Weighted Delaunay Triangulations of Piecewise-flat Surfaces

by

Yuliya Gorlina

A Dissertation Submitted to the Faculty of the

Department of Mathematics

In Partial Fulfillment of the Requirements
For the Degree of

Doctor of Philosophy

In the Graduate College

The University of Arizona

2 0 1 1

The University of Arizona
Graduate College

As members of the Dissertation Committee, we certify that we have read the dissertation prepared by Yuliya Gorlina
entitled Weighted Delaunay Triangulations of Piecewise-flat Surfaces
and recommend that it be accepted as fulfilling the dissertation requirement for the Degree of Doctor of Philosophy

_____ Date: 9/26/2011
David Glickenstein

_____ Date: 9/26/2011
Douglas Pickrell

_____ Date: 9/26/2011
Kevin Lin

_____ Date: 9/26/2011
Shankar Venkataramani

Final approval and acceptance of this dissertation is contingent upon the candidate's submission of the final copies of the dissertation to the Graduate College.
I hereby certify that I have read this dissertation prepared under my direction and recommend that it be accepted as fulfilling the dissertation requirement.

_____ Date: 9/26/2011
Dissertation Director: David Glickenstein

Statement by Author

This dissertation has been submitted in partial fulfillment of requirements for an advanced degree at The University of Arizona and is deposited in the University Library to be made available to borrowers under rules of the Library.

Brief quotations from this dissertation are allowable without special permission, provided that accurate acknowledgment of source is made. Requests for permission for extended quotation from or reproduction of this manuscript in whole or in part may be granted by the head of the major department or the Dean of the Graduate College when in his or her judgment the proposed use of the material is in the interests of scholarship. In all other instances, however, permission must be obtained from the author.

Signed: Yuliya Gorlina

Acknowledgements

I am grateful for the patience and support of my advisor, David Glickenstein, without whom this work would not have been possible.

I thank my other mentors and teachers, including Doug Ulmer, Joe Watkins, Doug Pickrell, Dimitri Psaltis, Ken McLaughlin, and Tom Kennedy for their roles in my graduate education; Paul Asimow, Vladimir Baranovsky, Joe Maher, and Nathan Dunfield for supporting me during my time as an undergraduate at Caltech; and my teachers at Niles North High School for their encouragement.

For helpful and inspiring conversations, I thank Tom Lagatta. For reading my dissertation and providing feedback, I thank my friends Danny Maienschein, Dave Savitt, Matt Pennybacker, Verónica Marino, and my brother Andrey Gorlin. For providing me the emotional support I needed to stay sane during this process, I thank my friends Victoria Frost, Matt Johnson, Chris Jewell, Nellie Gopaul, JD Salazar, Jen Balakrishnan, Lisa Berger, and my sister Yelena Gorlin. I would also like to acknowledge my former roommates and officemates.

During my time at the University of Arizona, I was supported by VIGRE and GTEAMS fellowships, provided by the National Science Foundation under Grants #0602173 and #0841234, respectively.

Finally, I would like to thank my parents and extended family.

Dedication

Roman Shostak 1949-2011

Table of Contents

List of Figures

Abstract

Given a triangulated piecewise-flat surface and a function on the vertices we can define the Dirichlet energy, which is related to the Dirichlet energy of a smooth function. To find the Delaunay triangulation of the piecewise flat surface, we modify the triangulation by a sequence of edge flips, called the flip algorithm, which transform an edge which is not Delaunay into one that is Delaunay. It is known that the flip algorithm works in the plane as well as for a piecewise-flat surface, where we have to ensure that only finitely many triangulations are possible.

When the vertices of a piecewise-flat surface have weights, we want to find the weighted Delaunay triangulation using a flip algorithm. In this dissertation, we prove that the maximum edge length during the algorithm is bounded, which guarantees that there are finitely many triangulations. Thus the flip algorithm terminates and the resulting triangulation is weighted Delaunay.

Additionally, we give a new way to find what we call the relaxed weighted Delaunay on a flat surface.

1. Introduction

1.1. Background

Consider a finite set of points in $\mathbb{R}^2$, which we will call the vertices. We split the plane into regions of influence of the vertices by assigning a point x to a vertex v if x is closer to v than to any other vertex. These regions form the Voronoi diagram of the set of vertices. The dual of the Voronoi diagram (say, in the graph-theoretic sense) creates a

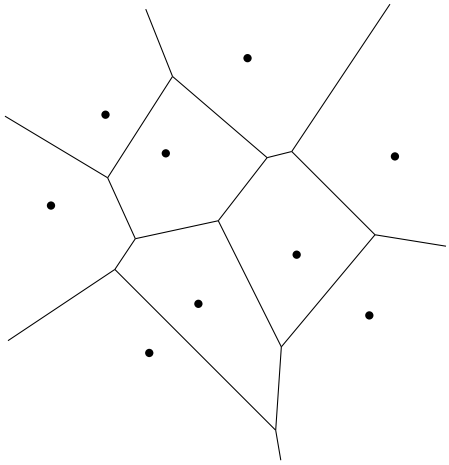

Figure 1.1. A Voronoi diagram.

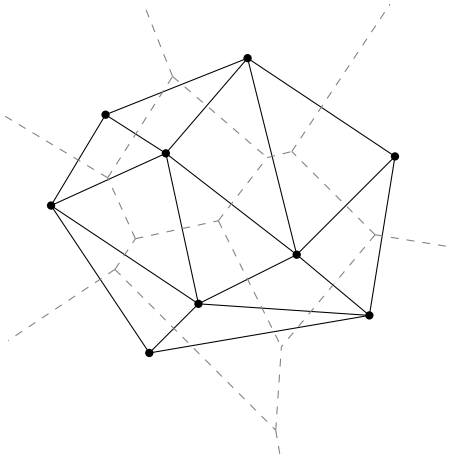

Figure 1.2. The Delaunay triangulation dual to the Voronoi diagram.

triangulation of the vertices called the Delaunay triangulation. Voronoi diagrams are named after Georges Voronoi, who was the first to consider the Delaunay triangulation as the dual of the Voronoi diagram [Vor08].

One of Voronoi's students, Boris Delaunay (also spelled Delone), was the first to use circumcircles to define Delaunay triangulations [Del34]. Recall that the circumcircle of a triangle is the unique circle which passes through its vertices. A triangulation is said to be Delaunay if none of the circumcircles of the triangles contain any vertices in their interiors. We would like to find the Delaunay triangulation without relying

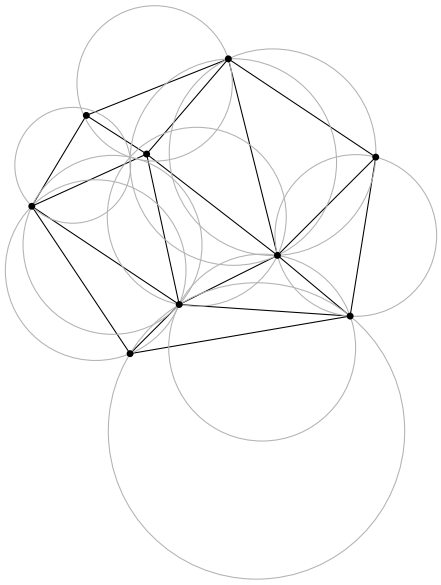

FIGURE 1.3. The Delaunay triangulation and the circumcircles.

on the Voronoi diagram. If we start with an arbitrary triangulation, we can modify it via a flip. We remove two triangles adjacent to an edge and forming a convex quadrilateral, and replace them by the triangles formed using the other diagonal of the quadrilateral. The procedure of repeatedly choosing an edge which is not

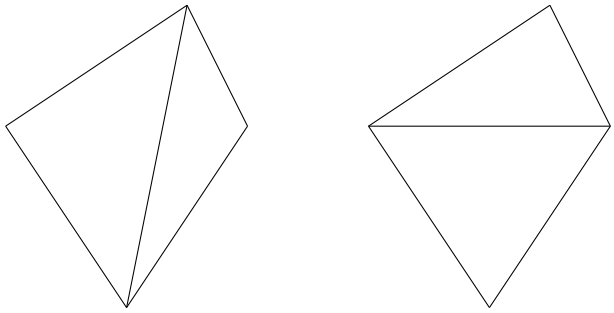

FIGURE 1.4. An illustration of a flip.

Delaunay and flipping it is called a flip algorithm. Lawson [Law72] was the first to use a flip algorithm to find the Delaunay triangulation in $\mathbb{R}^2$.

One of the ways we can generalize the concept of the Delaunay triangulation in $\mathbb{R}^2$ is to apply it to piecewise-flat surfaces, which was first done by Rivin [Riv94], who

also used a flip algorithm to find it.

Weighted Delaunay triangulations are another generalization of Delaunay triangulations in which vertices are equipped with weights. These weights may be thought of as modifying the vertices' regions of influence which form the (weighted) Voronoi diagram; see Section 3.1 for a precise definition. Circle packings can be viewed as a type of weighted Delaunay triangulations with positive weights, where the weights correspond to squared radii [Ste05]. Note that some definitions of weighted Delaunay triangulations allow deleting of vertices from the original vertex set; in our definition of the weighted Delaunay triangulation, all vertices must be used.

1.2. Applications

Gauss and Dirichlet were the first to use Voronoi diagrams in mathematics in the context of quadratic forms. Outside of mathematics, Delaunay triangulations and Voronoi diagrams have applications in various fields of natural science, such as crystallography and molecular physics. In addition, algorithms for computing the Delaunay triangulations and Voronoi diagrams are of interest in computer science. See survey papers [Aur91, AK00] for more details. Weighted Voronoi diagrams are used in the study of equal-volume foams [KS96], while weighted Delaunay triangulations are used in modeling molecules with weighted α-shapes [EKS83, EM92].

More recently, Delaunay triangulations have been used to define an intrinsic discrete Laplace operator on a piecewise-flat surface [BS07]. Given any triangulation we can define a discrete Laplace operator as follows,

$$(\Delta f)(a) = \sum_{e_{ab}} \frac{\ell^*_{ab}}{\ell_{ab}}(f(b) - f(a)),$$

where f is a function on the vertices, ℓ_{ab} is the length of the edge e_{ab}, and ℓ^*_{ab} is the dual of the edge e_{ab} with respect to weights (which can be set to zero in the case

without weights). We can then ask whether the operator has the usual properties that we expect of a Laplace operator:

- Δ is a negative semidefinite operator.
- Δ satisfies the weak maximum principle.

One way to ensure that these properties are satisfied is by ensuring that we are using a weighted Delaunay triangulation [Gli11, Prop. 39], and therefore it is useful to be able to find a weighted Delaunay triangulation. A discrete Laplacian is used in studying conformal variation of piecewise-flat manifolds [CGY11, Gli11].

1.3. Results

It is known that on piecewise-flat surfaces starting with an arbitrary triangulation, a flip algorithm can be used to find the Delaunay triangulation. See, for example, [BS07], in which Bobenko and Springborn show two facts:

(a) The algorithm does not get stuck.

(b) The algorithm terminates after a finite number of steps.

We say that the algorithm is stuck when all the edges that are not Delaunay cannot be flipped. In fact, we know that the algorithm will get stuck as soon as we encounter just one such edge. Our goal is to generalize these results to the case of weighted Delaunay triangulations on piecewise-flat surfaces. The natural generalization of (a) is false; in fact, even in the plane, the weighted Delaunay triangulation may not exist which results in an obstruction to (a). We prove the following generalization of (b) for weighted triangulations of piecewise-flat surfaces:

Theorem A (See Theorem 32)**.** The flip algorithm either terminates after a finite number of steps or gets stuck.

This follows from the following lemma together with the fact that there are finitely many triangulations with all edge lengths smaller than a fixed constant.

Lemma B (See Lemma 29)**.** There exists a constant M, which depends only on the piecewise flat manifold and the weights, so that if the longest edge is longer than M, then it is not weighted Delaunay.

We also prove that in the case of piecewise-flat surfaces of non-negative curvature, as well as in the plane, (a) generalizes. We conjecture that this is true for all piecewise-flat surfaces.

Theorem C (See Theorem 31)**.** If the weighted Delaunay triangulation exists, then the flip algorithm does not get stuck.

