A discrete Laplace-Beltrami operator for simplicial surfaces

Alexander I. Bobenko   Boris A. Springborn

February 24, 2006

Institut für Mathematik, Technische Universität Berlin,
Strasse des 17. Juni 136, 10623 Berlin, Germany

bobenko@math.tu-berlin.de
springb@math.tu-berlin.de

Abstract

We define a discrete Laplace-Beltrami operator for simplicial surfaces (Definition 16). It depends only on the intrinsic geometry of the surface and its edge weights are positive. Our Laplace operator is similar to the well known finite-elements Laplacian (the so called "cotan formula") except that it is based on the intrinsic Delaunay triangulation of the simplicial surface. This leads to new definitions of discrete harmonic functions, discrete mean curvature, and discrete minimal surfaces. The definition of the discrete Laplace-Beltrami operator depends on the existence and uniqueness of Delaunay tessellations in piecewise flat surfaces. While the existence is known, we prove the uniqueness. Using Rippa's Theorem we show that, as claimed, Musin's harmonic index provides an optimality criterion for Delaunay triangulations, and this can be used to prove that the edge flipping algorithm terminates also in the setting of piecewise flat surfaces.

Keywords: Laplace operator, Delaunay triangulation, Dirichlet energy, simplicial surfaces, discrete differential geometry

1 Dirichlet energy of piecewise linear functions

Let $\mathcal{S}$ be a *simplicial surface* in 3-dimensional Euclidean space, *i.e.* a geometric simplicial complex in $\mathbb{R}^3$ whose carrier S is a 2-dimensional submanifold,

Research for this article was supported by the DFG Research Unit 565 "Polyhedral Surfaces" and the DFG Research Center MATHEON "Mathematics for key technologies" in Berlin.

possibly with boundary. We assume $\mathcal{S}$ to be finite. Let $V = \{x_1, \ldots, x_{|V|}\}$, E, and F be the sets of vertices, edges and (triangular) faces of $\mathcal{S}$. Let $f : S \to \mathbb{R}$ be a piecewise linear (PL) function on S (linear on each simplex of $\mathcal{S}$). Then the gradient ∇f is constant on each triangle. The *Dirichlet energy* $E(f) = \frac{1}{2}\int_S \|\nabla f\|^2$ is

$$E(f) = \frac{1}{2} \sum_{(x_i,x_j)\in E} w_{ij}\,(f(x_i) - f(x_j))^2,$$

where the edge weights are

$$w_{ij} = \begin{cases} \frac{1}{2}(\cot\alpha_{ij} + \cot\alpha_{ji}) & \text{for interior edges} \\ \frac{1}{2}\cot\alpha_{ij} & \text{for boundary edges} \end{cases}$$

and α_{ij}, α_{ji} are the angle(s) opposite edge (x_i, x_j) in the adjacent triangle(s) (see Figure 1). This formula was, it seems, first derived by Duffin [7], who considers triangulated planar regions. It follows (by summation over the triangles) from the observation that the Dirichlet energy of a linear function on a triangle (x_1, x_2, x_3) with angles α_1, α_2, α_3 is $E(f_{|(x_1,x_2,x_3)}) = \frac{1}{4}\sum_{i\in\mathbb{Z}/3\mathbb{Z}} \cot\alpha_i\,(f(x_{i+1}) - f(x_{i+2}))^2$.

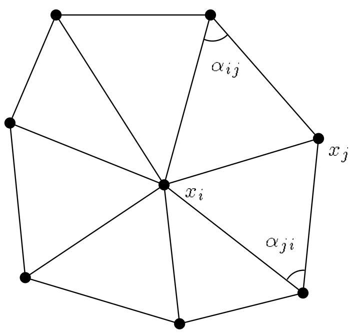

Figure 1: The α-angles of an internal edge.

In analogy to the smooth case, the Laplace operator is defined as the gradient of the Dirichlet energy. (We identify the vector space of PL functions $S \to \mathbb{R}$ with the the vector space $\mathbb{R}^V$ of functions on the vertices.) By differentiating $E(f)$ with respect to the value of f at a vertex $x_i \in V$ one obtains the "cotan formula" for the Laplace operator:

$$\Delta f(x_i) = \sum_{x_j\in V:(x_i,x_j)\in E} w_{ij}\,(f(x_i) - f(x_j)).$$

Dziuk was the first to treat a finite element approach for the Laplace operator on simplicial surfaces, but without stating the cotan formula explicitely [9]. It seems to have been rediscovered by Pinkall and Polthier in their investigation of discrete minimal surfaces [17], and turned out to be extremely important in geometry processing where it found numerous applications, *e.g.* [6], [4] to name but two. In particular, harmonic parameterizations $u : V \to \mathbb{R}^2$ are used in computer graphics for texture mapping. The cotan-formula also forms the basis for a theory of discrete holomorphic functions and discrete Riemann surfaces [8] [15].

Two important *disadvantages* of this definition of a discrete Laplace operator are:

1. The weights may be negative. The properties of the discrete Laplace operator with *positive weights* ($w_{ij} > 0$) are analogous to the properties of the classical Laplace-Beltrami operator on a surface with Riemannian metric. In particular the maximum principle holds. But some weights w_{ij} may be negative, and this leads to unpleasant phenomena: The maximum principle does not hold. As a consequence, a vertex of a discrete minimal surface (as defined by Pinkall & Polthier [17]) may not be contained in the convex hull of its neighbors. In texture mapping applications negative weights are undesirable because they lead to "flipped triangles". In practice various tricks are used to avoid negative weights.

2. The definition is not purely intrinsic. The classical Laplace-Beltrami operator is *intrinsic* to a Riemannian manifold: It depends only on the Riemannian metric. This is not the case with the discrete Laplace operator defined above. Two simplicial surfaces which are isometric but which are not triangulated in the same way give in general rise to different Laplace operators. As the simplest example, consider the two triangulations of a planar quadrilateral. They lead to different discrete Laplace operators. (Planarity is not what causes the problem since the quadrilateral may also be folded along either of its diagonals.)

The key idea of this paper is that one can avoid both the above shortcomings by using the *intrinsic* Delaunay triangulation of the surface S to define the discrete Laplace operator (Definition 16) instead of the triangulation that comes from the simplicial complex $\mathcal{S}$.

2 Delaunay triangulations of piecewise flat surfaces

This section provides the necessary background on Delaunay tessellations of piecewise flat surfaces. We decided to give a detailed exposition because not all necessary proofs can be found elsewhere.

