On Voronoi-Delaunay Duality and Delaunay Meshes

Ramsay Dyer*
GrUVi Lab
School of Computing Science
Simon Fraser University, Canada

Hao Zhang†
GrUVi Lab
School of Computing Science
Simon Fraser University, Canada

Torsten Möller‡
GrUVi Lab
School of Computing Science
Simon Fraser University, Canada

Abstract

In this paper, we are concerned with Delaunay triangulations of the vertex set of a piecewise flat (pwf) surface. We first propose the notion of *well-formed* Voronoi diagrams and establish a precise dual relationship between them and proper Delaunay triangulations on pwf surfaces. Then we provide an algorithm which, given any input manifold triangle mesh, constructs a *Delaunay mesh*: a manifold triangle mesh whose edges form an intrinsic Delaunay triangulation of its vertex set. Rather than relying on a geodesic Delaunay triangulation on the input mesh, our algorithm swaps the physical mesh edges based on the locally Delaunay criterion. We prove that when a physical edge that is not locally Delaunay is swapped, the surface area of the mesh is reduced. In order to ensure a proper Delaunay triangulation, some new vertices may need to be introduced, leading to a refinement scheme, and we detail the cases involved.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representations

Keywords: Delaunay triangulation, Voronoi diagram, Delaunay mesh, edge swap, refinement, area minimization

1 Introduction

The Voronoi diagram of a point set P in $\mathbb{R}^2$ is a partition of the plane into cells, one for each point, often called a site, in P. The Voronoi cell of a site $p \in P$ is the collection of all points in the plane that are closer to p than any other site in P. The dual of the Voronoi diagram is obtained by connecting sites in P if and only if they lie in adjacent Voronoi cells. If the sites in P are in general position, i.e. no four of them are cocircular, then the resulting tessellation is a triangulation of the point set, called the Delaunay triangulation, and it is unique (see [de Berg et al. 1998]).

The concept of Delaunay triangulations can be extended to higher dimensions, e.g., in $\mathbb{R}^3$ we are concerned with a Delaunay tetrahedralization [Shewchuk 1997]. Under certain conditions it can also be extended to non-Euclidean geometries. In particular the intrinsic Delaunay triangulation of a sufficiently dense set of points on a Riemannian manifold is well defined in terms of geodesic curves [Leibon and Letscher 2000]. In this paper, we focus on the case presented in [Bobenko and Springborn 2005]: Delaunay triangulations of the vertex set of piecewise flat surfaces [Aleksandrov and Zalgaller 1967; Bobenko and Springborn 2005].

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

We show that the empty circumdisk property of Delaunay triangulations can be used to establish a Voronoi-Delaunay duality on pwf surfaces, which, to the best of our knowledge, has not been done before. Unlike the planar case, ensuring that the dual of the Voronoi diagram of a pwf surface forms a proper triangulation [Fisher et al. 2006] is non-trivial and is dictated by sampling. To this end, we propose the notion of *well-formed Voronoi diagrams* and prove its precise dual relationship to a proper Delaunay triangulation.

Despite the sampling issues and cone-like singularities, the mostly flat metric of pwf surfaces allows many results concerning Delaunay triangulations to be relayed from the planar case. For example, the well-known Delaunay edge swapping algorithm [Lawson 1977] is shown to work on a pwf surface as well [Bobenko and Springborn 2005], with essentially the same termination proof as in the planar case via the harmonic index [Musin 1997].

The second part of this paper deals with *Delaunay meshes*, which offers a further departure from its 2D counterpart. A Delaunay mesh is a manifold triangle mesh [Floriani et al. 2004] whose edges form an intrinsic, proper Delaunay triangulation of its vertex set and as such an edge swap can change the manifold itself. In this paper we will introduce an algorithm, based on edge swaps, that constructs Delaunay meshes. The main insight we offer and prove is that swapping mesh edges based on the locally Delaunay criterion [Bobenko and Springborn 2005] corresponds to *area minimization*. However, we still need to deal with topological constraints to ensure that the manifold property is maintained. We detail the potential problems involved and show that they can be manifestations of a lack of well-formedness, which we relate to poor sampling. Finally, we design appropriate refinement schemes to handle these difficulties and experimentally confirm fast termination and improvement of triangle qualities for the Delaunay meshes produced.