When a weighted Delaunay triangulation does not exist, one may consider a relaxed weighted Delaunay triangulation, which is a weighted Delaunay triangulation with the largest possible subset of the vertices. To find such a triangulation, we use bistellar flips, which are more general than the flip described above (see page 42). We show the following.

Proposition D (See Propositions 23 and 34)**.** In the flat case, there is a sequence of bistellar flips which reaches the relaxed weighted Delaunay triangulation.

This improves on a result of Edelsbrunner-Shah [ES96] in two dimensions; their algorithm involves adding vertices incrementally.

The non-flat case presents additional subtlety since vertices with curvature cannot be removed. We hope to address this in future work.

2. Delaunay triangulations

2.1. Delaunay triangulations in the plane

Consider a finite set of points V in the Euclidean plane denoted by $\mathbb{R}^2$. Assume that no three points in V are collinear. Let S be the convex hull of V, and consider the pair (S, V). In order to define a triangulation of (S, V), we use a simplicial complex $\mathcal{T}$ with the vertex set V, as in [Thu97]. Recall that an edge, or 1-simplex, e_{ab}, is the convex hull of two distinct vertices $a, b \in V$; a triangle, or 2-simplex, f_{abc}, is the convex hull of three distinct vertices $a, b, c \in V$. Let $\mathcal{T}_i$ denote the set of i-simplices of $\mathcal{T}$, i.e., $\mathcal{T}_2$ is the set of triangles, $\mathcal{T}_1$ is the set of edges, and $\mathcal{T}_0 = V$ is the set of vertices.

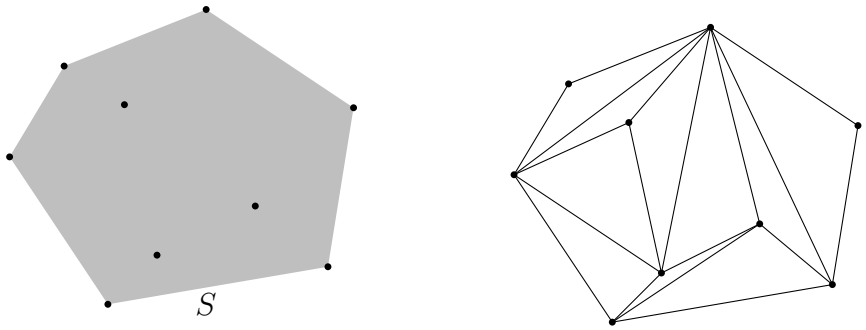

Figure 2.1. The pair (S, V) and a triangulation $\mathcal{T}$.

Definition 1. Consider a simplicial complex $\mathcal{T}$ with $\mathcal{T}_0 = V$. We say that $\mathcal{T}$ is a *triangulation* of (S, V) if

$$\bigcup_{f_{abc} \in \mathcal{T}_2} f_{abc} = S.$$

The circumcircle of f_{abc}, i.e., the unique circle through the points a, b, and c, is the boundary of an open disk, which we will denote by B_{abc}.

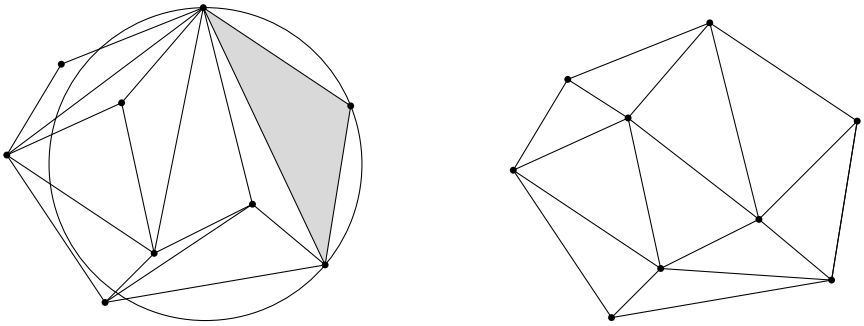

FIGURE 2.2. The triangulation on the left is not Delaunay, while the one on the right is Delaunay.

Definition 2. A triangulation $\mathcal{T}$ of (S, V) is said to be *Delaunay* if for each triangle $f_{abc} \in \mathcal{T}_2$, we have $B_{abc} \cap V = \varnothing$.

Remark 1. In a generic case, any three vertices define a unique circle, and given four vertices, the circle through any three of them does not intersect the fourth. So, generically, $\partial B_{abc} \cap V = \{a, b, c\}$. When V is not generic, we say it is degenerate; then it is possible that $\partial B_{abc} \cap V$ contains more than three vertices.

A triangulation consisting of just one triangle is automatically Delaunay. A triangulation consisting of two triangles is Delaunay if the circumcircle of each triangle does not contain the remaining vertex in its interior. In a larger triangulation, we can check whether the triangles adjacent to each edge form a triangulation which is Delaunay.

Definition 3. Consider an edge $e_{ab} \in \mathcal{T}_1$, with vertices $a, b \in V$. The *hinge* of e_{ab} is the pair $(e_{ab}, \mathcal{T}_{ab})$ of the edge e_{ab} and the subcomplex of $\mathcal{T}$ which contains the triangles which contain e_{ab}. A hinge is said to be *convex* if the interior angles of the underlying quadrilateral are all less than π.

Definition 4. An edge $e_{ab} \in \mathcal{T}_1$ is said to be *Delaunay* if the triangulation of the hinge at e_{ab} is Delaunay.

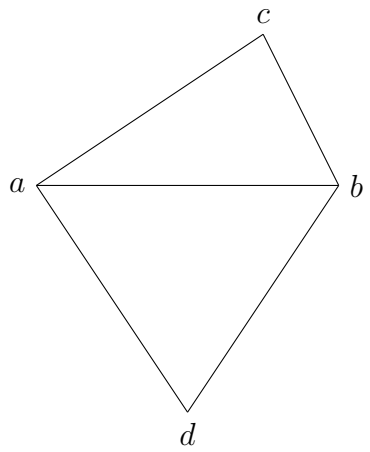

FIGURE 2.3. Example of a hinge.

Remark 2. This means that every boundary edge is automatically Delaunay, since it belongs to only one triangle. An interior edge e_{ab} adjacent to f_{abc} and f_{abd} is Delaunay if and only if the circumcircle of f_{abc} does not contain the vertex d in its interior. (When d lies on the circumcircle of f_{abc}, we say that the hinge of e_{ab} is degenerate, i.e., not generic. However, such a hinge is still Delaunay.) These two conditions can be taken as the definition of a (locally) Delaunay edge ([Ede01]). The latter is also called the circle criterion ([Rip90, Pow92]).

There is a very simple way to check whether an edge is Delaunay given in the following lemma.

Lemma 1. *Let e_{ab} be an interior edge of triangles f_{abc} and f_{abd}. Let θ_c and θ_d denote the angles at c in f_{abc} and at d in f_{abd}, respectively. The edge e_{ab} is Delaunay if and only if*

$$\theta_c + \theta_d \leq \pi,$$

with equality only if the hinge at e_{ab} is degenerate.

Proof. Notice that if $\theta_c + \theta_d = \pi$, then the hinge of the edge e_{ab} is a cyclic quadrilateral, i.e., it can be inscribed in a circle. Thus the circumcircles of f_{abc} and f_{abd}

coincide and the hinge is degenerate.

From here it is easy to see that d is outside the circumcircle of f_{abc} if and only if $\theta_c + \theta_d < \pi$. □

It turns out that if all the edges in a triangulation are Delaunay, then the entire triangulation is also Delaunay.

Lemma 2 (Delaunay [Del34])**.** *A triangulation is Delaunay if and only if each edge is Delaunay.*

To prove this lemma, we will use the notion of a power of a point with respect to a circle.

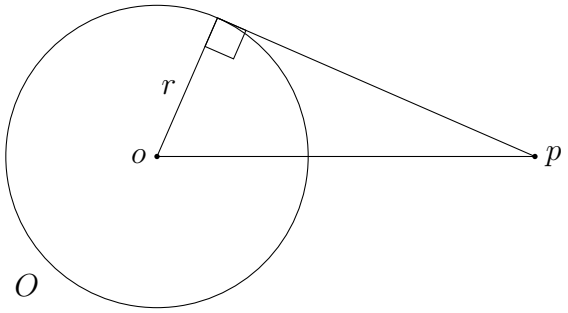

FIGURE 2.4. A geometric interpretation of the power of the point p with respect the circle O, when p is outside of O.

Definition 5. The *power* of a point $p \in \mathbb{R}^2$ with respect to a circle O with center o and radius r is

$$\pi_O(p) = ||p - o||^2 - r^2,$$

where $|| \cdot ||$ is the standard norm in $\mathbb{R}^2$.

The power is positive, zero, or negative if p lies outside, on, or inside the circle O, respectively.

Remark 3. We can also think of a circle as a point (representing the center), and a weight (representing the radius), which we will do in Chapter 3. Instead of using just

the radius as the weight, we use radius squared. However, we will also allow weights to be negative, in which case they cannot be interpreted as circles.

Let $\pi_{abc}(p)$ denote the power of a point $p \in \mathbb{R}^2$ with respect to the circumcircle of f_{abc}. Given two adjacent triangles f_{abc} and f_{abd}, as long as the hinge of e_{ab} is generic, the line containing the edge e_{ab} is exactly the set $\{p \mid \pi_{abc}(p) = \pi_{abd}(p)\}$, so we get the following lemma, which is a special case of the more general result in [Ede90].

Lemma 3. *Suppose an interior edge e_{ab} is Delaunay and adjacent to triangles f_{abc} and f_{abd}. If a ray starting at $p \in \mathbb{R}^2$ and crossing the interior of e_{ab} passes through f_{abc} first and through f_{abd} second, then*

$$\pi_{abd}(p) \geq \pi_{abc}(p),$$

with equality if and only if the hinge of e_{ab} is degenerate.

We can now prove Lemma 2. We give the complete proof because we will generalize it to piecewise-flat surfaces and to weighted Delaunay triangulations.

Proof of Lemma 2. We closely follow the argument given in [Ede01].

It is clear that if the triangulation is Delaunay, then each edge is Delaunay.

Suppose each edge is Delaunay.

Consider a triangle f_{abc} and a vertex d. We will show that the point d lies outside B_{abc}. If d is a vertex of f_{abc}, then it lies on the boundary, and so outside of, B_{abc}. Thus we only need to consider the case when d is distinct from a, b, c.

Choose a point $p \in f_{abc}$ so that the line segment from d to p contains no vertices of V in its interior. Let $\tau_0 = f_{abc}, \tau_1, \ldots, \tau_k$ be the sequence of triangles that intersect the segment pd. Let $\pi_i(d)$ be the power of d with respect to the circumcircle of τ_i. Since all the edges are Delaunay, we have

$$\pi_0(d) \geq \pi_1(d) \geq \cdots \geq \pi_k(d).$$

Since d is one of the vertices of τ_k, $\pi_k(d) = 0$ and so $\pi_0(d) \geq 0$, which means that d does not lie inside B_{abc}. □

2.1.1. Voronoi diagrams and duals

Delaunay triangulations can also be viewed as duals of Voronoi diagrams.

Definition 6. Given $V \subset \mathbb{R}^2$ a finite set of points, define the *Voronoi region* $\mathcal{V}_a$ of $a \in V$ as the set of points $p \in \mathbb{R}^2$ that are at least as close to a as to any other point in V; that is

$$\mathcal{V}_a = \{p \in \mathbb{R}^2 \mid ||p - a|| \leq ||p - b||, \forall b \in V\}.$$

The Voronoi regions together with their shared edges and vertices form the *Voronoi diagram* of V.

Alternatively, we can define $\mathcal{V}_a$ using half-planes H_{ab} of points that are at least as close to a as to b:

$$H_{ab} = \{p \in \mathbb{R}^2 \mid ||p - a|| \leq ||p - b||\}.$$

Then we get

$$\mathcal{V}_a = \bigcap_{b \neq a} H_{ab}.$$

We can recover the Delaunay triangulation dual to the Voronoi diagram by first connecting two vertices $a, b \in V$, if their Voronoi regions intersect along a line segment. If any non-triangular faces result, subdivide them to get a Delaunay triangulation. This may be necessary when more than three Voronoi cells meet at a point because more than three vertices lie on the same circle, i.e., when V is degenerate. Note that when more than three Voronoi cells meet at a point, the Delaunay triangulation is not unique.