The concept of a Delaunay triangulation in n-dimensional Euclidean space goes back to Delaunay [5]. Piecewise flat surfaces (Definition 1) were studied by (his student) Alexandrov [1] and more recently by Troyanov [20]. The idea of a Delaunay triangulation of a piecewise flat surface was apparently first considered by Rivin [19, Sec. 10]. The vertex set of the Delaunay triangulation is assumed to contain the set of cone points of the piecewise flat surface, so that the surface is flat away from the vertices. (This is very different from considering Delaunay triangulations in surfaces with Riemannian metric [13].) Rivin claims but does not prove an existence and uniqueness theorem for Delaunay triangulations in piecewise flat surfaces. His proof that the edge flipping algorithm terminates is flawed (see the discussion after Proposition 12 below). A correct proof was given by Indermitte *et al.* [11]. (They seem to miss a small detail, a topological obstruction to edge-flipability. See our proof of Proposition 11.) Furthermore, for our definition of the discrete Laplace-Beltrami operator we also need the *uniqueness* of the Delaunay tessellation, and this question has not been addressed properly. Rivin and Indermitte *et al.* define Delaunay triangulations by the local Delaunay criterion (see Definition 8), and infer existence via the edge flipping algorithm (see Proposition 12). To obtain uniqueness, we will define the Delaunay tessellation (whose faces are generically but not always triangular) via a global empty circle criterion (Definition 3). In a piecewise flat surface, the "empty circumcircles" are immersed empty disks (Definition 2) which may overlap themselves. Consequently, some work is required to show that this actually defines a (not necessarily regular) cell decomposition of the surface. Uniqueness, on the other hand, is immediate from this definition. We will also show that the local Delaunay criterion implies the global empty circumcircle condition. A Delaunay triangulation is obtained from *the* Delaunay tessellation by triangulating the non-triangular faces. It follows that a Delaunay triangulation, while in general not unique, differs from another Delaunay triangulation only by edges with vanishing cot-weights. This is important because it means that the discrete Laplace-Beltrami operator that will be defined in Section 3 depends only on the intrinsic geometry of the surface.

Definition 1. *A* piecewise flat surface (PF surface) (S, d) *is a 2-dimensional differential manifold S, possibly with boundary, equipped with a metric d which is flat except at isolated points, the* cone points*, where d has cone-like singularities.*

In other words, every interior point of a piecewise flat surface has a neighborhood which is isometric to either a neighborhood of the Euclidean plane or to a neighborhood of the apex of a Euclidean cone. The cone angle at the apex may be greater than 2π. (For a more detailed definition of closed PF surfaces see Troyanov [20].) In this paper, we consider only *compact surfaces* and we require the boundary (if there is a boundary) to be piecewise geodesic. (The interior angle at a corner of the boundary may be greater than 2π.)

A *tessellation* of a PF surface is a cell decomposition such that the faces are Euclidean polygons which are glued together along their edges. This implies that the cone points and the corners of the boundary are vertices of the cell decomposition. A *triangulation* is a tessellation where the faces are triangles. For the following it is essential that we do *not* require tessellations (and in particular triangulations) to be a *regular* cell complexes, *i.e.* a glueing homomorphism may identify points on the boundary of a cell. For example, it is allowed that two edges of a face may be glued to each other; and an edge may connect a vertex with itself. *A forteriori*, we do not require a tessellation to be *strongly regular*, *i.e.* the intersection of two closed cells may not be a single closed cell.

Remark. From the intrinsic point of view, the carrier S of a simplicial surface $\mathcal{S}$ with the metric induced by the ambient Euclidean space is a piecewise flat surface. The simplicial surface $\mathcal{S}$ also provides S with a triangulation whose vertex set includes the cone points and the corners of the boundary. However, this triangulation is not intrinsically distinguished from other triangulations with the same vertex set.

First, we will consider surfaces without boundary. To define the Delaunay tessellation of a PF surface in terms of empty disks, we must allow an empty disk to overlap with itself:

Definition 2 (immersed empty disk). *Let (S, d) be a compact PF surface* without boundary*, and let $V \subset S$ be a finite set of points which contains all cone points. An* immersed empty disk *is continuous map $\varphi : \bar{D} \to S$, where D is an open round disk in the Euclidean plane and $\bar{D}$ is its closure, such that the restriction $\varphi|_D$ is an isometric immersion (i.e. every $p \in D$ has a neighborhood which is mapped isometrically) and $\varphi(D) \cap V = \emptyset$.*

Hence any points in $\varphi^{-1}(V)$ are contained in the boundary of D: $\varphi^{-1}(V) \subset \partial D$.

Definition 3 (Delaunay tessellation, no boundary). *Let (S, d) be a compact PF surface* without boundary*, and let $V \subset S$ be a finite set of points which contains all cone points. The* Delaunay tessellation *of (S, d) on the vertex set V is the cell decomposition with the following cells: A subset $C \subset S$ is a closed cell of the Delaunay tessellation iff there exists an immersed empty disk $\varphi : \bar{D} \to S$ such that $\varphi^{-1}(V)$ is non-empty and C is the image of the convex hull of $\varphi^{-1}(V)$: $C = \varphi(\operatorname{conv} \varphi^{-1}(V))$. The cell-attaching map is $\varphi|_{\operatorname{conv} \varphi^{-1}(V)}$; and the cell is a 0-cell (vertex), 1-cell (edge), or 2-cell (face) if $\varphi^{-1}(V)$ contains one, two, or more points; respectively.*

Proposition 4. *The Delaunay tessellation as defined above is really a tessellation of (S, d).*

Proof. Let us first remark that the vertex set of the Delaunay tessellation is obviously V. An edge e is a geodesic segment such that there exists an immersed empty disk $\varphi : \bar{D} \to S$ with $\varphi^{-1}(V)$ containing exactly two points, $\varphi^{-1}(V) = \{p_1, p_2\}$ such that $e = \varphi([p_1, p_2])$, where $[p_1, p_2]$ is the line segment joining p_1 and p_2 in D. That the vertices and edges form a 1-dimensional cell complex then follows from Lemma 5 below.

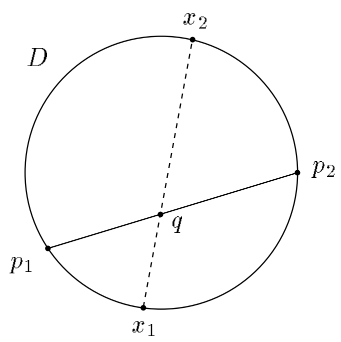

Figure 2: The Intersecting Chord Theorem says that $\|p_1 - q\| \cdot \|p_2 - q\| = \|x_1 - q\| \cdot \|x_2 - q\|$.

