* to each other if the line segment xy respects $\mathcal{P}$; recall Definition 2.11. We also say that x and y can *see* each other. A linear cell in $\mathcal{P}$ that intersects the relative interior of xy but does not include xy is said to *occlude* the visibility between x and y.

Definition 2.14 (constrained Delaunay)**.** In the context of a PLC $\mathcal{P}$, a simplex σ is *constrained Delaunay* if $\mathcal{P}$ contains the vertices of σ, σ respects $\mathcal{P}$, and there is an open circumdisk of σ that contains no vertex in $\mathcal{P}$ that is visible from a point in the relative interior of σ.

Figure 2.16 illustrates examples of a constrained Delaunay edge e and a constrained Delaunay triangle τ. Bold lines indicate PLC segments. Although e has no empty circumdisk, the depicted open circumdisk of e contains no vertex that is visible from the relative interior of e. There are two vertices in the disk, but both are hidden behind segments. Hence, e is constrained Delaunay. Similarly, the open circumdisk of τ contains two vertices, but both are hidden from the interior of τ by segments, so τ is constrained Delaunay.

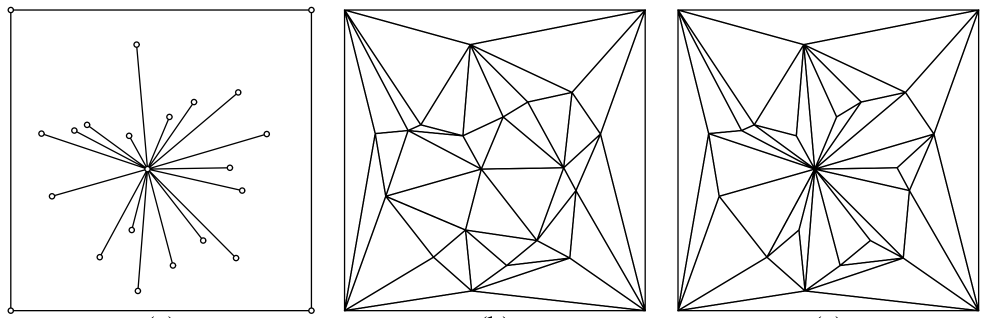

Figure 2.17: (a) A piecewise linear complex. (b) The Delaunay triangulation of its vertices. (c) Its constrained Delaunay triangulation.

Definition 2.15 (constrained Delaunay triangulation)**.** A *constrained Delaunay triangulation* (CDT) of a PLC $\mathcal{P}$ is a triangulation of $\mathcal{P}$ in which every triangle is constrained Delaunay.

Figure 2.17 illustrates a PLC, a Delaunay triangulation of its vertices, and a constrained Delaunay triangulation of the PLC. In the CDT, every triangle is constrained Delaunay, every edge that is not a PLC segment is constrained Delaunay, and every vertex is trivially constrained Delaunay.

CDTs and Steiner Delaunay triangulations are two different ways to force a triangulation to conform to the boundary of a geometric domain. CDTs partly sacrifice the Delaunay property for the benefit of requiring no new vertices. For mesh generation, new vertices are usually needed anyway to obtain good triangles, so many Delaunay meshing algorithms use Steiner Delaunay triangulations. But some algorithms use a hybrid of CDTs and Steiner Delaunay triangulations because it helps to reduce the number of new vertices. A *Steiner CDT* or *conforming CDT* of $\mathcal{P}$ is a Steiner triangulation of $\mathcal{P}$ in which every triangle is constrained Delaunay.

2.10.3 Properties of the constrained Delaunay triangulation

For every property of Delaunay triangulations discussed in this chapter, there is an analogous property of constrained Delaunay triangulations. This section summarizes them. Proofs are omitted, but each of them is a straightforward extension of the corresponding proof for Delaunay triangulations.

The Delaunay Lemma generalizes to CDTs, and provides a useful alternative definition: a triangulation of a PLC $\mathcal{P}$ is a CDT if and only if every one of its edges is locally Delaunay *or* a segment in $\mathcal{P}$.

Lemma 2.14 (Constrained Delaunay Lemma)**.** *Let $\mathcal{T}$ be a triangulation of a PLC $\mathcal{P}$. The following three statements are equivalent.*

- *Every triangle in $\mathcal{T}$ is constrained Delaunay (i.e. $\mathcal{T}$ is constrained Delaunay).*
- *Every edge in $\mathcal{T}$ not in $\mathcal{P}$ is constrained Delaunay.*

- *Every edge in $\mathcal{T}$ not in $\mathcal{P}$ is locally Delaunay.*

One way to construct a constrained Delaunay triangulation of a PLC $\mathcal{P}$ is to begin with any triangulation of $\mathcal{P}$. Apply the flip algorithm, modified so that it never flips a segment: repeatedly choose any edge of the triangulation that is not in $\mathcal{P}$ and not locally Delaunay, and flip it. When no such edge survives, the Constrained Delaunay Lemma tells us that the triangulation is constrained Delaunay.