Conversely, we can get the dual Voronoi diagram from the Delaunay triangulation. For every edge e_{ab} with adjacent faces f_{abc} and f_{abd}, the line segment between the

circumcenters C_{abc} and C_{abd} of f_{abc} and f_{abd} respectively will be the intersection of the Voronoi regions of a and b. In the case of a boundary edge e_{ab} of f_{abc}, we just extend the ray from C_{abc} to infinity along the perpendicular bisector of e_{ab}.

In addition to the Voronoi diagram being dual to the Delaunay triangulation, it is useful to define the dual of any triangulation.

Definition 7. Define $h_{ab,c}$ as the *signed distance* from the circumcenter C_{abc} of f_{abc} to the midpoint C_{ab} of the edge e_{ab}. It is equal in absolute value to the distance from C_{ab} to C_{abc} and is defined to be positive if C_{abc} is in the same half-plane determined by e_{ab} as c is, and negative if C_{abc} and c lie on opposite sides of e_{ab}. Given an interior edge e_{ab} that belongs to two triangles f_{abc} and f_{abd}, the *dual edge* is the line segment connecting C_{abc} and C_{abd} with a signed length equal to $h_{ab,c} + h_{ab,d}$.

As a consequence of Lemma 1, we get the following proposition.

Proposition 4. *The signed length of the dual edge is nonnegative if and only if the edge e_{ab} is Delaunay.*

2.1.2. Lifting triangulations

There is a relationship between Delaunay triangulations in $\mathbb{R}^2$ and convex hulls in $\mathbb{R}^3$.

Definition 8. Given a point $p = (x, y) \in \mathbb{R}^2$, define its *lifted point* $p^+ = (x, y, z) \in \mathbb{R}^3$ where $z = x^2 + y^2$. Similarly, given a set V, define $V^+ = \{p^+ \mid p \in V\}$. Finally, a facet of a convex polytope in $\mathbb{R}^3$ will be called a *lower facet* if the plane that contains it is nonvertical and the polytope lies vertically above this plane.

Lemma 5. *An interior edge e_{ab} is Delaunay if and only if the lifted hinge of e_{ab} is concave up or flat.*

Note that the lifted hinge of e_{ab} is flat when the hinge of e_{ab} is degenerate.

Proof. We closely follow the argument given in [Ede01].

Let f_{abc} and f_{abd} be the triangles adjacent to e_{ab}. The edge e_{ab} is not Delaunay if and only if $d \in B_{abc}$. The lifted hinge is strictly concave down (not flat) if and only if d^+ is strictly below the plane through a^+, b^+, c^+.

Let P be the plane through a^+, b^+, c^+ and let Q be the graph of the paraboloid $z = x^2 + y^2$. First we show that the circumcircle of f_{abc} is the vertical projection of $P \cap Q$. Transform the space via

$$(x, y, z) \mapsto (x, y, z - x^2 - y^2).$$

This map sends Q to the xy-plane. The plane P becomes a concave down paraboloid through a, b, c, which intersects the xy-plane in the circumcircle of f_{abc}. The plane P partitions Q into the part below P, the intersection $P \cap Q$, and the part above P. Thus, the intersection projects down to the circumcircle of f_{abc} and the part below P projects down to B_{abc}.

It follows that d^+ is below the plane P if and only if $d \in B_{abc}$. □

Lemma 6 ([ES96])**.** *Let* $\mathrm{conv}(V^+)$ *be the convex hull of the lifted points of* V*. The vertical projection of the lower facets of* $\mathrm{conv}(V^+)$ *into the plane* $\{z = 0\}$ *gives the faces of the Delaunay triangulation.*

Proof. This follows from Lemmas 2 and 5. □

So we have three equivalent ways of looking at Delaunay triangulations: using circumcircles, using duals, and using convex hulls in $\mathbb{R}^3$.

2.1.3. Flip algorithm for finding the Delaunay triangulation in $\mathbb{R}^2$

Definition 9. Consider an interior edge e_{ab} adjacent to f_{abc} and f_{abd}. If the union of f_{abc} and f_{abd} is a convex quadrilateral, we can *flip* e_{ab} to e_{cd}. This means that we remove the triangles f_{abc} and f_{abd} as well as the edge e_{ab} and replace them with triangles f_{acd} and f_{bcd} and the edge e_{cd}.

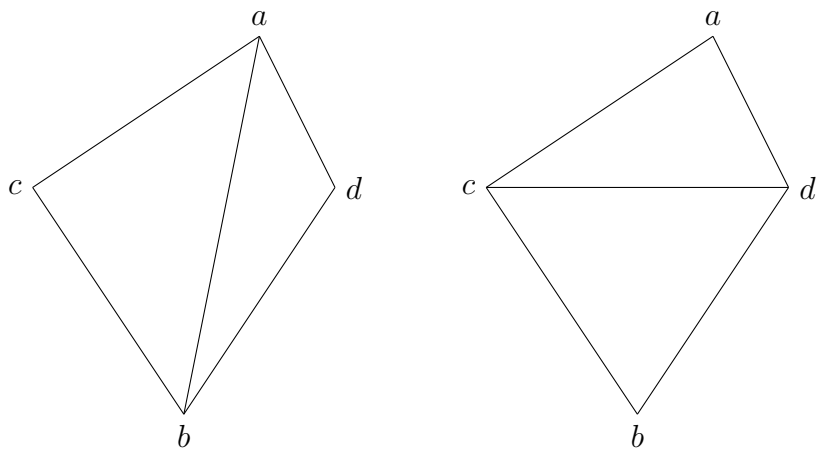

FIGURE 2.5. An example of a flip.

The following two lemmas are consequences of Lemma 1.

Lemma 7. *If an interior edge e_{ab} is not Delaunay, then flipping it results in an edge that is Delaunay.*

Lemma 8. *If an edge has a hinge which is not convex, then it is Delaunay.*

We can try to use edge flips to transform an arbitrary triangulation to the Delaunay triangulation. This was first proposed by Lawson [Law72]. Consider the following algorithm, which we call the flip algorithm:

1. Start with any triangulation $\mathcal{T}$ of (S, V).
2. If all interior edges of $\mathcal{T}$ are Delaunay, stop.
3. Flip an interior edge that is not Delaunay. Go to Step 2.

To prove that this algorithm works, we will use the notion of Dirichlet energy.

Definition 10. The *Dirichlet energy*, $E(f)$ of a piecewise C^1 function $f : U \to \mathbb{R}$ is defined as

$$E(f) = \frac{1}{2} \int_U \|\nabla f\|^2.$$

Let $g : V \to \mathbb{R}$ be a function and for each triangulation $\mathcal{T}$ of (S, V) let $g_{\mathcal{T}} : S \to \mathbb{R}$ be the piecewise-linear interpolation of g that is linear on each $f_{abc} \in \mathcal{T}_2$. The gradient $\nabla g_{\mathcal{T}}$ is constant on each triangle.

Theorem 9 ([Rip90])**.** *If triangulation $\mathcal{T}'$ is obtained from $\mathcal{T}$ by performing one flip that results in an edge that is Delaunay, then*

$$E(g_{\mathcal{T}}) \geq E(g_{\mathcal{T}'}), \tag{2.1}$$

where E denotes the Dirichlet energy. Equality holds only if $g_{\mathcal{T}} = g_{\mathcal{T}'}$, e.g., if the edge was Delaunay before the flip.

Remark 4. When V is not degenerate, there always exists a function that makes the inequality (2.1) strict.

We can now prove the following proposition.

Proposition 10. *The flip algorithm terminates after a finite number of steps, and the final triangulation is the Delaunay triangulation.*

Proof. We first need to know every edge that is not Delaunay can be flipped, which follows from Lemma 8.

Then we need to prove that the algorithm terminates after a finite number of steps. First note that there are only finitely many possible triangulations of (S, V), since simplices are uniquely determined by their vertices. Next, we use the Dirichlet energy. According to Theorem 9 [Rip90], the Dirichlet energy decreases each time we perform a flip. This guarantees that we do not encounter any infinite loops, and thus the algorithm terminates after a finite number of steps.

The final triangulation does not contain any edges which are not Delaunay. Thus, by Lemma 2, the final triangulation is Delaunay. □

2.2. Delaunay triangulations of piecewise-flat surfaces

In this section, we discuss Delaunay triangulations of piecewise-flat surfaces. We review the proof that a flip algorithm can be used to find the Delaunay triangulation starting from an arbitrary triangulation, and give a new proof that the algorithm terminates.

Consider a topological two-manifold S without boundary. Let $V \subset S$ be n distinct points. Let δ be a flat Riemannian metric on $S \setminus V$.

Definition 11. We say that (S, V, δ) is a *piecewise-flat surface* (PF surface) if the induced distance function induced by δ extends continuously to a distance function on all of S.

Remark 5. Let S be a sphere and V consist of two points. Then we can give $S \setminus V$ a flat Riemannian metric of a cone. However, the induced distance function cannot be extended to all of S.

Since there are infinitely many geodesic paths between two vertices on a PF surface, we cannot define an edge the way we did in the plane, just in terms of its endpoints. In addition, we need to use a more general object than a simplicial complex.

Definition 12. A *two-dimensional gluing* consists of a finite set of 2-simplices (triangles), a choice of pairs of 1-simplices (edges) such that each edge appears in exactly one of the pairs, and an identification map between the edges of each pair [Thu97].

We are interested in the quotient space of the union of the triangles by the equivalence relation generated by the identification maps, which is a Δ-complex [Hat02].

Let $\mathcal{T}$ be a Δ-complex which is the quotient space of a two-dimensional gluing. Let $\mathcal{T}_i$ denote the set of i-simplices of $\mathcal{T}$.

Definition 13. A *triangulation* of a piecewise-flat surface (S, V, δ) is a pair $(\mathcal{T}, \psi)$,

where $\mathcal{T}$ is a Δ-complex as above, $\psi : \mathcal{T} \to S$ is a homeomorphism, such that $\psi(\mathcal{T}_0) = V$, and given an edge $\epsilon \in \mathcal{T}_1$, $\psi(\epsilon)$ is a geodesic path between its endpoints.

Remark 6. Consider the pullback metric $\psi^*(\delta)$ on $\mathcal{T} \setminus \mathcal{T}_0$. This makes $(\mathcal{T}, \mathcal{T}_0, \psi^*(\delta))$ into a PF surface, and ψ into an isometry. Thus, we will so we will identify (S, V, δ) and $(\mathcal{T}, \mathcal{T}_0, \psi^*(\delta))$ when convenient.

We could have also defined a piecewise-flat metric on $\mathcal{T}$ by specifying edge lengths using a function $\ell : \mathcal{T}_1 \to \mathbb{R}_+ = (0, \infty)$, such that the edge lengths satisfy the triangle inequality for each triangle. In fact, ℓ and δ are interchangeable, and in practice, we define a PF surface by specifying $(\mathcal{T}, \ell)$.

Given a triangle $\tau \in \mathcal{T}_2$, let $\hat{\tau} \subset \mathbb{R}^2$ be a congruent triangle. We have a continuous map $\varphi_\tau : \hat{\tau} \to \mathcal{T}$, with $\varphi_\tau(\hat{\tau}) = \tau$, which is an isometry on $\hat{\tau}^\circ$, the interior of $\hat{\tau}$. Note that when $\mathcal{T}$ is a simplicial complex, then φ_τ is a bijection on all of $\hat{\tau}$. Let $\hat{B}_\tau \subset \mathbb{R}^2$ be the open disc whose boundary is the circumcircle of $\hat{\tau}$. The map φ can be extended to a local isometry on some larger open set $U \subset \mathbb{R}^2$. In addition, given two distinct triangles τ and σ which share an edge ϵ, it is clear that we can unfold the pair into a hinge in $\mathbb{R}^2$, i.e., we can extend φ_τ across $\varphi_\tau^{-1}(\epsilon)$ to a local isometry onto $(\tau \cup \sigma) \setminus \mathcal{T}_0$. Note that this since the triangles are distinct, this is actually an isometry on $\varphi_\tau^{-1}(\tau \cup \sigma)^\circ$, the interior of the hinge.