Then we have to show that the open faces are indeed homeomorphic to open disks. (A cell attaching map $\varphi : |_{\operatorname{conv} \varphi^{-1}(V)}$ is *a priori* only an immersion of the interior of the domain.) This follows from Lemma 6.

It is comparatively easy to see that the cell attaching homeomorphism φ maps the boundary $\partial \operatorname{conv}\{p_1, p_2, \ldots, p_n\}$ into the 1-skeleton; and that edges do not intersect open faces. We omit the details.

Finally, Lemma 7 asserts that every point in S is contained in a closed cell. □

Lemma 5. *The edges do not cross each other or themselves.*

Proof. Let $e = \varphi([p_1, p_2])$ and $\tilde{e} = \tilde{\varphi}([\tilde{p}_1, \tilde{p}_2])$ be edges contained in the empty immersed disks $\varphi : \bar{D} \to S$ and $\tilde{\varphi} : \bar{\tilde{D}} \to S$, respectively, such that p_1, p_2 are the only points in $\varphi^{-1}(V)$ and $\tilde{p}_1$, $\tilde{p}_2$ are the only points in $\tilde{\varphi}^{-1}(V)$. Suppose e and $\tilde{e}$ have an interior point in common: $\varphi(q) = \tilde{\varphi}(\tilde{q})$ with $q \in (p_1, p_2)$ and $\tilde{q} \in (\tilde{p}_1, \tilde{p}_2)$. We are going to show that

$$\varphi^{-1}(\tilde{e}) = [p_1, p_2]. \tag{1}$$

Since $\tilde{\varphi}(\tilde{p}_1), \tilde{\varphi}(\tilde{p}_2) \notin \varphi(D)$ the Intersecting Chord Theorem implies

$$\|p_1 - q\| \cdot \|p_2 - q\| \leq \|\tilde{p}_1 - \tilde{q}\| \cdot \|\tilde{p}_1 - \tilde{q}\|,$$

where $\|.\|$ denotes the Euclidean norm; see Figure 2. The opposite inequality follows equally. Hence $\varphi^{-1}(\tilde{\varphi}(\tilde{p}_1))$ and $\varphi^{-1}(\tilde{\varphi}(\tilde{p}_2))$ must be contained in ∂D. Since p_1 and p_2 are the only points in $\varphi^{-1}(V)$, this implies (1). □

Lemma 6. *Let $\varphi : \bar{D} \to S$ be an immersed empty disk and suppose $\varphi^{-1}(V) = \{p_1, p_2, \ldots, p_n\}$ with $n \geq 3$. Let $P = \operatorname{conv}\{p_1, p_2, \ldots, p_n\}$. Then the restriction $\varphi|_{\operatorname{int} P}$ of φ to the interior of P is injective (and hence a homeomorphism $\operatorname{int} P \to \varphi(\operatorname{int} P)$).*

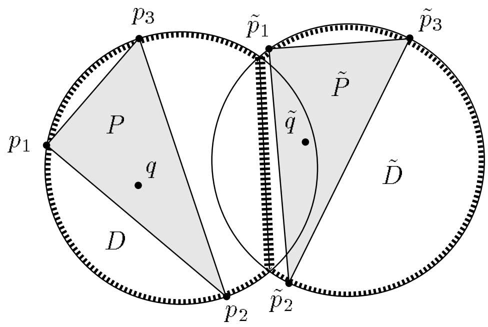

Figure 3: The dashed lines mark the boundaries of of the circular segments $\operatorname{conv}((\partial D) \setminus \tilde{D})$ and $\operatorname{conv}((\partial \tilde{D}) \setminus D)$.

Proof. Suppose $q \in \operatorname{int} P$ and $\tilde{q} \in D$ with $q \neq \tilde{q}$ but $\varphi(q) = \varphi(\tilde{q})$. We will show that $\tilde{q} \notin \operatorname{int} P$.

Because $\varphi|_D$ is an isometric immersion, there is a neighborhood $q \ni U \subset D$ and an isometry T of the Euclidean plane such that $T(q) = \tilde{q}$, $T(U) \subset D$, and $\varphi(T(x)) = \varphi(x)$ for all $x \in U$. Let $\tilde{D} = T(D)$ and $\tilde{\varphi} = \varphi \circ T^{-1} : \tilde{D} \to S$. Since φ and $\tilde{\varphi}$ agree on the intersection $D \cap \tilde{D}$, there is a continuous map $\hat{\varphi} : \overline{D \cup \tilde{D}} \to S$ such that $\hat{\varphi}|_{\bar{D}} = \varphi$ and $\hat{\varphi}|_{\bar{\tilde{D}}} = \tilde{\varphi}$.

Now let $\tilde{p}_i = T(p_i)$ $(i = 1, \ldots n)$ and $\tilde{P} = T(P) = \operatorname{conv}\{\tilde{p}_1, \tilde{p}_2, \ldots, \tilde{p}_n\}$. Then P and $\tilde{P}$ have no common interior points: $\operatorname{int} P \cap \operatorname{int} \tilde{P} = \emptyset$. Indeed, since D and $\tilde{D}$ are "empty", $\{p_1, p_2, \ldots, p_n\} \subset (\partial D) \setminus \tilde{D}$ and $\{\tilde{p}_1, \tilde{p}_2, \ldots, \tilde{p}_n\} \subset (\partial \tilde{D}) \setminus D$. Hence $P \subset \operatorname{conv}((\partial D) \setminus \tilde{D})$ and $\tilde{P} \subset \operatorname{conv}((\partial \tilde{D}) \setminus D)$. But the circular segments $\operatorname{conv}((\partial D) \setminus \tilde{D})$ and $\operatorname{conv}((\partial \tilde{D}) \setminus D)$ have no interior points in common. (See Figure 3.)