Proposition 2.15. *Given a triangulation of a PLC having n vertices, the modified flip algorithm (which never flips a PLC segment) terminates after $O(n^2)$ edge flips, yielding a constrained Delaunay triangulation.*

Proposition 2.16. *Every PLC has a constrained Delaunay triangulation.*

The CDT has the same optimality properties as the Delaunay triangulation, except that the optimality is with respect to a smaller set of triangulations—those that include the PLC's edges.

Theorem 2.17. *Among all the triangulations of a PLC, every constrained Delaunay triangulation maximizes the minimum angle in the triangulation, minimizes the largest circumdisk, and minimizes the largest min-containment disk.*

A sufficient but not necessary condition for the CDT to be unique is that no four vertices are cocircular.

Theorem 2.18. *If a PLC is generic—no four of its vertices lie on a common circle—then the PLC has one unique constrained Delaunay triangulation, which contains every constrained Delaunay simplex.*

2.11 Notes and exercises

Delaunay triangulations and the Delaunay Lemma were introduced by Boris Delaunay's seminal 1934 paper [69]. The relationship between Delaunay triangulations and convex hulls was discovered by Brown [34], who proposed a different lifting map that projects the points onto a sphere. The parabolic lifting map of Seidel [190, 90] is numerically better behaved than the spherical lifting map.

The flip algorithm, the incremental insertion algorithm, and the Delaunay triangulation's property of maximizing the minimum angle were all introduced in a classic paper by Charles Lawson [131]. D'Azevedo and Simpson [67] show that two-dimensional Delaunay triangulations minimize the largest circumdisk and the largest min-containment disk. Lambert [129] shows that the Delaunay triangulation maximizes the mean inradius (equivalently, the sum of inradii) of its triangles. Rippa [175] shows that it minimizes the roughness (defined in Section 2.6) of a piecewise linearly interpolated function, and Powar [172] gives a simpler proof. Ciarlet and Raviart [63] show that for Dirichlet boundary value problems on Laplace's equation, finite element discretizations with piecewise linear elements over a triangulation in the plane with no obtuse angles have solutions that satisfy a discrete maximum principle. The result extends easily to all Delaunay triangulations in the plane, even ones with obtuse angles, but it is not clear who first made this observation. Miller, Talmor, Teng, and Walkington [148] extend the result to a finite volume method that uses a three-dimensional Delaunay triangulation and its Voronoi dual, with Voronoi cells as control volumes.

The symbolic weight perturbation method of Section 2.9 originates with Edelsbrunner and Mücke [89, Section 5.4].

Constrained Delaunay triangulations in the plane were mathematically formalized by Lee and Lin [133] in 1986, though algorithms that unwittingly construct CDTs appeared much earlier [97, 159]. Lee and Lin extend to CDTs Lawson's proof that Delaunay triangulations maximize the minimum angle.

Exercises

1. Draw the Delaunay triangulation of the following point set.

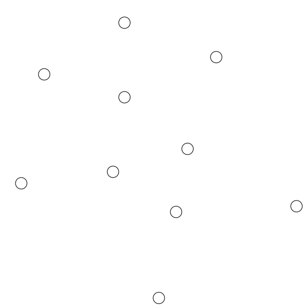

2. Let P and Q be two disjoint point sets in the plane. (Think of them as a red point set and a black point set.) Let $p \in P$ and $q \in Q$ be two points from these sets that minimize the Euclidean distance $d(p,q)$. Prove that pq is an edge of $\mathrm{Del}\,(P \cup Q)$. This observation leads easily to an $O(n \log n)$-time algorithm for finding p and q, the red-black closest pair.

3. Let S be a point set in the plane. S may have subsets of four or more cocircular points, so S may have many Delaunay triangulations.

   (a) Prove that it is possible to transform a triangulation of a convex polygon to any other triangulation of the same polygon by a sequence of edge flips.

   (b) Prove that it is possible to flip from any Delaunay triangulation of S to any other Delaunay triangulation of S, such that every intermediate triangulation is also Delaunay.

   (c) Prove that *every* Delaunay triangulation of S maximizes its minimum angle—there is no triangulation of S whose smallest angle is greater.

4. Show that a simplex can appear in a triangulation of a PLC $\mathcal{P}$ (Definition 2.10) only if it respects $\mathcal{P}$ (Definition 2.11).

5. Recall that every Delaunay triangulation of a point set contains every strongly Delaunay edge, but there is no such guarantee for Delaunay edges that are not strongly Delaunay. Show constructively that for any PLC $\mathcal{P}$, *every* constrained Delaunay edge is in at least one CDT of $\mathcal{P}$. Hint: See Exercise 3(b).

6. Prove Lemma 2.14, the Constrained Delaunay Lemma.

7. Recall that a triangle τ is *constrained Delaunay* with respect to a PLC $\mathcal{P}$ if its vertices are in $\mathcal{P}$, it respects $\mathcal{P}$, and the open circumdisk of τ contains no vertex in $\mathcal{P}$ that is visible from a point in τ's interior.

   Let τ be a triangle that satisfies the first two of those three conditions. Let q be a point in the interior of τ. Prove that if no vertex of $\mathcal{P}$ in τ's open circumdisk is visible from q, then no vertex of $\mathcal{P}$ in τ's open circumdisk is visible from *any* point in the interior of τ, so τ is constrained Delaunay.