In the case when the edge ϵ belongs to only one triangle τ, we can attach another congruent copy $\hat{\tau}'$ of τ along an edge of $\hat{\tau}$ according to the identification map of the gluing. Then we can extend φ_τ to a continuous map on $\hat{\tau} \cup \hat{\tau}'$, which will be a local isometry on the interior of the hinge in $\mathbb{R}^2$, $(\hat{\tau} \cup \hat{\tau}')^\circ$.

Definition 14. The *hinge* of an edge $\epsilon \in \mathcal{T}_1$ is a hinge in $\mathbb{R}^2$ which consists of a pair of triangles $\hat{\tau}$ and $\hat{\sigma}$ which share an edge, together with a continuous map $\varphi_\epsilon : \hat{\tau} \cup \hat{\sigma} \to \mathcal{T}$ so that $\varphi_\epsilon(\hat{\tau}) = \tau \in \mathcal{T}_2$, $\varphi_\epsilon(\hat{\sigma}) = \sigma \in \mathcal{T}_2$, so that φ_ϵ is a local isometry on the interior of the hinge in $\mathbb{R}^2$.

If the hinge in $\mathbb{R}^2$ forms a cyclic quadrilateral, i.e., the hinge can be inscribed in a circle, we say that it is degenerate (not generic).

Definition 15. A triangulation $\mathcal{T}$ of (S, V, δ) is said to be *Delaunay* if for each triangle τ in $\mathcal{T}_2$, the map φ_τ can be extended to a local isometry on $\hat{B}_\tau$,

$$\varphi_\tau : \hat{B}_\tau \to \mathcal{T} \setminus \mathcal{T}_0.$$

In the generic case, this definition of a Delaunay triangulation of a PF surface coincides with the definition of a Delaunay tessellation in [BS07] and [Riv94].

Definition 16. An edge $\epsilon \in \mathcal{T}_1$ is *Delaunay* if its hinge in $\mathbb{R}^2$ is Delaunay in the $\mathbb{R}^2$ sense, as in Definition 2.

Proposition 11. *A triangulation $\mathcal{T}$ of a piecewise flat surface (S, V, δ) is the Delaunay triangulation if and only if all edges of $\mathcal{T}_1$ are Delaunay.*

Proof. We closely follow the argument given in [BS07].

It is clear that if the triangulation is Delaunay, then each edge is Delaunay.

Suppose each edge is Delaunay. Given a triangle τ, we will show that φ_τ can be extended to a local isometry $\varphi_\tau : \hat{B}_\tau \to \mathcal{T} \setminus \mathcal{T}_0$.

Starting with τ, let us develop a part of $\mathcal{T}$ in $\mathbb{R}^2$ by laying out triangles the same way as in the definition of a hinge. Begin with $\tau_0 = \hat{\tau} \subset \mathbb{R}^2$. Next, lay out triangles neighboring τ. The new vertices will be outside $\hat{B}_\tau$ because the edges of τ are Delaunay. Keep laying out neighboring triangles at free edges which intersect $\hat{B}_\tau$. We want to show that none of the new vertices introduced during this process intersect $\hat{B}_\tau$. Thus, when the layout process stops, the triangles cover $\hat{B}_\tau$, which implies that we have a local isometry $\varphi_\tau : \hat{B}_\tau \to \mathcal{T} \setminus \mathcal{T}_0$.

Consider some sequence $\tau_0 = \hat{\tau}, \tau_1, \ldots, \tau_k$ of triangles in the development such that τ_{i-1} and τ_i share an edge. Let g_i be the line containing the common edge, and let H_i be the half space bounded by g_i and containing τ_i.

Let x_{i+1} be vertex τ_{i+1} which is not a vertex of τ_i. By induction, assume that $x_i \notin \hat{B}_\tau$ for all $i < k$. Let us show that $x_k \notin \hat{B}_\tau$.

Since we only add triangles across edges which intersect $\hat{B}_\tau$, we observe that

$$(H_1 \cap \hat{B}_\tau) \supset (H_2 \cap \hat{B}_\tau) \supset \cdots \supset (H_k \cap \hat{B}_\tau).$$

Thus we only need to consider the case where $x_k \in H_i$ for all $i = 0, \ldots, k-1$, because otherwise $x_k \notin \hat{B}_\tau$.

Let π_i denote the power with respect to the circumcircle of τ_i. Then, since all the edges are Delaunay, and as long as none of the hinges are degenerate,

$$H_i = \{x \mid \pi_{i-1}(x) \geq \pi_i(x)\}.$$

Thus,

$$\pi_0(x_k) \geq \pi_1(x_k) \geq \cdots \geq \pi_k(x_k) = 0,$$

so $\pi_0(x_k) \geq 0$ and $x_k \notin \hat{B}_\tau$. □

Thus we can try to use edge flips to transform an arbitrary triangulation of the piecewise-flat surface to the Delaunay triangulation. Consider the following flip algorithm:

1. Start with any triangulation $\mathcal{T}$ of (S, V, δ).
2. If edges of $\mathcal{T}$ are Delaunay, stop.
3. Flip an edge that is not Delaunay. Go to Step 2.

Proposition 12. *There exists a finite sequence of flips which transforms an arbitrary triangulation into the Delaunay triangulation.*

Proof. We will show that we can use a flip algorithm to find the Delaunay triangulation. First we show that if an edge is not Delaunay, then it can be flipped.

For piecewise flat surfaces, there are two situations where an edge cannot be flipped. We may have a hinge that is not convex, and as in the case of $\mathbb{R}^2$, such a hinge is guaranteed to be Delaunay. In addition, we may have two sides of an isosceles triangle glued together at an edge of the triangulation. Such an edge cannot be flipped either. However, it is again guaranteed to be Delaunay, because the sum of the angles opposite of the glued edge is less than π. Thus the algorithm does not get stuck.

We also know that there are no infinite loops, since Theorem 9 also applies for piecewise flat surfaces [BS07].

However, since there are infinitely many triangulations of a piecewise flat surface, extra work is needed to prove the algorithm terminates in a finite number of steps. Given any M, there are finitely many triangulations such that all the edge lengths are at most M. We show that there are finitely many possible triangulations by showing that the maximum edge length is bounded.

In [BS07], they use a function of the triangulation called the harmonic index, which is related to the Dirichlet energy. Like the Dirichlet energy, the harmonic index decreases when an edge which is not Delaunay is flipped. When the harmonic index is bounded, so is the maximum edge length in the triangulation.

We instead apply Corrolary 14. Thus there is a finite sequence of flips which transforms an arbitrary initial triangulation of the piecewise-flat surface into the Delaunay triangulation. □

The following lemma is important because the argument can be generalized to the case of weighted Delaunay triangulations of piecewise flat surfaces.

Lemma 13. *If the longest edge is sufficiently long, then it is not Delaunay. More precisely, there exists M, which depends only on the piecewise flat manifold (S, V, δ), so that if the longest edge in a triangulation $\mathcal{T}$ of (S, V, δ) is longer than M, then it is not Delaunay.*

Proof. Let

$$M = \frac{4A}{m},$$

where A is the total area of (S, V, δ) and

$$m = \min_{\alpha,\beta \in V} \{d_\delta(\alpha, \beta)\},$$

with $d_\delta(\alpha, \beta)$ denoting the length of the shortest geodesic path between two vertices $\alpha, \beta \in V$, or, if $\alpha = \beta$, the length of the shortest nonconstant geodesic path from α to itself.

Let $e_{ab} \in \mathcal{T}_1$ be the longest edge, and denote its length by ℓ_{ab}. We show that if $\ell_{ab} > M$, then any triangle containing e_{ab} must be obtuse.

Let f_{abc} be a triangle containing the edge e_{ab}. Since e_{ab} is the longest edge in the triangulation, the angles at a and at b must be acute. Consider the altitude from c to e_{ab}. Denote its length by h and note that it splits the angle at c into two positive parts. We will show that each part of the angle at c greater than $\pi/4$.

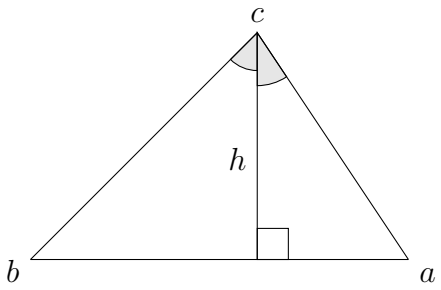

FIGURE 2.6. The triangle f_{abc} and the altitude from c.

The area of f_{abc} is

$$\frac{1}{2}\ell_{ab}h,$$

and since $\ell_{ab} > M$, we get

$$\begin{aligned} A &> \frac{1}{2}\ell_{ab}h \\ &> \frac{1}{2}Mh \\ &= \frac{1}{2}\left(\frac{4A}{m}\right)h \\ &= A\frac{2h}{m}, \end{aligned}$$

so $h < m/2$. Thus each part of the angle at c is greater than $\pi/3$ and so the angle at c is obtuse. (Note that we could have used $M = 2\sqrt{2}A/m$ instead to get a slightly better bound.)

So, if e_{ab} is the longest edge in the triangulation and $\ell_{ab} > M$, then any triangle containing e_{ab} have an obtuse angle opposite e_{ab}, and so the edge e_{ab} is not Delaunay.

□

Recall that when we flip an edge that is not Delaunay, the resulting edge is Delaunay. Thus we get the following corollary.

Corollary 14. *Let L be the length of the longest edge in a triangulation, let M be the bound from Lemma 13, and let $M' = \max\{L, M\}$. All the edges that occur during the flip algorithm will be at most M' long.*

This corollary guarantees that only finitely many triangulations can occur during the flip algorithm.

3. Weighted Delaunay triangulations

3.1. Weighted Delaunay triangulations in the plane

We start with a finite set of points V in $\mathbb{R}^2$ and let S be the convex hull of V. In addition, we assign a weight to each vertex via a function $w : V \to \mathbb{R}$ with $w(a) = w_a$. When the weight w_a is positive, we can think of it as representing a circle centered at a of radius $\sqrt{w_a}$. We generalize the notion of circle power from Section 2.1 to the case where w_a is allowed to be negative.

Definition 17. Given a weighted point p with weight w_p, we define $\pi_p : \mathbb{R}^2 \to \mathbb{R}$ so that $\pi_p(q) = \|q - p\|^2 - w_p$, where $\|q - p\|$ is the Euclidean distance between p and q. The value $\pi_p(q)$ is the *power* of the point q with respect to the weighted point p.

The set of points $x \in \mathbb{R}^2$ such that $\pi_p(x) = \pi_q(x)$ is a line.

Using the notion of power distance, we can define the weighted analog of the Voronoi diagram, also called the power cell diagram.

Definition 18. Given $V \subset \mathbb{R}^2$ a finite set of weighted points, define the *weighted Voronoi region* (or cell) $\mathcal{V}_a$ of $a \in V$ as the set of points $p \in \mathbb{R}^2$ whose power with respect to a is no larger than the power with respect to any other weighted point in V; that is

$$\mathcal{V}_a = \{p \in \mathbb{R}^2 \mid \pi_a(p) \leq \pi_b(p), \forall b \in V\}.$$

The weighted Voronoi regions together with their shared edges and vertices form the *weighted Voronoi diagram* of V. See Figure 3.3 for a partial illustration.

It is important to note that it is possible for a weighted Voronoi cell of a point to be empty. Recall that to get the Delaunay triangulation from the Voronoi diagram, two vertices were connected by an edge if their Voronoi cells shared an edge. Since

we would like the weighted case to be analogous to the one without weights, a point with an empty Voronoi cell would present a problem.

Definition 19. If the weighted Voronoi cell of a is empty, we say that a is a *redundant* point.

We will first describe what happens when there are no redundant vertices, since dealing with redundant vertices basically involves removing them from the triangulation.

We will define what it means for a triangulation to be weighted Delaunay using a weighted generalization of the circumcircle.