Now $\tilde{q} \in \operatorname{int} \tilde{P}$ implies $\tilde{q} \notin \operatorname{int} P$. □

Lemma 7. *Every point $x \in S$ is contained in a closed cell.*

Proof. Consider two immersed disks $\varphi : \bar{D} \to S$, $\tilde{\varphi} : \bar{\tilde{D}} \to S$ as equivalent if they differ only by a change of parameter, *i.e.* if $\tilde{\varphi} = \varphi \circ T$ for some isometry of the Euclidean plane. The manifold of equivalence classes is parameterized by the set

$$\mathcal{D} = \big\{(c, r) \in S \times \mathbb{R}_{>0} \,\big|\, d(c, V) \geq r\big\}$$

of center/radius pairs. If we adjoin degenerated immersed empty disks with radius 0, we obtain a compact manifold with boundary, parameterized by

$$\bar{\mathcal{D}} = \big\{(c, r) \in S \times \mathbb{R}_{\geq 0} \,\big|\, d(c, V) \geq r\big\}.$$

For a point $x \in S$ the *power function* with respect to x is the continuous function

$$\operatorname{pow}_x : \bar{\mathcal{D}} \to \mathbb{R}, \quad \operatorname{pow}_x(c, r) = \big(d(c, x)\big)^2 - r^2.$$

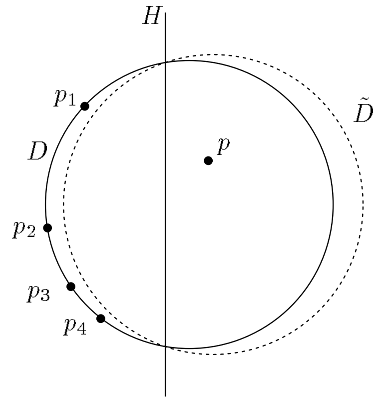

Figure 4: If p is not contained in the convex hull of $\varphi^{-1}(V) = \{p_1, \ldots, p_n\}$, then there is another immersed empty disk with smaller power. (See proof of Lemma 7.)

If $\varphi : \bar{D} \to S$ is an immersed empty disk with center c and radius r, then $\mathrm{pow}_x(c, r)$ is smaller than, equal to, or greater than zero depending on whether $x \in \varphi(D)$, $x \in \partial\varphi(D)$, or $x \in S \setminus \varphi(\bar{D})$.

Let $x \in S$. We have to show that x is contained in some closed cell. If $x \in V$ this is clear, because the points in V are 0-cells. So assume $x \in S \setminus V$. Since pow_x is continuous on the compact set $\bar{\mathcal{D}}$, there is a $(c_{\mathrm{min}}, r_{\mathrm{min}}) \in \bar{\mathcal{D}}$ where pow_x attains its minimum. Since $x \notin V$, there is an empty disk containing x and hence $\mathrm{pow}_x(c_{\mathrm{min}}, r_{\mathrm{min}}) < 0$. Let $\varphi : \bar{D} \to S$ be an immersed empty disk with center c_{min} and radius r_{min}; *i.e.* D is a disk in $\mathbb{R}^2$ with center $m \in \mathbb{R}^2$ and radius r_{min} and $\varphi(m) = c_{\mathrm{min}}$. There is a $p \in D$ with $\varphi(p) = x$ and $\|p - m\| = d(x, c_{\mathrm{min}})$. We show by contradiction that $p \in \mathrm{conv}\, \varphi^{-1}(V)$.

Suppose the opposite is true: $p \notin \mathrm{conv}\, \varphi^{-1}(V)$. Then there exists a closed half-space $H \subset \mathbb{R}^2$ with $\varphi^{-1}(V) \subset H$ but $p \notin H$. In that case there exists another immersed empty disk $\tilde{\varphi} : \tilde{D} \to M$ with $D \setminus \tilde{D} \subset \mathrm{int}\, H$, $\tilde{D} \setminus D \subset \mathbb{R}^2 \setminus H$, and $\varphi|_{D \cap \tilde{D}} = \tilde{\varphi}|_{D \cap \tilde{D}}$ (see Figure 4). Let $\tilde{m}$ and $\tilde{r}$ be center and radius of $\tilde{D}$. Then

$$\|p - \tilde{m}\|^2 - \tilde{r}^2 < \|p - m\|^2 - r^2.$$

(To see this, note that $q \mapsto (\|q - \tilde{m}\|^2 - \tilde{r}^2) - (\|q - m\|^2 - r^2)$ is an affine linear function $\mathbb{R}^2 \to \mathbb{R}$, which vanishes on ∂H and is positive on $\mathrm{int}\, H$, and negative on $\mathbb{R}^2 \setminus H$.) This implies $\mathrm{pow}_x(\tilde{c}, \tilde{r}) < \mathrm{pow}_x(c_{\mathrm{min}}, r_{\mathrm{min}})$, where $\tilde{c} = \tilde{\varphi}(\tilde{m})$. This contradicts the assumption that pow_x attains its minimum on $(c_{\mathrm{min}}, r_{\mathrm{min}})$. $\square$

Delaunay tessellations of PF surfaces with boundary. Now let (S, d) be a compact PF manifold with piecewise geodesic boundary. Let $V \subset S$

be a finite set of points which contains all cone points and all corners of the boundary. The boundary is then the union of geodesic segments connecting points in V but containing no points of V in their interior. Each connected component of the boundary is a closed geodesic polygon with vertices in V. To each such boundary polygon glue a PF surface obtained by cyclicly gluing together the appropriate number of isosceles triangles with appropriate base lengths and legs of length $R > 0$. Each of these caps contains a special point where the triangle apices are glued together. It is in general a cone point. The result of closing all wholes with such caps is a closed PF surface $(\hat{S}, d)$. Let $\hat{V}$ be the union of V and the set of special points of the caps. If R is chosen large enough, then the isosceles triangles in the caps will be faces of the Delaunay tessellation of $(\hat{S}, d)$ with respect to $\hat{V}$. (This is so because if R is large enough, the immersed circumdisks intersect S in lunes which are so small that they are empty.) The Delaunay tessellation of the bounded surface (S, d) with respect to V is defined to be the cell complex obtained by removing these triangles.

Delaunay triangulations. A *Delaunay triangulation* is a triangulation obtained from a Delaunay tessellation by triangulating all non-triangular faces. A Delaunay triangulation is characterized by the *empty circumcircle property:* A tessellation of (S, d) on the vertex set V is a Delaunay triangulation iff for each face f there exists an immersed empty disk $\varphi : \bar{D} \to S$ such that there are three points $p_1, p_2, p_3 \in \varphi^{-1}(V)$ with $f = \varphi(\operatorname{conv}\{p_1, p_2, p_3\})$. (But there may be more than three points in $\varphi^{-1}(V)$.)

In Proposition 10 we give a more local characterization of Delaunay triangulations.

Definition 8. *Let T be a geodesic triangulation of (S, d) with vertex set V, and let e be an interior edge of T. Since all faces of T are isometric to Euclidean triangles we can isometrically unfold to the plane the two triangles of T that are adjacent to e. We say that e is* locally Delaunay *if the vertices of the two unfolded triangles are not contained inside the circumcircles of these triangles.*

For our investigation of discrete Laplace operators we will need the following characterization of Delaunay edges.