1.1 Motivation

Intrinsic Delaunay triangulations of surfaces have made recent appearances in geometry processing literature. Most notably, [Bobenko and Springborn 2005] have observed that the linear finite element discretization of the Laplace-Beltrami operator (the cot operator [Meyer et al. 2003]) has no negative edge weights on a Delaunay triangulation; this is desirable, e.g., as the discrete Harmonic maps computed

*e-mail: rhdyer@cs.sfu.ca
†e-mail: haoz@cs.sfu.ca
‡e-mail: torsten@cs.sfu.ca

would be guaranteed to be one-to-one [Gu and Yau 2002]. They advocate constructing an intrinsic Delaunay triangulation of the piecewise linear mesh surface and defining the cot operator in terms of that. Such a construction was implemented and described in a later paper [Fisher et al. 2006], and it was shown that the condition number of the operator was significantly improved in most cases. This improvement is significant for any application that involves the numerical evaluation of elliptic PDEs on triangle mesh surfaces. Examples include parameterization [Desbrun et al. 2002] and reaction diffusion textures [Turk 1991].

It was primarily these results that have sparked our interest in the notion of Delaunay meshes. As explained in [Fisher et al. 2006], constructing an intrinsic Delaunay triangulation of a triangle mesh requires maintaining a data structure to record the connectivity describing the intrinsic triangulation in addition to the data structure describing the triangle mesh itself. This implementational burden would be unnecessary if the triangle mesh was itself an intrinsic Delaunay triangulation of its vertices: a Delaunay mesh.

In order to transform a given mesh into a Delaunay mesh, we usually need to either distort its geometry or add vertices. However, for most modelling and graphic applications, the triangle mesh is considered an approximation to a smooth or piecewise smooth surface. So there is no need to be strictly faithful to the geometry of the input mesh; small distortions may be tolerated. Further, many modern surface reconstruction and remeshing algorithms produce meshes that are quite close to being Delaunay, thus only small adjustments need be made. On the other hand, such distortions change the geometry and therefore the underlying Voronoi diagram. The impact of these changes on the Voronoi diagram (and its resulting Delaunay mesh) as well as resulting topological changes need to be considered in order to create a robust algorithm.

1.2 Contributions

The contributions of our paper can be summarized as

- established the duality relationship between the Voronoi diagram and the Delaunay tessellation of a piecewise flat (pwf) surface (Section 3);
- using this duality to characterize, in terms of the Voronoi diagram, meshes that admit a proper intrinsic Delaunay triangulation
- developed an algorithm (based on edge swaps and refinement) that produces a Delaunay mesh efficiently. (Section 4.2);
- a proof that the edge swaps of this algorithm lead to mesh surface area minimization (Section 4.2.1);
- a characterization of unswappable edges, their relation to inadequate sampling, and an associated refinement scheme (Section 4.2.4).

2 Related Work

In addition to the elegant work of [Bobenko and Springborn 2005] on intrinsic Delaunay triangulation of pwf surfaces, [Leibon and Letscher 2000] had earlier studied the same problem on general Riemannian manifolds. Both works relied on the empty circumdisk property to define Delaunay triangulations whose edges are given by appropriate geodesic arcs. In contrast to the planar case, sampling density of the point set becomes relevant to ensure that the Delaunay triangulations are well-defined over a manifold. To this end, [Leibon and Letscher 2000] resorted to the notion of strong convexity radius as a means of constraining sampling density. Weak bounds on the number of samples required to obtain a well defined Delaunay triangulation are presented in [Onishi and Itoh 2003], but the utility of this work in practical applications seems limited.

To the best of our knowledge, there was no previous work that produced Delaunay meshes as defined here. Although, in light of our new insight, the early surface area minimization algorithm of [O'Rourke 1981] turns out to be performing Delaunay edge swaps, it is not guaranteed to produce a Delaunay mesh due to possible topological constraints, as we note in Section 4.2.4. Appropriate refinement is necessary, which we develop in this work. Note that many Delaunay refinement techniques have been developed for planar triangulations [Ruppert 1995; Rivara and Inostroza 1997] and 3D tetrahedralizations [Shewchuk 1997].