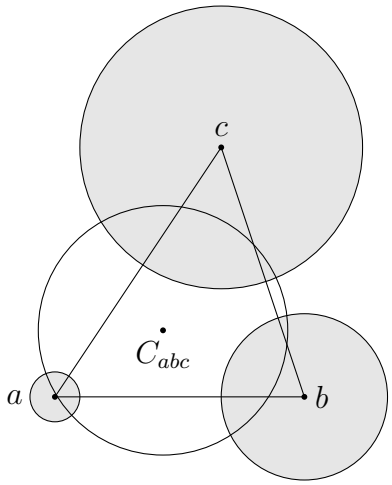

FIGURE 3.1. The weighted center C_{abc} of the triangle f_{abc}. The weights of the vertices are represented by circles.

Let f_{abc} be a triangle in a triangulation. If the weights w_a, w_b, w_c are all nonnegative, they represent circles centered at a, b, c with radii $\sqrt{w_a}, \sqrt{w_b}, \sqrt{w_c}$ respectively. There is a unique circle which intersects the three circles orthogonally; denote its center by C_{abc} and its radius by $\sqrt{w_{abc}}$. We can define the point C_{abc} without requiring the weights to be nonnegative as the unique point such that

$$\pi_a(C_{abc}) = \pi_b(C_{abc}) = \pi_c(C_{abc});$$

that is, C_{abc} is the common point of the weighted Voronoi regions of a, b, c if we restrict our attention to just those three vertices. Also, we let $w_{abc} = \pi_a(C_{abc}) = \pi_b(C_{abc}) = \pi_c(C_{abc})$. This allows us to define the center of f_{abc} and its weight even in the case when the weights are negative and thus cannot be interpreted as circles.

Note that when the weights of all the vertices are zero, the center C_{abc} is the circumcenter of f_{abc} and the weight w_{abc} is the circumradius squared.

The analog of a vertex d lying on the circumcircle of f_{abc} is the circle centered at d with radius $\sqrt{w_d}$ being orthogonal to the circle centered at C_{abc} with radius $\sqrt{w_{abc}}$. In terms of power, this corresponds to $\pi_{abc}(d)$ being equal to w_d, or, equivalently, $\pi_d(C_{abc})$ being equal to w_{abc}.

Definition 20. A triangulation $\mathcal{T}$ of (S, V, w) is said to be *weighted Delaunay* if for each triangle $f_{abc} \in \mathcal{T}_2$ and for each $d \in V$, we have $\pi_{abc}(d) \geq w_d$.

Note that we recover the definition of the Delaunay triangulation by setting the weights of the vertices equal to zero.

Definition 21. An edge $e_{ab} \in \mathcal{T}_1$ is said to be *weighted Delaunay* if the triangulation of the hinge at e_{ab} is weighted Delaunay.

Remark 7. Consider a triangle $f_{abc} \in \mathcal{T}_2$ and $d \notin \{a, b, c\}$. If $\pi_{abc}(c) = w_d$, we say that the set of weighted points (V, w) is degenerate, or not generic. If $f_{abd} \in \mathcal{T}_2$ is another triangle and $\pi_{abc}(c) = w_d$, then we say that the hinge at e_{ab} is degenerate.

The following is a weighted generalization of Lemma 3.

Lemma 15. *Suppose an interior edge e_{ab} is weighted Delaunay and adjacent to triangles f_{abc} and f_{abd}. If a ray starting at $p \in \mathbb{R}^2$ and crossing the interior of e_{ab} passes through f_{abc} first and through f_{abd} second, then*

$$\pi_{abd}(p) \geq \pi_{abc}(p),$$

with equality if and only if the hinge of e_{ab} is degenerate.

Lemma 16 ([ES96])**.** *Let $\mathcal{T}$ be a triangulation (S, V). If all the edges in $\mathcal{T}_1$ are weighted Delaunay, then the triangulation $\mathcal{T}$ is weighted Delaunay.*

Proof. This proof is almost identical to the proof of Lemma 2.

Suppose each edge in $\mathcal{T}_1$ is weighted Delaunay.

Consider a triangle $f_{abc} \in \mathcal{T}_2$ and a vertex $d \in V$. We will show that $\pi_{abc}(d) \geq w_d$. If d is a vertex of f_{abc}, $\pi_{abc}(d) = w_d$. Thus we only need to consider the case when d is distinct from a, b, c.

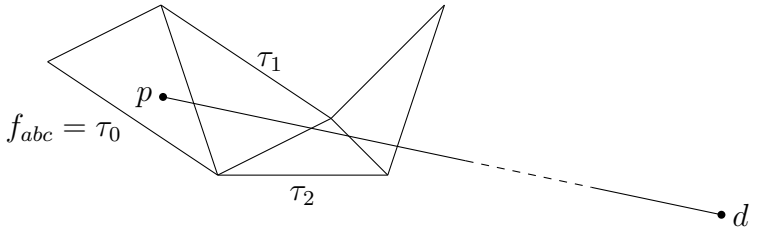

FIGURE 3.2. The segment pd and the sequence of triangles along it.

Choose a point $p \in f_{abc}$ so that the line segment from d to p contains no vertices of V in its interior. Let $\tau_0 = f_{abc}, \tau_1, \ldots, \tau_k$ be the sequence of triangles that intersect the segment pd. Let $\pi_i(d)$ be the power of d with respect to the weighted center of τ_i.

Since all the edges are weighted Delaunay, we have

$$\pi_0(d) \geq \pi_1(d) \geq \cdots \geq \pi_k(d).$$

Since d is one of the vertices of τ_k, $\pi_k(d) = w_d$ and so $\pi_0(d) = \pi_{abc}(d) \geq w_d$. □

3.1.1. Duals of triangulations with weights

Just as Delaunay triangulations are dual to Voronoi diagrams, weighted Delaunay triangulations are dual to weighted Voronoi diagrams. In addition, we want to define the dual of an arbitrary triangulation with weighted vertices, analogously to the way we defined the dual of an arbitrary triangulation in Section 2.1. In fact, we can define the dual of a triangulation even if some of the vertices are redundant. The definition

is similar to the one in Section 2.1, except instead of circumcenters and midpoints, we use weighted centers of the triangles and edges.

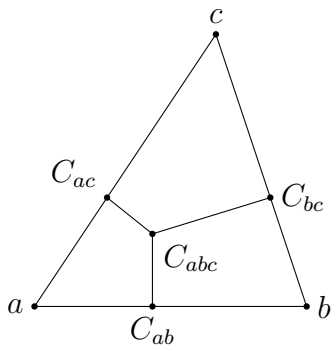

FIGURE 3.3. The weighted center C_{abc} of the triangle f_{abc} and the weighted centers of its edges.

Recall that C_{abc} is the unique point satisfying

$$\pi_a(C_{abc}) = \pi_b(C_{abc}) = \pi_c(C_{abc}),$$

and let C_{ab} denote the intersection of the line $\{p \in \mathbb{R}^2 \mid \pi_a(p) = \pi_b(p)\}$ and the line containing e_{ab}. Note that the intersection point does not have to be on e_{ab}. Let $h_{ab,c}$ denote the signed length from the C_{abc} to C_{ab}. It is equal in absolute value to the distance from C_{ab} to C_{abc} and is defined to be positive if C_{abc} is in the same half-plane determined by e_{ab} as c is, and negative if C_{abc} and c lie on opposite sides of e_{ab}.

Definition 22. Given an interior edge e_{ab} which belongs to two triangles f_{abc} and f_{abd}, the dual edge is the line segment connecting C_{abc} and C_{abd} which has a *signed length* ℓ^*_{ab} equal to $h_{ab,c} + h_{ab,d}$.

Proposition 17. *The signed length of the dual edge is nonnegative if and only if the edge e_{ab} is weighted Delaunay.*

Proof. Let e_{ab} be an interior edge which belongs to triangles f_{abc} and f_{abd}. If e_{ab} is weighted Delaunay, then the line segment from C_{abc} to C_{abd} is part of the weighted

Voronoi diagram of the hinge, and the signed length of the dual edge to e_{ab} is non-negative.

Suppose e_{ab} is not weighted Delaunay. Then at least one of $\pi_d(C_{abc}) < w_{abc}$ and $\pi_c(C_{abd}) < w_{abd}$ must hold. Without loss of generality, assume that $\pi_d(C_{abc}) < w_{abc}$ holds. We know that

$$\pi_d(C_{abc}) = \|C_{abc} - d\|^2 - w_d,$$

and

$$\begin{aligned} w_{abc} &= \pi_d(C_{abd}) \\ &= \|C_{abd} - d\|^2 - w_d. \end{aligned}$$

Thus

$$\|C_{abc} - d\| < \|C_{abd} - d\|,$$

i.e., C_{abc} is closer to d than C_{abd} is. This is only possible if the dual edge has negative signed length. □

3.1.2. Lifting triangulations with weights

Let us now generalize the definition of a lifted point to points with weights.

Definition 23. Given a point $p = (x, y) \in \mathbb{R}^2$ with weight w_p, define its *weighted lifted point* $p^+ = (x, y, z) \in \mathbb{R}^3$ where $z = x^2 + y^2 - w_p$.

Consider an interior edge e_{ab} adjacent to triangles f_{abc} and f_{abd}. Let f^+_{abc} be the triangle formed by a^+, b^+, c^+ and similarly for f^+_{abd}. Let H^+_{ab} be the union of f^+_{abc} and f^+_{abd}, forming the lifted hinge of e_{ab}.

Lemma 18. *Consider an interior edge e_{ab} and its lifted hinge H^+_{ab}. The edge e_{ab} is weighted Delaunay if and only if H^+_{ab} is concave up (or flat, in the degenerate case).*

Proof. Let f_{abc} and f_{abd} be adjacent to e_{ab}. If $C_{abc} = (\kappa, \nu) \in \mathbb{R}^2$ is the center of f_{abc} with weight w_{abc}, then

$$z = 2\kappa x + 2\nu y - \kappa^2 - \nu^2 + w_{abc}$$

is the plane which contains f^+_{abc}. We can check this algebraically. If $a^+ = (x_a, y_a, z_a)$, then,

$$\begin{aligned} z_a &= x_a^2 + y_a^2 - w_a \\ &= x_a^2 + y_a^2 - \pi_{abc}(a) \\ &= x_a^2 + y_a^2 - (\|a - C_{abc}\|^2 - w_{abc}) \\ &= x_a^2 + y_a^2 - (x_a - \kappa)^2 - (y_a - \nu)^2 + w_{abc} \\ &= 2\kappa x_a + 2\nu y_a - \kappa^2 - \nu^2 + w_{abc}. \end{aligned}$$

The edge e_{ab} is weighted Delaunay if and only if $\pi_{abc}(d) \geq w_d$. Thus, the lifted hinge is concave up if and only if

$$z_d \geq 2\kappa x_d + 2\nu y_d - \kappa^2 - \nu^2 + w_{abc},$$

which we can again verify similarly. □

Remark 8. Consider the case when the hinge is a convex quadrilateral. A useful consequence of this lemma is that flipping an edge that is not weighted Delaunay results in a new edge that is weighted Delaunay. This is true because flipping a hinge whose lift is concave down must result in a hinge which is concave up. A consequence of the proof is that the only way that a hinge and its flipped version can both be weighted Delaunay is when the lift of the hinge is flat; this holds if and only if the centers of the two triangles coincide.

Another consequence of Lemma 18 is the following lemma.

Lemma 19 ([ES96])**.** *Let* $\mathrm{conv}(V^+)$ *be the convex hull of the weighted lifted points of* V*. The vertical projection of the lower facets of* $\mathrm{conv}(V^+)$ *into the plane* $\{z = 0\}$ *gives the faces of the weighted Delaunay triangulation of* (S, V, w)*.*

Remark 9. The above lemma only holds when V contains no redundant points. If a is a redundant point, then a^+ will be vertically above a lower facet of $\mathrm{conv}(V^+)$, and so a will not be part of the projection. Because of that, any triangulation which includes a must have an edge whose lift is concave down. Thus, if any of the points in V are redundant, there does not exist a weighted Delaunay triangulation $\mathcal{T}$ of (V, S, w).

So, as in previous section, we have three equivalent ways of looking at weighted Delaunay triangulations: using power, using duals, and using lifts to $\mathbb{R}^3$.