Lemma 9. *An interior edge e of a triangulation T of a piecewise flat surface (S, d) is locally Delaunay if and only if the sum of the angles opposite e in the adjacent triangles does not exceed π.*

This follows immediately from the fact that opposite angles in a circular quadrilateral sum to π.

Clearly all interior edges of a Delaunay triangulation are locally Delaunay. This property actually characterizes Delaunay triangulations:

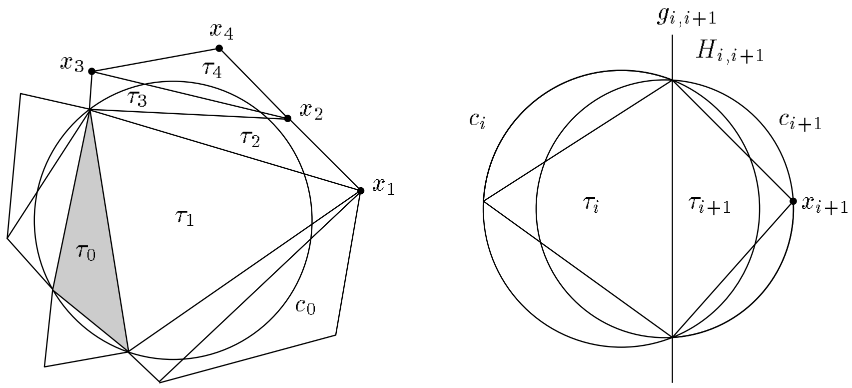

Figure 5: *Left:* The layed out triangles. *Right:* The power line of c_i and c_{i+1} is $g_{i,i+1}$. Since the edge $e_{i,i+1}$ is assumed to be Delaunay, $p_{c_i}(x) \geq p_{c_{i+1}}(x)$ in $H_{i,i+1}$.

Proposition 10. *A triangulation T of a piecewise flat surface (S, d) is a Delaunay triangulation if and only if all interior edges of T are locally Delaunay.*

The following proof is an adaptation of Delaunay's original argument [5] for Delaunay triangulations in $\mathbb{R}^n$. (See also Edelsbrunner [10, p. 8] for a more easily available modern exposition.) Alternatively, one could also adapt the argument of Aurenhammer and Klein [2] for Delaunay triangulations in the plane.

Proof. If S is a manifold with boundary, construct a closed PF surface by glueing piecewise flat disks to the boundary components in the manner described above. If R is chosen large enough, all edges will be locally Delaunay. It remains to prove the Proposition for closed PF surfaces.

Suppose that all interior edges of the triangulation T of the closed PF surface (S, d) are locally Delaunay. We want to show that T is a Delaunay triangulation. To this end we will show that the empty circumcircle property holds. Let t_0 be a triangle of T. Starting with t_0 we develop a part of T in the Euclidean plane (see Figure 5, left). Begin with a triangle τ_0 in the Euclidean plane that is congruent to t_0. Let c_0 be the circumcircle of τ_0. Next, lay out congruent copies of the triangles neighboring t_0. This introduces new vertices in the plane, which are on or outside c_0 because the edges of t_0 are Delaunay by assumption. Keep laying out neighboring triangles in the plane at free edges but only if the free edges intersect c_0. (Different layed out triangles may correspond to the same triangle in (S, d).) Each new triangle introduces a new vertex and we will show that they do not lie inside c_0. Hence, when the layout process stops (when all free edges do not intersect c_0), the triangles simply cover the inside of c_0. It follows that t_0 has an empty circumcircle in (S, d).

Let $\tau_0, \tau_1, \ldots, \tau_n$, be a sequence of layed out triangles such that τ_i and τ_{i+1} share an edge $e_{i,i+1}$. Let x_{i+1} be the vertex opposite $e_{i,i+1}$ in τ_{i+1}. Assuming that x_i is not inside c_0 for all $i < n$ we will show that the same holds for x_n. Let $g_{i,i+1}$ be the straight line containing $e_{i,i+1}$ and let $H_{i,i+1}$ be the half space bounded by $g_{i,i+1}$ and containing τ_{i+1}. Then

$$(H_{0,1} \cap D_0) \supset (H_{1,2} \cap D_0) \supset \ldots \supset (H_{n-1,n} \cap D_0),$$

where D_0 is the open disk bounded by c_0. Hence it remains to consider the case where $x_n \in H_{i,i+1}$ for all $i = 0, \ldots, n-1$, because otherwise $x_n \notin D_0$.

Now consider the power of a point $x \in \mathbb{R}^2$ with respect to a circle $c \subset \mathbb{R}^2$ with center x_c and radius r as a function of x:

$$p_c(x) = \|x - x_c\|^2 - r^2.$$

It is positive, zero, or negative if x lies outside, on, or inside c, respectively. The *power line* of two different circles c and c' is the locus of points x with $p_c(x) = p_{c'}(x)$. It is a straight line because $p_c(x) - p_{c'}(x)$ is linear in x. The power line of two intersecting circles is the line through the intersection points. Let c_i be the circumcircle of τ_i. Either $c_i = c_{i+1}$ or the power line of c_i and c_{i+1} is $g_{i,i+1}$ and

$$H_{i,i+1} = \{x : p_{c_i}(x) \geq p_{c_{i+1}}(x)\}. \tag{2}$$

(See Figure 5, right.) Indeed, $p_{c_{i+1}}(x_{i+1}) = 0$ and since the edge $e_{i,i+1}$ is locally Delaunay by assumption, $p_{c_i}(x_{i+1}) \geq 0$. Hence

$$p_{c_0}(x_n) \geq p_{c_1}(x_n) \geq \ldots \geq p_{c_n}(x_n) = 0,$$

and therefore x_n lies on or outside c_0. This concludes the proof. □

The *edge flipping algorithm* may be used to construct a Delaunay triangulation of a piecewise flat surface (S, d) with marked points $V \subset S$:

1. Start with any triangulation T of (S, d) with vertex set V.
2. If all interior edges of T are locally Delaunay, stop.
3. Otherwise there is an interior edge e of T which is not locally Delaunay. Perform an *intrinsic edge flip*: Replace e by the other diagonal of the quadrilateral formed by the two triangles adjacent to e. Go to Step 2.

The following two propositions show that this is indeed an algorithm.