Many works exist that employ Delaunay concepts and describe algorithms which produce meshes that are close to being Delaunay. [Chew 1993] adapted his Delaunay refinement technique to curved surfaces. This algorithm produced an approximate geodesic triangulation of the surface with a guaranteed angle bound of $[30^\circ, 120^\circ]$; intersections of the surface with a sphere are used in place of geodesic disks. Later, [Chen and Bishop 1997] used Delaunay refinement in a planar parametric space from which a triangulation is mapped back in 3D. Here the geodesic discs on the surface were approximated by ellipses in the parameter space.

More recently remeshing algorithms have appeared that are based on geodesic distances on a surface. [Peyré and Cohen 2003] describe a farthest point sampling method based on geodesic distances computed by the fast marching method of [Kimmel and Sethian 1998]. By taking the dual of the Voronoi diagram of the sampled points, a triangulation is produced that is a good approximation to the intrinsic Delaunay triangulation of the sample points. Another more complicated technique was based on triangulating strips between equidistant curves [Sifri et al. 2003]. This algorithm also uses the fast marching method to compute geodesic distances and finishes by taking the dual of the Voronoi diagram of the samples as the triangulation defining the mesh.

There is another class of surface reconstruction and remeshing algorithms that are based on the notion of a restricted Delaunay triangulation. There are different definitions of the restricted Delaunay triangulation of a surface [Edelsbrunner and Shah 1994; Amenta and Bern 1998; Amenta et al. 2000] and they are not all equivalent [Dey et al. 2005]. Nonetheless, the restricted Delaunay triangulation does not define a Delaunay mesh in general. Perhaps the simplest definition of a restricted Delaunay triangulation is the dual of the restricted Voronoi diagram [Amenta et al. 2000]. The latter is the restriction of the 3D Voronoi diagram to the surface. This definition corresponds to the "good triangles" defined by [Amenta and Bern 1998]. In Section 4.1, we discuss how the good triangles may not yield a Delaunay mesh. But these meshes are again generally close to being Delaunay and can be good candidates for post processing via edge swapping.

Edge swapping for geometry processing is not new. Early work on producing Delaunay triangulations in the plane traces back to [Lawson 1977]. Following the assumption that a smoother mesh is better, edge flipping algorithms designed to minimize curvature measures have been proposed [van Damme and Alboul 1995; Dyn et al. 2001]. These algorithms

produce visually pleasing meshes, but in general they are far from being Delaunay meshes. A recent paper focuses more on triangle element quality by edge swapping so as to minimize the discrete Willmore energy [Alboul et al. 2006]. All of these mesh edge swapping algorithms tend to converge to local rather than global minima. This includes the surface area minimization algorithm of O'Rourke [O'Rourke 1981] as well as the locally Delaunay edge swapping scheme we describe in this paper.

3 Delaunay Triangulations and Voronoi Diagrams on pwf Surfaces

In this section we review, in Section 3.1, the Delaunay triangulation on a piecewise flat (pwf) surface as defined by [Bobenko and Springborn 2005] and then show in Section 3.2 that the duality relationship with the Voronoi diagram can be extended to this setting. This allows us to characterize in terms of Voronoi diagrams the meshes that will admit a proper Delaunay triangulation, in Section 3.3.

3.1 Delaunay Triangulations on pwf Surfaces

Defining a Delaunay triangulation of a discrete set P of points, called *samples*, on a Riemannian surface S requires much more care than is needed in the planar setting. Difficulties arise as there may not be a unique shortest geodesic between two points, or since there may not be a unique geodesic disk that has three given points on its boundary.

One approach is to put constraints on the density of samples on S. The idea is that in a sufficiently small neighbourhood a manifold is well approximated by a plane. Thus if the samples are sufficiently close to each other, the obstacles to defining a Delaunay triangulation will be avoided. This is the approach developed by [Leibon and Letscher 2000].

Another approach is to constrain the types of surfaces and samples under consideration. This is the approach that was taken by [Bobenko and Springborn 2005] and is the one that we will follow. This approach requires no explicit constraints on the sampling density, but uses a weaker definition of a triangulation than is traditional in differential geometry.