3.1.3. Flip algorithm for finding the weighted Delaunay triangulation in $\mathbb{R}^2$

Whenever (S, V, w) contains no redundant vertices, practically the same flip algorithm as in Section 2.1 can be used to find a weighted Delaunay triangulation:

1. Start with any triangulation $\mathcal{T}$ of (S, V, w).
2. If all interior edges of $\mathcal{T}$ are weighted Delaunay, stop.
3. Flip an interior edge that is not weighted Delaunay. Go to Step 2.

Proposition 20. *If none of the vertices in* V *are redundant, then the flip algorithm terminates after a finite number of steps, and the final triangulation is the weighted Delaunay triangulation.*

Proof. We need to prove that the algorithm does not get stuck and terminates after a finite number of steps. In Lemma 21, we show that it is always possible to flip an edge that is not weighted Delaunay.

Since there are only finitely many triangulations of (S, V, w), to prove that the algorithm terminates after a finite number of steps, we use Glickenstein's generalization [Gli07] of Rippa's theorem [Rip90]. Glickenstein proved that the Dirichlet energy of a weighted triangulation decreases whenever an edge which is not weighted Delaunay is flipped. This implies that the algorithm does not result in any infinite loops, and so terminates after a finite number of steps with all edges weighted Delaunay. By Lemma 16, this means that the final triangulation is weighted Delaunay. □

Lemma 21. *Let e_{ab} be an interior edge adjacent to f_{abc} and f_{abd}. Suppose e_{ab} is not weighted Delaunay. If the angle at a in the hinge of e_{ab} is at least π, then a must be redundant.*

Proof. Suppose the angle in the hinge at a is at least π. Let f^+_{bcd} be the triangle in $\mathbb{R}^3$ with vertices b^+, c^+, d^+. Since e_{ab} is not weighted Delaunay, the lifted hinge H^+_{ab} is concave down. Thus, a^+ lies vertically above f^+_{bcd} and so is redundant. □

This implies that when V contains no redundant vertices, any edge that is not weighted Delaunay can be flipped.

Next, let us define the Dirichlet energy for the case of triangulations with weights [Gli07].

Definition 24. Given a triangulation $\mathcal{T}$ of (V, S, w) and a function $g : V \to \mathbb{R}$, the *Dirichlet energy*, $E_w(g, \mathcal{T})$, is defined as

$$E_w(g, \mathcal{T}) = \frac{1}{2} \sum_{e_{ab}} \frac{\ell^*_{ab}}{\ell_{ab}} (g_b - g_a)^2,$$

where ℓ_{ab} is the length of the edge e_{ab} and ℓ^*_{ab} is the signed length of the dual edge to e_{ab}.

Note that when the weights are zero, this Dirichlet energy is equal to the piecewise-linear interpolation Dirichlet energy in Rippa's theorem.

Theorem 22 ([Gli07])**.** *Consider a function $g : V \to \mathbb{R}$. If triangulation $\mathcal{T}'$ is obtained from $\mathcal{T}$ by performing one flip that results in an edge that is weighted Delaunay, then the following inequality holds*

$$E_w(g, \mathcal{T}) \geq E_w(g, \mathcal{T}').$$

Remark 10. If V contains no redundant vertices, the minimum Dirichlet energy among all possible triangulations is attained on the weighted Delaunay triangulations of (S, V, w), which can always be found using the flip algorithm.

3.1.4. Dealing with redundant vertices

Edelsbrunner and Shah [ES96] give a slightly different definition of the weighted Delaunay triangulation of (S, V, w), which agrees with our definition for the case when V does not contain redundant vertices. We will call this the relaxed weighted Delaunay triangulation. This definition only works when (S, V, w) is generic.

Definition 25. A triangle f_{abc} is said to be *weighted Delaunay* if $\pi_{abc}(d) > w_d$ for all $d \notin a, b, c$. This is equivalent to C_{abc} being one of the vertices of the weighted Voronoi diagram of (V, w). Then, the *relaxed weighted Delaunay triangulation* of (S, V, w) consists of all the weighted Delaunay triangles, along with their edges and vertices.

In fact, the relaxed weighted Delaunay triangulation is exactly the dual of the weighted Voronoi diagram as well as the vertical projection of the lower facets of the convex hull of V^+. The vertices of the relaxed weighted Delaunay triangulation are the nonredundant vertices of V. So, if $V' \subset V$ is the set of vertices that are not redundant, the relaxed weighted Delaunay triangulation of (S, V, w) is the weighted Delaunay triangulation of $(S, V', w|_{V'})$. Notice that the lower facets of the convex hull of V and the lower facets of the convex hull of V' are equal.

So, we somehow need to get from the vertex set V to the vertex set V'. If we knew which vertices are redundant ahead of time, we could just throw them all out before

creating the original triangulation. However, we usually do not discover that a vertex is redundant until we are in the middle of the flip algorithm. That is, we find an edge that is not weighted Delaunay, yet it cannot be flipped because its hinge is not convex (see Lemma 21). We modify the flip algorithm:

1. Start with any triangulation $\mathcal{T}$ of (S, V, w).
2. If all interior edges of $\mathcal{T}$ are weighted Delaunay, stop.
3. If an interior edge that is not weighted Delaunay has a convex hinge, flip it. Go to Step 2.
4. Get rid of the redundant vertex. Go to Step 2.

One approach to Step 4 is to simply throw the redundant vertex out of the triangulation and re-triangulate as necessary [She10]. We prefer to change the triangulation using flips.

Recall that in the lift a flip involved replacing the top facets of a tetrahedron with the lower ones or vice versa. However, given four points in $\mathbb{R}^3$, the tetrahedron equal to their convex hull may have three top facets and just one lower facet or vice versa. In fact, if a redundant vertex is in exactly three triangles, then in the lift the three faces become the top facets of a tetrahedron. So, we introduce a new type of of flip, called the "3 to 1" flip (as opposed to the usual "2 to 2" flip) which replaces the three triangles adjacent to a single vertex by just one triangle. The "3 to 1" flip and the "2 to 2" flip are both types of bistellar flips, and are illustrated in Figure 3.4.

The "3 to 1" flip gives us a way to get rid of a redundant vertex. Unfortunately, when we find an edge that is not weighted Delaunay with a hinge that is not convex, the redundant vertex may be in more than three faces, preventing us from performing the flip.

Edelsbrunner and Shah's algorithm solves this problem by finding the relaxed weighted Delaunay triangulation incrementally. They start with just an auxiliary

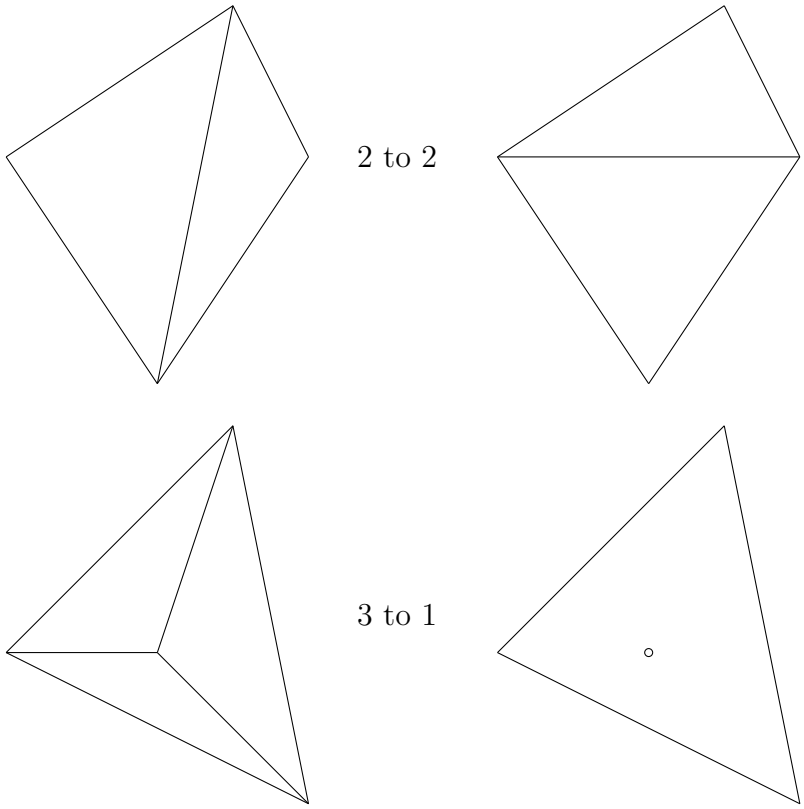

FIGURE 3.4. Illustration of "2 to 2" and "3 to 1" bistellar flips.

triangle, which has all of V in its interior, but they temporarily ignore V. Then they add back one vertex from V at a time, and find the relaxed weighted Delaunay triangulation before adding the next vertex. They prove that this guarantees that whenever we come across an edge that is not weighted Delaunay with a hinge that is not convex, the redundant vertex is exactly in three faces. In the end, they get rid of the auxiliary vertices and all the simplices incident on them. Note that their algorithm works for triangulations in any dimension, not just in $\mathbb{R}^2$.

Our solution for Step 4 of algorithm is as follows.

4. Get rid of the redundant vertex:

   (a) Use edge flips to make the redundant vertex be in exactly three faces.

   (b) Perform a "3 to 1" flip. Go to Step 2.

In Step 4a, we flip edges which contain the redundant vertex until there are exactly three of them.

Proposition 23. *Given an arbitrary triangulation $\mathcal{T}$ of (S, V, w), there exists a finite sequence of "2 to 2" and "3 to 1" flips which results in the relaxed weighted Delaunay triangulation of (S, V, w).*

Proof. Since there are only finitely many vertices in V, it is only possible to perform Step 4 finitely many times, thus it remains to prove that it is always possible to perform Step 4a. We do this in Lemma 26, which states that whenever a vertex is in more than three triangles, one of the edges it is in can be flipped, thus iteratively reducing the number of adjacent triangles.

So, we know that for each redundant vertex, Step 4 is performed through a finite sequence of "2 to 2" flips followed by a single "3 to 1" flip. Between the times when we encounter redundant vertices, there are no infinite loops via the same argument as in Proposition 20. Thus there is a finite sequence of flips which results in the relaxed weighted Delaunay triangulation of (S, V, w). □

Note that in the case when V contains no redundant vertices, this algorithm finds the weighted Delaunay triangulation of (S, V, w).

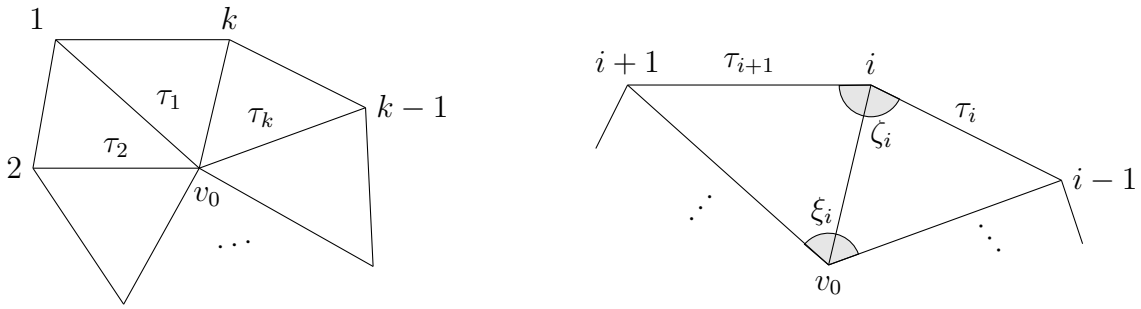

FIGURE 3.5. The triangles adjacent to vertex v_0, and the angles ξ_i and ζ_i.

Choose any vertex, and call it v_0. Let $\tau_1, \ldots, \tau_k$ be the triangles adjacent to v_0, so that τ_i and τ_{i+1} share the edge e_{0i}, and τ_1 and τ_k share the edge e_{0k}. Let θ_i be the

angle at v_0 in τ_i and let $\xi_i = \theta_i + \theta_{i+1}$ for $i < k$ and $\xi_k = \theta_1 + \theta_k$. So, the ξ_i is the angle at v_0 in the hinge of e_{0i}. Let ζ_i denote the angle opposite ξ_i in the hinge of e_{0i}.