Proposition 11. *If an edge is not locally Delaunay, then it can be flipped.*

Proof. In $\mathbb{R}^2$, an edge is flipable iff the two adjacent triangles form a convex quadrilateral. In a PF surface, there is an additional topological obstruction to flipability: Since the triangulation may not be regular, an edge may be adjacent to the same triangle on both sides. So suppose the edge e is not locally Delaunay, *i.e.* the sum of opposite angles exceeds π. Then there are two different triangles adjacent to e, because the sum of all angles in one triangle is π. These two triangles form a Euclidean quadrilateral (possibly with some of the boundary edges identified with each other), which is convex by the usual argument. Hence e can be flipped. $\square$

Proposition 12 (Indermitte et al. [11]). *The edge flipping algorithm terminates after a finite number of steps.*

Together with Proposition 10 this implies that the edge flipping algorithm produces a Delaunay triangulation (in the global empty-circumcircle sense). To prove Proposition 12, one has to show that the algorithm cannot loop infinitely. In the setting of *planar* Delaunay triangulations, it is enough to define a suitable real valued function on the set of triangulations on the given vertices which decreases (or increases) with each edge flip. Because this set of triangulations is finite, the algorithm has to terminate. As a further consequence such a function attains its minimal (or maximal) value on the Delaunay triangulations. Several such functions are known, see for example Lambert [12], Musin [16], Rivin [19, Sec. 10], and the survey article [2]. When we consider Delaunay triangulations of PF surfaces, however, the set of triangulations on the marked points may be infinite. (The fact that the number of combinatorial types of triangulations is finite [19] is not sufficient to make the argument.) For example, the surface of a cube has infinitely many geodesic triangulations on the eight vertices. To prove Proposition 12 by means of a function which decreases with every flip, one has to show that it has the following additional property.

Definition 13. *Let $\mathcal{T}$ be the set of triangulations of a PF surface on a given set of marked points and let $f : \mathcal{T} \to \mathbb{R}$. We say that f is* proper *if for any $M \in \mathbb{R}$ the number of triangulations $T \in \mathcal{T}$ with $f(T) \leq M$ is finite.*

In their proof of Proposition 12, Indermitte *et al.* [11] use the sum of squared circumcircle radii as proper function which decreases with every flip. (For a proof that it decreases with every flip, they refer to an unpublished PhD thesis. However, see Musin [16], his Theorem 3 and the following Lemma, for a hint on how to prove this.) Another possible choice is Musin's harmonic index. Below we show that it is proper and decreases with every flip. The latter fact we deduce from Rippa's Theorem [18], which is also of independent interest in connection with the Laplace-Beltrami operator. Rippa's proof holds without change also for piecewise flat surfaces.

Rippa's Theorem. *Let (S,d) be a piecewise flat surface and let $V \subset S$ be a set of marked points which contains the cone points and the corners of the boundary. Let $f: V \to \mathbb{R}$ be a function on the marked points. For each triangulation T of (S,d) with vertex set V let $f_T : S \to \mathbb{R}$ be the PL interpolation of f that is linear on the faces of the triangulation T.*

Suppose T_1 is a triangulation with an interior edge e and T_2 is obtained from T_1 by flipping e. If the edge e is a Delaunay edge after the flip, i.e. in T_2, then

$$E(f_{T_1}) \geq E(f_{T_2}),$$

where E denotes the Dirichlet energy as in Section 1. Equality holds only if $f_{T_1} = f_{T_2}$ or if e was also a Delaunay edge in T_1.

As a consequence, the minimum of the Dirichlet energy among all possible triangulations is attained on the Delaunay triangulations (S,d):

$$\min_T \int_S \mid \nabla f_T \mid^2 = \int_S \mid \nabla f_{T_D} \mid^2,$$

where T_D is any Delaunay triangulation.

Moreover, for generic $f: V \to \mathbb{R}$, this property of Delaunay triangulations is characteristic.

Rippa's proof is based on the following comparison formula for the Dirichlet energies of two possible triangulations of a quadrilateral

$$E(f_{T_1}) - E(f_{T_2}) = \frac{(f_1 - f_2)^2}{2\sin\theta} \frac{(r_1 + r_3)(r_2 + r_4)}{r_1 r_2 r_3 r_4}(r_1 r_3 - r_2 r_4).$$

Here T_1 and T_2 are the two triangulations of the convex quadrilateral $Q = (x_1, x_2, x_3, x_4)$ obtained by addition of the diagonals (x_1, x_3) and (x_2, x_4) respectively, f_1 and f_2 are the values of f_{T_1} and f_{T_2} at the intersection point x_0 of the diagonals, $r_1, \ldots, r_4$ are the distances from x_0 to the vertices $x_1, \ldots, x_4$ of the quadrilateral and θ is the intersection angle of the diagonals.

For a Euclidean triangle t with sides a, b, c and area A, Musin [16] defines the *harmonic index* as

$$hrm(t) = \frac{a^2 + b^2 + c^2}{A}.$$

The harmonic index of a triangulation T with face set F is the sum of the harmonic indices of all triangles:

$$hrm(T) = \sum_{t \in F} hrm(t).$$

Proposition 14. *The harmonic index is proper.*

Proof. In a PF surface, there may be infinitely many geodesic lines connecting two points, but for any $L \in \mathbb{R}$ only a finite number of them have length $\leq L$ [11]. Hence there are only a finite number of triangulations of a PF surface with marked points such that all edges are not longer than L. Now the Proposition follows from the inequality

$$hrm(T) \geq \frac{l_{\max}(T)}{A_{\mathrm{tot}}},$$

where $l_{\max}(T)$ is the largest length of an edge of T and A_{tot} is the total area of the PF surface. Indeed, if $hrm(T) \leq M$, then $l_{\max}(T) \leq MA_{\mathrm{tot}}$, and there are only finitely many triangulations satisfying this bound on edge lengths. $\square$

The following theorem was stated by Musin without proof [16].