Definition 2. *A* tessellation *of a compact Riemannian surface S with respect to a finite discrete point set P is as follows.*

Let E be a collection of curves on S, which form a connected graph $\mathcal{G}$ whose vertex set is P, such that $S - \mathcal{G}$ is a disjoint union of open subsets f_i, each homeomorphic to a disk. The elements of E are called edges *of the tessellation.*

The f_i's are called faces, and for each face there exists a continuous map $\varphi_i : \gamma_i \to \bar{f}_i$, where γ_i is a closed planar polygon and $\bar{f}_i$ is the closure of f_i. The map φ_i is a homeomorphism on the interior of γ_i, and is continuous on the boundary and such that vertices of γ_i get mapped to elements of P that lie on the boundary of f_i. If γ_i is an n-gon, we call f_i an n-gon face, and in particular, if γ_i is a triangle, then we also call f_i a triangle face.

A triangulation *is a tessellation in which all the faces are triangle faces. A* geodesic tessellation *is a tessellation in which all the edges are geodesics.*

Note that edges cannot cross in a tessellation. The mappings φ_i are not required to be injective on the boundary of γ. In particular, two edges of γ_i may be mapped onto a single edge in E. Likewise, the restriction of φ_i to the vertices of γ_i is not required to be injective.

We will confine our attention to compact pwf surfaces without boundary (see Definition 1). On a pwf surface every point has a neighbourhood that is either isometric to a neighbourhood in $\mathbb{R}^2$ or to a neighbourhood of the apex of a single cone. A manifold triangle mesh can be considered a pwf surface that is isometrically immersed in $\mathbb{R}^3$. However pwf surfaces are a more general class of objects, and they do not necessarily admit an isometric immersion in $\mathbb{R}^3$; the flat torus is a well known counter example [do Carmo 1976].

We require that the finite set of sample points P includes all of the cone points of M. In the pwf setting we often refer to the elements of P as *vertices*, emphasizing that the model pwf surface we have in mind is a mesh. The Delaunay tessellation is defined in terms of empty disks.

An *immersed empty disk* is a continuous map $\phi : \bar{D} \to M$, where D is an open round disk in the Euclidean plane and $\bar{D}$ is its closure, such that the restriction $\phi|_D$ is an isometric immersion (i.e. every $p \in D$ has a neighbourhood which is mapped isometrically) and $\phi(D) \cap P = \emptyset$.

We can think of ϕ as wrapping D on M, but it may wrap around onto itself: ϕ is not injective in general. It should be emphasized that ϕ is defined on the closure of D and that only the image of D itself is required to be empty. Most of the time we are working with empty disks that have elements of P on their boundary, so that $\phi^{-1}(P)$ is non-empty.

Immersed empty disks are more convenient to work with than geodesic disks because they allow us to work with ordinary disks in the plane, with the caveat that the mapping ϕ is not injective in general. Since M is flat in a neighbourhood not containing cone points, we can always find an isometric immersion ϕ for an empty disk and if two immersions ϕ and ϕ' have the same geodesic disk as their image, then there will be a planar isomorphism $T : \mathbb{R}^2 \to \mathbb{R}^2$ such that $\phi = \phi' \circ T$. So working with immersed empty disks is really equivalent to working with geodesic disks.

Thus we can place D wherever is convenient on the plane. In particular, we have the following useful lemma, whose proof is indicated in [Bobenko and Springborn 2005, Lemma 6].

Lemma 1. *Suppose that $\phi : \bar{D} \to M$ and $\phi' : \bar{D}' \to M$ are two immersed empty disks with $\phi(D) \cap \phi'(D') \neq \emptyset$. Then there exists a disk $\tilde{D}$, $\tilde{D} \cap D \neq \emptyset$ an isometry $T : \mathbb{R}^2 \to \mathbb{R}^2$, with $T(\tilde{D}) = D'$ and an isometric immersion $\hat{\phi} : \overline{D \cup \tilde{D}} \to M$ such that $\hat{\phi}|_{\bar{D}} = \phi$ and $\hat{\phi}|_{\bar{\tilde{D}}} = \phi' \circ T$.*