Lemma 24. *If $k > 3$, then $\xi_i > \pi$ for at most two values of i. In fact, these two values must be consecutive.*

Proof. Without loss of generality, suppose $\xi_1 > \pi$. Then for $i = 3, \ldots, k-1$, $\xi_1 + \xi_i \leq 2\pi$ and so $\xi_i < \pi$. Thus, the only edges whose hinges we need to check, are the two adjacent to e_{01}, i.e., e_{02} and e_{0k}.

We also have that $\xi_k + \xi_2 \leq 2\pi$, so they cannot both be greater than π. □

Lemma 25. *There are at least three edges e_{0i} so that the angle $\zeta_i < \pi$.*

Proof. The star of v_0 is a polygon P. Consider the external angles of P. The sum of

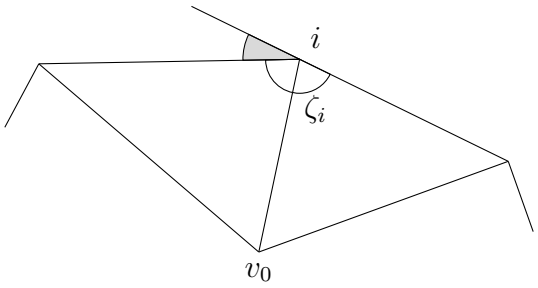

FIGURE 3.6. The external angle at the vertex i is shaded.

(signed) external angles of P is 2π. Since external angles are in $(-\pi, \pi)$, there must be at least 3 positive ones. If the external angle of P at i is positive, then $\zeta_i < \pi$, and so there are at least three edges e_{0i} for which the angle $\zeta_i < \pi$. □

Lemma 26. *If a vertex v_0 is in more than three triangles, then at least one of the edges e_{0i} can be flipped.*

Proof. The hinge at e_{0i} can be flipped if both $\xi_i < \pi$ and $\zeta_i < \pi$. We know that $\zeta_i < \pi$ for at least three values of i. Assuming no three points of V are collinear, we get that $\xi_i < \pi$ for all but at most two values of i. Thus at least one hinge has both $\zeta_i < \pi$ and $\xi_i < \pi$, and so at least one edge can be flipped. □

3.2. Weighted Delaunay triangulations of piecewise-flat surfaces

Consider a piecewise-flat surface (S, V, δ) and let $w : V \to \mathbb{R}$ be a function, as in Section 3.1, which assigns a weight to each vertex $w(v_i) = w_i$.

Definition 26. A piecewise-flat surface (S, V, δ) together with a function $w : V \to \mathbb{R}$ is a *piecewise-flat surface with weights*, (S, V, δ, w).

Definition 27. Let $e_{ab} \in \mathcal{T}_1$ be an edge of a triangulation $\mathcal{T}$ of (S, V, δ, w). The edge e_{ab} is said to be *weighted Delaunay* if the hinge of e_{ab} in $\mathbb{R}^2$ is weighted Delaunay with the same weights.

So we have the three equivalent ways of looking at weighted Delaunay edges from the plane case.

Now let's define power on a PF manifold with weighted vertices.

Definition 28. Given a weighted point p with weight w_p, we define $\pi_p : S \to \mathbb{R}$ as

$$\pi_p(q) = d_\delta(q, p)^2 - w_p,$$

where $d_\delta(q, p)$ is the length of the shortest geodesic path from q to p. The value $\pi_p(q)$ is the *power* of q with respect to p.

Once we defined power, the definition of the weighted Voronoi diagram and the definition of a redundant vertex are identical to the ones in the previous section.

We now turn our attention to the notion of curvature.

Definition 29. Consider a piecewise-flat surface (S, V, δ). Given a point $p \in S$, the *cone angle* $\theta(p)$ is defined as

$$\theta(p) = \frac{C_r}{r},$$

for small $r > 0$ and C_r the circumference of a small disk of radius r centered at p. The *curvature* $\kappa(p)$ is defined as

$$\kappa(p) = 2\pi - \theta(p).$$

Remark 11. If $p \in S \setminus V$, then $\theta(p) = 2\pi$ and $\kappa(p) = 0$. When the PF surface is triangulated, then for $p \in V$, we can define $\theta(p)$ instead as the sum of the angles at p of the triangles which contain p.

Consider the following tentative flip algorithm:

1. Start with any triangulation $\mathcal{T}$ of (S, V, δ, w).
2. If all edges are weighted Delaunay, stop.
3. Flip an edge that is not weighted Delaunay. Go to Step 2.

Recall that in the plane, a weighted Delaunay triangulation exists if and only if none of the vertices are redundant, i.e., all the vertices have nonempty weighted Voronoi cells. In the plane, Lemma 21 implies that when a weighted Delaunay triangulation exists, any hinge that cannot be flipped is already weighted Delaunay. We show that the situation is similar for surfaces of nonnegative curvature. In the case of a vertex of positive curvature, another reason why we may not be able to flip an edge is when it belongs to only one triangle (twice). In that situation, if the edge is not weighted Delaunay, then the vertex has an empty weighted Voronoi cell.

Lemma 27. *Suppose (S, V, δ, w) is a piecewise-flat surface with weighted vertices, and the curvature at every point is nonnegative. Let ϵ be an edge that is adjacent to two distinct faces. Let the hinge of ϵ in $\mathbb{R}^2$ consist of f_{abc} and f_{abd} together with a map $\varphi_\epsilon : f_{abc} \cup f_{abd} \to S$. Suppose e_{ab} is not weighted Delaunay and the angle at a in the hinge of e_{ab} is at least π, then the vertex $\varphi_\epsilon(a)$ has an empty weighted Voronoi cell.*

Proof. Using Lemma 21, we already know that the vertex a in the hinge of e_{ab} is redundant, i.e., has an empty weighted Voronoi cell. In other words, for every $x \in \mathbb{R}^2$, $\pi_a(x) > \pi_v(x)$ for at least one of $v = b, c, d$.

We need to show that a similar fact is true in the piecewise-flat surface. Let $\tilde{a}$, $\tilde{b}$, $\tilde{c}$, and $\tilde{d}$ denote the images of a, b, c, and d under φ_ϵ.

The images of some of the vertices may coincide and so may the images of some of the edges.

Let y be an arbitrary point in the PF surface.

If y belongs to the image of the hinge, then let x be a preimage of y under φ_ϵ. We know that $\pi_a(x) > \pi_v(x)$ for at least one of $v = b, c, d$. Also, $\pi_{\tilde{a}}(y) = \pi_a(x)$ and $\pi_{\tilde{v}}(y) \leq \pi_v(x)$, so $\pi_{\tilde{a}}(y) > \pi_{\tilde{v}}(y)$ for at least one of $\tilde{v} = \tilde{b}, \tilde{c}, \tilde{d}$.

Now suppose that y does not belong to the hinge. Consider a shortest geodesic segment γ between $\tilde{a}$ and y. Develop the triangulation into $\mathbb{R}^2$ along γ starting with the image of the hinge. As in the previous paragraph, we get $\pi_{\tilde{a}}(y) = \pi_a(x)$ and $\pi_{\tilde{v}}(y) \leq \pi_v(x)$, so $\pi_{\tilde{a}}(y) > \pi_{\tilde{v}}(y)$ for at least one of $\tilde{v} = \tilde{b}, \tilde{c}, \tilde{d}$.

Thus, the weighted Voronoi cell of $\tilde{a}$ is empty. □

Lemma 28. *Let ϵ be an edge that belongs to only one triangle. Let a be a vertex of ϵ which only belongs to ϵ. If it is not weighted Delaunay, then the power cell of a is empty.*

Proof. This is a consequence of Proposition 2.5 from [BI08]. □

In the case when vertices have negative curvature, things are different. It is possible have a hinge that is not convex and not weighted Delaunay, yet the vertex still has a nonempty weighted Voronoi cell, as in Figures 3.7 and 3.8.

Finally, in addition to showing that the algorithm does not get stuck, we need to show that it terminates after a finite number of steps. The proof is somewhat similar to that of Lemma 13, but significantly more involved.

Lemma 29. *If the longest edge is sufficiently long, then it is not weighted Delaunay. Also, it can be flipped. More precisely, there exists M, which depends only on the*

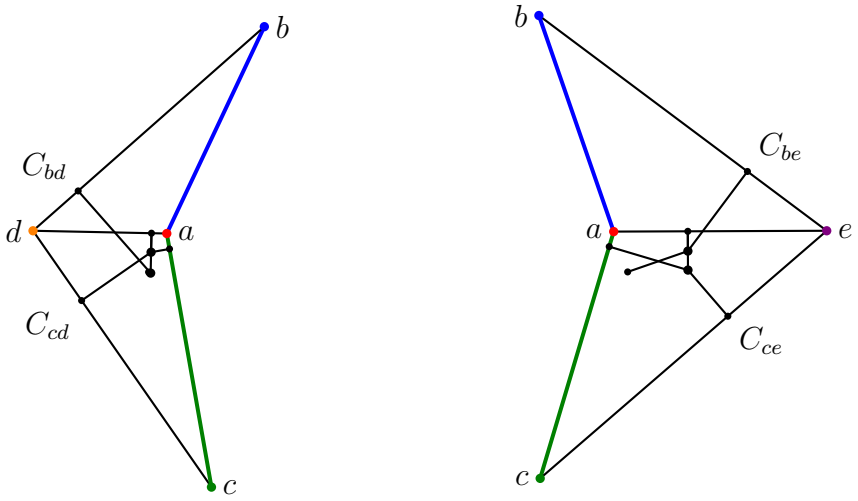

FIGURE 3.7. Glue the two hinges above along ab and ac. Dual edges are drawn.

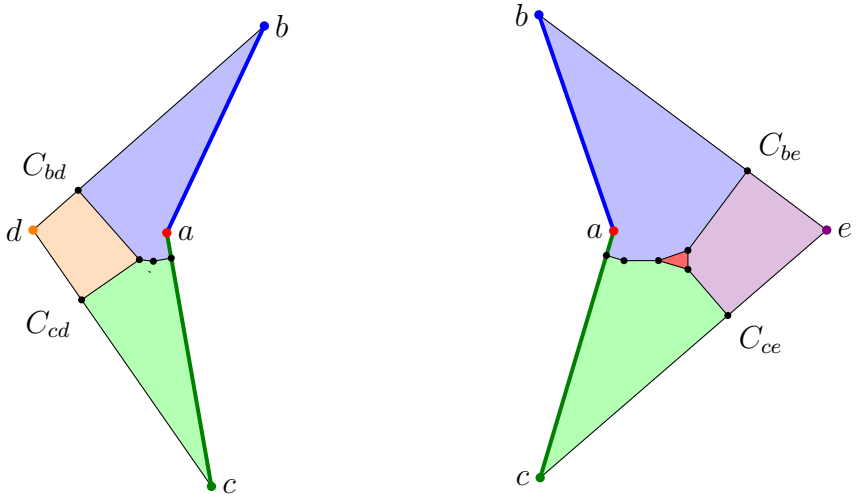

FIGURE 3.8. The weighted Voronoi diagram for the pair of hinges.

piecewise flat manifold with weights (S, V, δ, w), so that if the longest edge in a triangulation $\mathcal{T}$ of (S, V, δ, w) is longer than M, then it is not weighted Delaunay and can be flipped.

Proof. Let

$$M = \frac{4A + 2W}{m},$$

where

$$W = \max_{\alpha,\beta \in V} |w_\alpha - w_\beta|,$$

is the largest difference in weights, and A and m are as in the proof of Lemma 13, i.e., A is the total area of (S, V, δ) and

$$m = \min_{\alpha,\beta \in V} \{d_\delta(\alpha, \beta)\}.$$

Let $e_{ab} \in \mathcal{T}_1$ be the longest edge with length ℓ_{ab}. Suppose $\ell_{ab} > M$. Then the edge e_{ab} must be adjacent to two distinct triangles. Let c and d be the two other vertices in the hinge of e_{ab}. Since e_{ab} is the longest edge, the hinge must be convex. Thus the edge e_{ab} can be flipped.