Theorem 15. *With the notation and under the conditions of Rippa's Theorem*

$$hrm(T_1) \geq hrm(T_2)$$

and equality holds only if e is a Delaunay edge in T_1 as well. This implies that the harmonic index is minimal for a Delaunay triangulation T_D (and hence for all of them):

$$\min_T hrm(T) = hrm(T_D).$$

Proof. The harmonic index of a triangle is

$$hrm(t) = 4(\cot\alpha + \cot\beta + \cot\gamma),$$

where α, β, γ are the angles of the triangle. (Because the area is $A = \frac{1}{2}ah_a = \frac{1}{2}bh_b = \frac{1}{2}ch_c$, where h_a, h_b, h_c are the heights of the triangle; and $a/h_a = \cot\beta + \cot\gamma$, *etc.*)

For a triangulation T of (S, d) with vertex set V and $x \in V$ let $\delta_{x,T} : S \to \mathbb{R}$ be the function that is linear on the triangles of T, equal to 1 at x, and equal to 0 at all other marked points in V. Then

$$\sum_{x \in V} E(\delta_{x,T}) = \frac{1}{2} \sum_{\text{angles } \alpha \text{ in } T} \cot\alpha = \frac{1}{8}\, hrm(T).$$

Hence the theorem follows from Rippa's Theorem. $\square$

3 The discrete Laplace-Beltrami operator and discrete harmonic functions

We are now in a position to define the discrete Laplace operator on a simplicial surface in an intrinsic way.

Definition 16. *Let $\mathcal{S}$ be a simplicial surface with vertex set V and let S be its carrier, which is a piecewise flat surface. The* discrete Laplace-Beltrami operator Δ of a simplicial surface $\mathcal{S}$ *is defined as follows. For a function $f : V \to \mathbb{R}^n$ on the vertices, the value of $\Delta f : V \to \mathbb{R}^n$ at $x_i \in V$ is*

$$\Delta f(x_i) = \sum_{x_j \in V : (x_i, x_j) \in E_D} \nu(x_i, x_j)(f(x_i) - f(x_j)), \tag{3}$$

where E_D is the edge set of a Delaunay triangulation of S and the weights are given by

$$\nu(x_i, x_j) = \begin{cases} \frac{1}{2}(\cot \alpha_{ij} + \cot \alpha_{ji}) & \text{for interior edges} \\ \frac{1}{2} \cot \alpha_{ij} & \text{for boundary edges} \end{cases}. \tag{4}$$

Here α_{ij} (and α_{ji} for interior edges) are the angles opposite the edge (x_i, x_j) in the adjacent triangles of the Delaunay triangulation (see Figure 1).

The discrete Dirichlet energy of f *is*

$$\mathcal{E}_D = \frac{1}{2} \sum_{(x_i, x_j) \in E_D} \nu(x_i, x_j)(f(x_i) - f(x_j))^2.$$

Due to Lemma 9 and the formula $\cot \alpha + \cot \beta = \frac{\sin(\alpha+\beta)}{\sin \alpha \sin \beta}$ the discrete Laplace operator has non-negative weights. The edges with vanishing weights are diagonals of non-triangular cells of the Delaunay tessellation. Erasing such edges in (3) we obtain a discrete Laplace operator on the Delaunay tessellation of S with positive weights $\nu(x, x_i)$. Moreover, this property is characteristic for Delaunay triangulations: Consider a piecewise flat surface (S, d) with a triangulation T. Denote by Δ_T the Laplace operator of the triangulation T: it is given by the same formula (3) with the weights ν_T determined by the triangulation T by the same formulas (4) as for the Delaunay triangulation. The following observation is elementary.

Proposition 17. *The Laplace operator Δ_T of the triangulation T has non-negative weights ν_T if and only if the triangulation T is Delaunay.*

Laplace operators with positive weights on graphs possess properties analogous to the smooth theory.

Definition 18. *A discrete function $f : V \to \mathbb{R}^n$ on a simplicial surface is* harmonic *if $\Delta f(x) = 0$ for all interior vertices x.*

Discrete harmonic functions satisfy the maximum principle: A real valued harmonic function attains its maximum on the boundary. This implies:

Proposition 19. *For each interior vertex x, the value $f(x)$ of a harmonic function $f : V \to \mathbb{R}^n$ lies in the convex hull of the values $f(x_i)$ on its neighbors.*

We conclude this section with some standard facts regarding boundary value problems for the discrete Laplace operator. Let us fix a subset $V_\partial \subset V$ of vertices–which may but need not be the set of boundary vertices–and a function $g : V_\partial \to \mathbb{R}$. The problem of finding a function that satisfies $\Delta f(x) = 0$ for $x \in V \setminus V_\partial$ and $f(x) = g(x)$ for $x \in V_\partial$ is called a *Dirichlet boundary value problem*. The problem of finding a function that satisfies $\Delta f(x) = 0$ for $x \in V \setminus V_\partial$ and $\Delta f(x) = g(x)$ for $x \in V_\partial$ is a *Neumann boundary value problem*.

Theorem 20. *For arbitrary V_∂ and $g(x)$ the solutions of the Dirichlet and Neumann boundary value problems exist and are unique (up to an additive constant if $V_\partial = \emptyset$ and in the Neumann case). The solution minimizes the Dirichlet energy $\mathcal{E}_D$.*

Indeed, solutions of the Dirichlet and Neumann boundary value problems are critical points of the Dirichlet energy. Since the energy is a positive definite quadratic form, its only critical point is the global minimum.

More complicated boundary conditions, such as so called "natural boundary conditions" (see Desbrun et al. [6]), are also intensively used in geometry processing. The corresponding existence and uniqueness results are still missing.

4 Discrete mean curvature and minimal surfaces

In this section we adapt the definitions of the mean curvature vector for simplicial surfaces and minimal surfaces originally suggested by Pinkall and Polthier [17] to the discrete Laplace operator that we introduced in Section 3.

For a *smooth* immersed surface $f : \mathbb{R}^2 \supset U \to \mathbb{R}^3$ the mean curvature vector is given by the formula $H = \Delta f$, where Δ is the Laplace-Beltrami operator of the surface. For a simplicial surface we define the mean curvature vector by the same formula, following [17], but we use a different Laplace operator:

Definition 21. *Let $\mathcal{S}$ be a simplicial surface with carrier S. The* discrete mean curvature vector *at a vertex x is*

$$\mathcal{H}(x) = \Delta f(x),$$

where $f : S \to \mathbb{R}^3$ is the restriction of the identity map on $\mathbb{R}^3$ to S, and Δ is the discrete Laplace-Beltrami operator of Definition 16.

The discrete mean curvature vector $\mathcal{H}(x)$ of a simplicial surface corresponds to the integral of the mean curvature vector of a smooth surface over a neighborhood of the point x. (Note when we scale the simplicial surface, $\mathcal{H}$ varies as the integral of the mean curvature over some domain.) This suggests the following alternative definition:

Definition 22. *In the setup of Definition 21, let $C(x)$ be a Voronoi cell of the vertex x of a simplicial surface. The* discrete mean curvature vector density *at x is defined by*

$$H(x) = \frac{\Delta f(x)}{A(C(x))},$$

where Δ is the discrete Laplace-Beltrami operator of Definition 16 and $A(C(x))$ is the area of the Voronoi cell $C(x)$.