The *Delaunay tessellation* of (M, d_M) on the vertex set P is defined by the immersed empty disks $\phi : \bar{D} \to M$ such that $\phi^{-1}(P)$ is non-empty. If $\phi^{-1}(P)$ contains three or more points, then the convex hull $\operatorname{conv} \phi^{-1}(P)$ is a polygon and $\phi|_{\operatorname{conv} \phi^{-1}(P)}$ defines a face of the tessellation. If $\phi^{-1}(P)$ contains exactly two points, then the image of $\phi|_{\operatorname{conv} \phi^{-1}(P)}$ is an edge. It was established [Bobenko and Springborn 2005] that this does indeed describe a tessellation, something that is not obvious a priori.

Note that if a face contains more than three vertices, the diagonals of the face are not included in the tessellation. To obtain a *Delaunay triangulation*, we triangulate all non-triangular faces. A face of the Delaunay triangulation is still contained in an immersed empty disk, but there may be more than three vertices on the disk's boundary.

We say that the vertices are in *general position* if there exists no empty disk with more than three vertices on its boundary. In this case the Delaunay tessellation is the unique Delaunay triangulation of the vertices. The Delaunay triangulation of the surface is often referred to as the *intrinsic* Delaunay triangulation to emphasize that it is based

on geodesic distances and not distances in the ambient Euclidean space.

Now, consider an arbitrary geodesic triangulation $\mathcal{T}$ of the vertices of M. Since the triangles are empty of cone points, they are intrinsically planar. Given an edge e of $\mathcal{T}$, we can map the two triangular faces adjacent to e isometrically onto the plane forming a quadrilateral with e as its diagonal. We say that e is *locally Delaunay* if it is contained in a disk that does not have the other two vertices of the quadrilateral in its interior. This is different from the immersed empty disk criteria in that we are only considering two additional vertices of M.

A convenient characterization of a locally Delaunay edge is presented in [Bobenko and Springborn 2005]:

Lemma 2. *An edge is locally Delaunay if and only if the sum of the two angles it subtends does not exceed π.*

This follows from the fact that in a quadrilateral whose vertices lie on a circle, opposite angles sum to π. As in the planar case, the intrinsic Delaunay triangulation can be obtained by systematically swapping the geodesic edges that are not locally Delaunay. An edge e that is not locally Delaunay is replaced by the edge e' that is the other diagonal of the quadrilateral defined by the triangles adjacent to e. This algorithm runs in $\mathcal{O}(n^2)$ time, n being the number of vertices in the mesh. The proof described in Shewchuck's thesis [Shewchuk 1997] holds without modification to the case of a fixed piecewise flat surface.

For our purposes we are primarily concerned with Delaunay triangulations on the vertex set of meshes. A triangle mesh comes with an inherent triangulation defined by its faces and edges. We refer to this as the *physical triangulation* of the mesh, and in particular, the edges of the mesh are *physical edges*. This is to distinguish it from the Delaunay triangulation of its vertices, which consists of geodesic edges that do not correspond to the physical edges in general. We define a *Delaunay mesh* as a triangle mesh whose physical triangulation coincides with the Delaunay triangulation of its vertices.

3.2 Voronoi Diagrams on pwf Surfaces

In this section we examine the Voronoi diagram of a pwf surface and its relationship with the Delaunay tessellation. Recall that we are restricting our attention to compact pwf surfaces without boundaries. The *Voronoi diagram* of P divides M into *Voronoi cells*, one for each $p \in P$, defined by $\mathcal{V}(p) = \{q \in M | d_M(p,q) \leq d_M((s,q) \forall s \in P\}$.

Definition 3. *A* Voronoi vertex *is a point $q \in M$ that has three or more distinct geodesics realizing the shortest distance from q to P. A* Voronoi edge *is a curve C terminating at Voronoi vertices and such that every point q on C has exactly two geodesics realizing the minimum distance from q to P. C is called an* internal Voronoi edge *if both the minimal geodesics connect with the same vertex.*

An equivalent view of Voronoi edges and Vertices is via the immersed empty disk property: If $\phi : \bar{D} \to M$ is an immersed empty disk with centre c and with $\phi^{-1}(P)$ containing three or more points, then $\phi(c)$ is a Voronoi vertex. If $\phi^{-1}(P)$ contains exactly two points, p and q, then c lies on a Voronoi edge, and it is an internal edge if $\phi(p) = \phi(q)$.