Map the hinge of e_{ab} into $\mathbb{R}^2$ so that the midpoint of e_{ab} is mapped to the origin, a is on the positive x-axis, and c is in the upper half-plane. Denote the coordinates of a and c by $(x_a, 0)$ and (x_c, y_c), respectively. Then, in the lift, the z-coordinates

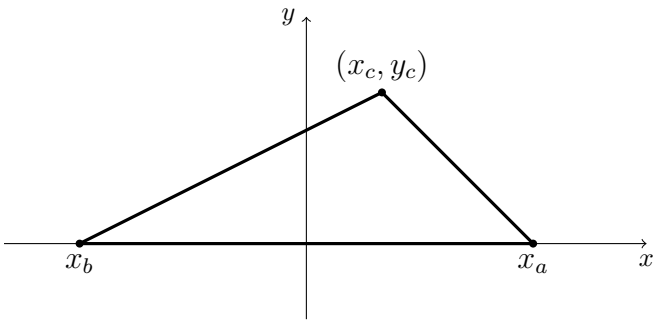

FIGURE 3.9. The triangle f_{abc} in $\mathbb{R}^2$.

of a^+ and c^+ are $z_a = x_a^2 - w_a$ and $z_c = x_c^2 + y_c^2 - w_c$, respectively. Without loss of generality, assume that $x_c \geq 0$.

The area of f_{abc} is

$$\frac{1}{2}\ell_{ab}h < A,$$

where $h = y_c$, so, as in the proof of Lemma 13 $y_c < m/2$. Thus we have

$$\begin{aligned}
(x_a - x_c)^2 + y_c^2 &\geq m^2, \\
(x_a - x_c)^2 &\geq m^2 - y_c^2, \\
x_a - x_c &\geq \sqrt{m^2 - y_c^2}, \\
x_a - \sqrt{m^2 - y_c^2} &\geq x_c, \\
x_a^2 - 2x_a\sqrt{m^2 - y_c^2} + m^2 - y_c^2 &\geq x_c^2, \\
x_a^2 - x_a m + m^2 &\geq x_c^2 + y_c^2.
\end{aligned}$$

So,

$$\begin{aligned}
z_a - z_c &= (x_a^2 - w_a) - (x_c^2 + y_c^2 - w_c) \\
&\geq x_a^2 - w_a - (x_a^2 - x_a m + m^2 - w_c) \\
&= x_a m - m^2 + w_c - w_a \\
&\geq \frac{2A + W}{m} m - m^2 - W \\
&> 2A - m^2 > 0.
\end{aligned}$$

Thus, z-coordinate a^+ is larger than the z-coordinate of c^+. Similarly, b^+ is above c^+, and both a^+ and b^+ are above d^+. Thus the lifted hinge is concave down, and so e_{ab} is not weighted Delaunay.

We have already shown it can be flipped. □

As in the previous chapter, we get a useful corollary.

Corollary 30. *Let L be the length of the longest edge in a triangulation, let M be the bound from Lemma 29, and let $M' = \max\{L, M\}$. All the edges that occur during the flip algorithm will be at most M' long.*

Thus we get the following theorem.

Theorem 31. *Let (S, V, δ, w) be a piecewise-flat surface of nonnegative curvature so that all weighted Voronoi cells are nonempty. Then the flip algorithm above finds the weighted Delaunay triangulation which is unique in the generic case.*

Proof. If all weighted Voronoi cells are nonempty, then it follows from Lemmas 27 and 28 that any hinge that is not weighted Delaunay can be flipped. Thus the algorithm does not get stuck. Also, we know from Lemma 29 and Corollary 30, together with Theorem 22 [Gli07] that the algorithm terminates. □

In the general case, we get the following weaker result.

Theorem 32. *The flip algorithm either terminates after a finite number of steps or gets stuck.*

3.2.1. Redundant vertices on a flat surface

In the plane, the concept of removing redundant vertices from a triangulation makes sense, so we can define the concept of a relaxed weighted Delaunay triangulation. In a piecewise-flat surface, we cannot indiscriminantly remove vertices because some of them are necessary in order to triangulate the surface using Euclidean triangles. Thus, it only makes sense to remove vertices with zero curvature. In this section, we discuss surfaces in which all vertices have zero curvature, i.e., the flat torus and the flat Klein bottle.

We need the surface equivalent of Proposition 23. The three lemmas that are used to prove this proposition hold true even on a flat surface, but in the case of a surface they are not sufficient. This is because in a surface, flipping an edge adjacent to a vertex does not have to result in the number of edges adjacent to the vertex being reduced. To guarantee that this number is decreased each time we flip an edge, we need to make sure that the redundant vertex is first not adjacent to itself.

Lemma 33. *Let a be a vertex in a flat PF surface. If there is a vertex $b \neq a$ in V, then there exists a sequence of edge flips so that there are no edges from a to itself.*

Proof. If all the edges containing a are closed loops, then there is no other vertex in the surface.

Thus, if there is more than one vertex in V, then we have at least one edge ϵ which connects a to another vertex. If there are no edges from a to itself, we are done.

When there are two or more distinct edges from a to itself, they all must have convex hinges, and so all but one of them can be eliminated by a sequence of flips as follows.

Develop the star of a into $\mathbb{R}^2$, so that the middle vertex is v_0, i.e., $\varphi(v_0) = a$. Choose the map φ so that $\varphi(v_1) = b \neq a$ and $\varphi(v_2) = a$. There are two cases. If $\varphi(v_3) \neq a$, then flipping the edge $\varphi(e_{02})$ will clearly reduce the number of edges from a to itself. If $\varphi(v_3) = a$, we have a triangle f_{023} whose image under φ has all three vertices equal to a. Since a has no curvature the three edges of $\varphi(f_{023})$ must be distinct. Thus if we flip the edge $\varphi(e_{02})$, we eliminate one of the edges from a to itself without introducing another such edge.

If there is exactly one edge from a to itself, we have to ensure that its hinge is convex. If its hinge is not convex, then the hinge of the longest edge containing a must be convex, so we can flip it. Repeat until we have that the hinge of the unique edge from a to itself is convex, and flip it.

Thus we can make a adjacent only to other vertices. □

Once a vertex is adjacent to only other vertices, the proof of Proposition 23 carries over directly, and we have the following proposition for the case of a flat surface.

Proposition 34. *Given an arbitrary triangulation $\mathcal{T}$ of (S, V, δ, w), there exists a finite sequence of "2 to 2" and "3 to 1" flips which results in the relaxed weighted Delaunay triangulation of (S, V, δ, w).*

3.2.2. Other consequences

A flip algorithm may fail for two reasons: we may not be able to flip any of the hinges which are not weighted Delaunay, or the algorithm may never terminate. Lemma 29 and Corollary 30 guarantee as long as the algorithm does not get stuck, it will always terminate. In Theorem 31, we used them to show that a flip algorithm can be used to find the weighted Delaunay triangulation of a piecewise-flat surface of non-negative curvature which has no vertices with empty weighted Voronoi cells. Another situation where the algorithm does not get stuck is when in the initial triangulation, the weighted center C_{ab} of each edge e_{ab} belongs to the edge e_{ab}. This condition is preserved by flips and guarantees that any hinge which cannot be flipped is already weighted Delaunay.

4. Future Work

We want to use the algorithms described in Chapter 3 to find weighted Delaunay triangulations. We have already proved that the algorithms terminate (see Propositions 20, 23, and Theorem 31). In the future, we would like to address the implementation and performance of these algorithms.

We know some of the necessary conditions and some of the sufficient conditions for the existence of the weighted Delaunay triangulation on a piecewise-flat surface. For example, we know that each vertex must have a nonempty weighted Voronoi cell, but we only show that this is sufficient when the curvatures are all non-negative (Lemma 27). In addition, we know that it is sufficient for all the weighted centers of the edges to belong to their respective edges. We would like to find the necessary and sufficient conditions for the existence of a weighted Delaunay triangulation on any piecewise-flat surface.

One of the advantages of the Delaunay triangulation and the weighted Delaunay triangulation is that they are intrinsic to the piecewise-flat surface (or piecewise-flat surface with weighted vertices). This allows us to define a Laplace operator that is both well-conditioned and intrinsic. In a situation where a weighted Delaunay triangulation does not exist, e.g. when we have a vertex with an empty weighted Voronoi cell, we would like to define some other canonical triangulation which also results in a well-conditioned Laplace operator.

In addition to studying piecewise-flat surfaces, we would like to study triangulations of higher dimensional piecewise-flat manifolds. Unfortunately, in three dimensions, weighted Delaunay triangulations do not necessarily result in well-conditioned Laplace operators. It would be useful to find another intrinsic triangulation for higher dimensional piecewise-flat manifots that gives rise to a Laplace operator with desirable properties.

References

[AK00] Franz Aurenhammer and Rolf Klein. Voronoi diagrams. In *Handbook of computational geometry*, pages 201–290. North-Holland, Amsterdam, 2000.

[Aur91] Franz Aurenhammer. Voronoi diagrams–a survey of a fundamental geometric data structure. *ACM Comput. Surv.*, 23:345–405, September 1991.

[BI08] Alexander I. Bobenko and Ivan Izmestiev. Alexandrov's theorem, weighted Delaunay triangulations, and mixed volumes. *Ann. Inst. Fourier (Grenoble)*, 58(2):447–505, 2008.

[BS07] Alexander I. Bobenko and Boris A. Springborn. A discrete Laplace-Beltrami operator for simplicial surfaces. *Discrete Comput. Geom.*, 38(4):740–756, 2007.

[CGY11] Daniel Champion, David Glickenstein, and Andrea Young. Regge's Einstein-Hilbert functional on the double tetrahedron. *Differential Geom. Appl.*, 29(1):108–124, 2011.

[Del34] Boris N. Delaunay. Sur la sphère vide. *Izv. Akad. Nauk SSSR, Otd. Mat. Estestv. Nauk*, (7):793–800, 1934.

[Ede90] H. Edelsbrunner. An acyclicity theorem for cell complexes in d dimension. *Combinatorica*, 10(3):251–260, 1990.

[Ede01] Herbert Edelsbrunner. *Geometry and topology for mesh generation*, volume 7 of *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press, Cambridge, 2001.

[EKS83] Herbert Edelsbrunner, David G. Kirkpatrick, and Raimund Seidel. On the shape of a set of points in the plane. *IEEE Trans. Inform. Theory*, 29(4):551–559, 1983.

[EM92] Herbert Edelsbrunner and Ernst P. Mücke. Three-dimensional alpha shapes. In *Proceedings of the 1992 workshop on Volume visualization*, VVS '92, pages 75–82, New York, NY, USA, 1992. ACM.

[ES96] H. Edelsbrunner and N. R. Shah. Incremental topological flipping works for regular triangulations. *Algorithmica*, 15(3):223–241, 1996.

[Gli07] David Glickenstein. A monotonicity property for weighted Delaunay triangulations. *Discrete Comput. Geom.*, 38(4):651–664, 2007.

[Gli11] David Glickenstein. Discrete conformal variations and scalar curvature on piecewise flat two- and three-dimensional manifolds. *J. Differential Geom.*, 87(2):201–237, 2011.

[Hat02] Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.

[KS96] Rob Kusner and John M. Sullivan. Comparing the Weaire-Phelan equal-volume foam to Kelvin's foam. *Forma*, 11(3):233–242, 1996.

[Law72] Charles L. Lawson. Transforming triangulations. *Discrete Math.*, 3:365–372, 1972.

[Pow92] P. L. Powar. Minimal roughness property of the Delaunay triangulation: a shorter approach. *Comput. Aided Geom. Design*, 9(6):491–494, 1992.

[Rip90] Samuel Rippa. Minimal roughness property of the Delaunay triangulation. *Comput. Aided Geom. Design*, 7(6):489–497, 1990.

[Riv94] Igor Rivin. Euclidean structures on simplicial surfaces and hyperbolic volume. *Ann. of Math. (2)*, 139(3):553–580, 1994.

[She10] Jonathan Richard Shewchuk. Personal communication, May 2010.

[Ste05] Kenneth Stephenson. *Introduction to circle packing*. Cambridge University Press, Cambridge, 2005. The theory of discrete analytic functions.

[Thu97] William P. Thurston. *Three-dimensional geometry and topology. Vol. 1*, volume 35 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 1997. Edited by Silvio Levy.

[Vor08] Georges Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. Deuxième mémoire. Recherches sur les parallélloèdres primitifs. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 1908(134):198–287, jan 1908.