This definition is similar to the definition of mean curvature suggested by Meyer, Desbrun, Schröder and Barr [14]. Again, the difference is that we contend that one should use the discrete Laplace-Beltrami operator.

Definition 23 (Wide definition of simplicial minimal surfaces). *A simplicial surface is called* minimal *if its mean curvature vector vanishes identically.*

Since the embedding $f : S \to \mathbb{R}^3$ of a simplicial minimal surface is harmonic, Proposition 19 implies

Proposition 24. *Every interior vertex of a simplicial minimal surface lies in the convex hull of its neighbors.*

The following stricter definition is also natural.

Definition 25 (Narrow definition of simplicial minimal surfaces). *A simplicial surface $\mathcal{S}$ is called* minimal *if its mean curvature vector vanishes identically and the intrinsic Delaunay triangulation of the carrier of $\mathcal{S}$ coincides with the triangulation induced by the simplicial complex $\mathcal{S}$.*

Such a simplicial minimal surface is a critical point of the area functional under variations of the vertex positions [17]. We would like to note that there exists also a non-linear theory of discrete minimal surfaces based on the theory of circle patterns [3].

The mean curvature flow for simplicial surfaces is given by the equation

$$\frac{df}{dt}(x) = H(x).$$

This flow may change the Delaunay triangulation of the surface. At some moment two Delaunay circles coincide and the diagonal of the corresponding quadrilateral flips. However, since the weights of the diagonals vanish at the flip moment, the discrete Laplace-Beltrami operator and mean curvature vectors are continuous functions of time t.

For the numerical computation of discrete minimal surfaces one should use the algorithm of [17] with an extra step to adapt the Delaunay triangulation: Start with a simplicial surface $\mathcal{S}_0$ which respects the given boundary conditions. Calculate the Delaunay triangulation of the carrier S_0 and the weights ν. Find an $f : S_0 \to \mathbb{R}^3$ which respects the boundary conditions and minimizes the Dirichlet energy (see Definition 16). You may start with the embedding of S_0 as initial guess. You obtain a new simplicial surface $\mathcal{S}_1$ which is combinatorially equivalent to $\mathcal{S}_0$ but geometrically different. Calculate the Delaunay triangulation and weights ν of S_1 and find an $f : S_1 \to \mathbb{R}^3$ that minimizes the Dirichlet energy. Iterate.

Acknowledgments. We would like to thank Peter Schröder for useful discussions and for pointing out Rippa's paper to us. Günter M. Ziegler and Ivan Ismestiev have made valuable comments for which we are grateful. Last but not least we would like to thank the referees for their helpful suggestions, in particular for informing us of the reference [7].

References

[1] A. D. Alexandrov, *Convex Polyhedra*, Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2005. English translation with additional material of the 1950 Russian monograph.

[2] F. Aurenhammer and R. Klein, Voronoi diagrams, pp. 201–290 in J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, North-Holland, Amsterdam, 2000.

[3] A. I. Bobenko, T. Hoffmann and B. A. Springborn, Minimal surfaces from circle patterns: Geometry from combinatorics, Preprint arXiv:math.DG/0305184, 2003. To appear in *Ann. of Math.*

[4] M. Botsch and L. Kobbelt, An intuitive framework for real time freeform modeling, *ACM Transactions on Graphics* 23:3 (2004), 630–634.

[5] B. N. Delaunay, Sur la sphère vide, *Izvestia Akademii Nauk SSSR, Otdelenie Matematicheskikh i Estestvennykh Nauk (7)* 1934, 793–800.

[6] M. Desbrun, M. Meyer, P. Alliez, Intrinsic parameterizations of surface meshes, *Computer Graphics Forum* 21 (2002) (Proc. Eurographics Conf. 2002), 209–218.

[7] R. J. Duffin, Distributed and lumped networks, *Journal of Mathematics and Mechanics* 8:5 (1959), 793–826.

[8] R. J. Duffin, Potential theory on a rhombic lattice, *Journal of Combinatorial Theory* 5 (1968), 258–272.

[9] G. Dziuk, Finite elements for the Beltrami operator on arbitrary surfaces, pp. 142–155 in S. Hildebrandt and R. Leis, editors, *Partial Differential Equations and Calculus of Variations*, Lecture Notes in Mathematics 1357, Springer-Verlag, Berlin, 1988.

[10] H. Edelsbrunner, *Geometry and Topology for Mesh Generation*, Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press, Cambridge, 2001.

[11] C. Indermitte, Th. M. Liebling, M. Troyanov and H. Clemençon, Voronoi diagrams on piecewise flat surfaces and an application to biological growth, *Theor. Computer Science* 263 (2001), 263–274.

[12] T. Lambert, The Delaunay triangulation maximizes the mean inradius, *Proc. of the 6th Canadian Conf. on Computational Geometry (CCCG'94)*, 201–206, 1994.

[13] G. Leibon and D. Letscher, Delaunay triangulations and Voronoi diagrams for Riemannian manifolds, *Proceedings of the Sixteenth Annual Symposium on Computational Geometry (Hong Kong, 2000)*, 341–349, ACM, New York, 2000.

[14] M. Meyer, M. Desbrun, P. Schröder and A.H. Barr, Discrete differential-geometry operators for triangulated 2-manifolds, pp. 35–57 in H.-C. Hege and K. Polthier, editors, *Visualization and Mathematics III*, Springer-Verlag, Heidelberg, 2003.

[15] C. Mercat, Discrete Riemann surfaces and the Ising model, *Comm. Math. Phys.* 218 (2001), 177–216.

[16] O. Musin, Properties of the Delaunay triangulation, in *Proc. of the 13th Annual Symposium on Computational Geometry (SCG '97)*, 424–426, ACM, New York, 1997.

[17] U. Pinkall, K. Polthier, Computing discrete minimal surfaces and their conjugates, *Exper. Math.* 2:1 (1993), 15–36.

[18] S. Rippa, Minimal roughness property of the Delaunay triangulation, *Computer Aided Geometric Design* 7 (1990), 489–497.

[19] I. Rivin, Euclidean structures on simplicial surfaces and hyperbolic volume, *Ann. of Math.* 139 (1994), 553–580.

[20] M. Troyanov, Les surface Euclidiennes à singularités coniques, *Enseign. Math.* 32 (1986), 79–94.