According to this view each Voronoi vertex is associated with a face in the Delaunay tessellation via the immersed empty disk that defines them both. Thus there is a finite number of Voronoi vertices. However, a Voronoi vertex is not necessarily associated with distinct samples and a Voronoi edge may terminate at the same Voronoi vertex at both ends. For example an interior edge will generally terminate at a Voronoi vertex together with a loop Voronoi edge.

Voronoi edges are geodesics between Voronoi vertices. To see this, let $\phi : \bar{D} \to M$ be an immersed empty disk with $\{p,q\} = \phi^{-1}(P)$ and $c \in D$ the centre. So $\phi(c)$ lies on some Voronoi edge C. Since there are only two vertices on the boundary of $\phi(D)$, we can find some ϵ and (exploiting Lemma 1) another immersed empty disk $\phi' : \bar{D}' \to M$ with centre c', $d_{\mathbb{R}^2}(c,c') = \epsilon$ and with $\{p',q'\} = \phi'^{-1}(P)$ such that $\phi'(p') = \phi(p)$ and $\phi'(q') = \phi(q)$. Then any point $\tilde{c}$ on the line segment $[c,c']$ will be the centre of an immersed empty disk $\tilde{\phi} : \tilde{\bar{D}} \to M$ whose image is contained in $\phi(\bar{D}) \cup \phi'(\bar{D}')$ and thus has $\phi(p)$ and $\phi(q)$ as the only points of P on its boundary. i.e. $[c,c']$ lies on the Voronoi edge C. The image of $[c,c']$ under the joint mapping $\hat{\phi}$ (lemma 1) is geodesic, since $[c,c']$ is a geodesic in the plane.

Lemma 3. *A Voronoi cell is topologically a disk if and only if it contains no internal edges.*

Proof. Let $q \in \mathcal{V}(p)$ and assume that there are two minimal length geodesics, α and β connecting p with q. Suppose $\mathcal{V}(p)$ were a topological disk. Together α and β define a closed curve contained in $\mathcal{V}(p)$. Let U be the region bounded by α and β. Then there is an isometric embedding $\varphi : U \hookrightarrow \mathbb{R}^2$. But then $\varphi(U)$ would be a region in the plane bounded by two geodesics (line segments) between $\varphi(p)$ and $\varphi(q)$. Thus U must be empty and $\alpha = \beta$.

Conversely, if $\mathcal{V}(p)$ is not a disk then since it is compact it has a smallest closed geodesic through p in a non-trivial homotopy class [Leibon and Letscher 2000]. The midpoint on this loop then has two distinct geodesics realizing the minimal distance to p and so lies on an internal edge. □

Since a minimal closed geodesic in $\mathcal{V}(p)$ must pass through an internal edge, the interior of $\mathcal{V}(p)$ – that part which remains when we remove all Voronoi edges – is a topological open disk. Note also that we cannot have a Voronoi edge that is a closed loop not containing any Voronoi vertices. If such a loop were to exist, it would have to be the unique boundary between two Voronoi cells that were both topologically disks (Otherwise an internal or other edge would create a Voronoi vertex). Therefore M must have only two vertices and be topologically a sphere. If such a pwf surface exists, it certainly cannot be realized as a mesh and it will not concern us here.

These observations demonstrate that the Voronoi diagram can be viewed as a tessellation. The faces of the tessellation are the interiors of the Voronoi cells; those points $q \in M$ for which there is a single geodesic realizing the shortest distance from q to P.

We now turn our attention to the duality relationship between the Delaunay tessellation and the Voronoi diagram. A nice thing about pwf surfaces is that if $\phi : \bar{D} \to M$ is an immersed empty disk, and $\phi^{-1}(P) = \{p,q\}$, then there is a unique geodesic between $\phi(p)$ and $\phi(q)$ contained in the image of ϕ; it is the image of the line segment between p and q. In other words there is only one possible edge contained in an empty disk with two samples on its boundary. This is not true for more general surfaces.

The image of the centre of D lies on a Voronoi edge C. If $e = [\phi(p), \phi(q)]$ is the Delaunay edge defined by ϕ, then we say C is the Voronoi edge *associated* with e and vise versa. The following lemmas demonstrate that this association is exclusive.

Lemma 4. *There is a unique Delaunay edge associated with each Voronoi edge.*

Proof. Suppose that $e = [a, b]$ and $e' = [a, b]$ are two Delaunay edges associated with the Voronoi edge C. Let u and u' be the centres of the empty geodesic disks containing e and e' respectively. Now centred at every point between u and u' on C there is an empty immersed disk with a and b on its boundary. Two such disks, if they are sufficiently close to each other, must contain the same Delaunay edge (we can appeal to Lemma 1). Thus we can push the disk centre from u to u' always maintaining e in the empty disk and thus we must have $e' = e$. □

Lemma 5. *Different Voronoi edges are associated with distinct Delaunay edges.*

Proof. Let $e = [a, b]$ be a Delaunay edge and suppose that it is contained in two different empty immersed disks $\phi : \bar{D} \to M$ and $\phi' : \bar{D}' \to M$. By Lemma 1 we can assume that $D \cap D'$ contains a line segment whose image under the combined map $\hat{\phi}$ is e. We have $p, q \in \partial D \cap \partial D'$ with $\hat{\phi}(p) = a$ and $\hat{\phi}(q) = b$. Let c and c' be the centres of D and D' respectively. Then at any point between c and c' there is a disk $\tilde{D}$ that is contained in $D \cup D'$ and touching p and q on its boundary. The restriction of $\hat{\phi}$ to $\tilde{D}$ defines an immersed empty disk. Therefore there is no Voronoi vertex between c and c' and thus they lie on the same Voronoi edge. □

The results of this section can be summarized in the following theorem, which establishes a Voronoi-Delaunay duality on pwf surfaces.

Theorem 1. *Considered together with its internal edges, the Voronoi diagram of the vertices of a pwf surface is a tessellation. Further, the empty circumdisk property defines a one to one correspondence between the edges of the Voronoi diagram and the edges of the Delaunay tessellation.*

3.3 Proper Triangulations

It has been emphasized [Bobenko and Springborn 2005; Fisher et al. 2006] that the Delaunay tessellation of a piecewise flat surface may not yield a proper triangulation. A geodesic triangulation of a manifold surface S is *proper* if:

1. each triangle has three distinct vertices,
2. each vertex has valence at least three, and
3. two vertices are connected by at most one edge.

Proper triangulations correspond to what was called a strongly regular triangulation[1] in [Fisher et al. 2006]. Note that the physical triangulation defined by a manifold triangle mesh is necessarily proper[2].

[1]We choose not to employ this use of the word regular since it already has meaning in terms of the connectivity of a mesh. Further, computational geometers use the word regular to describe Delaunay triangulations in non-Euclidean metrics.

[2]The traditional definition of a triangulation of a surface S is a homeomorphism $K \to S$ where K is a simplicial complex. This means that triangles share no more than one edge and edges share no more than one vertex. A manifold triangle mesh is then a pwf surface triangulated on its vertex set and isometrically immersed in $\mathbb{R}^3$. It is quick to check that a triangulation so defined must be proper.

A Delaunay triangulation that is proper is called a proper Delaunay triangulation. The correspondence, established in the previous section, between the Voronoi diagram of the vertex set of a pwf surface and its Delaunay triangulation provides an alternate characterization of a proper Delaunay triangulation.

Let S be a Riemannian surface endowed with a set P of sample points. We say the Voronoi diagram of P on S is *well formed* if each Voronoi cell:

1. is topologically a disk (is not neighbours with itself),
2. has at least three distinct Voronoi neighbours, and
3. meets each Voronoi neighbour at a single contiguous edge.

In the plane, a Delaunay triangulation is always a proper triangulation. Likewise the Voronoi diagram of points on the plane is always well formed. It was shown in [Leibon and Letscher 2000] that if the set of samples on a smooth surface is in general position and is sufficiently dense, then there will be a unique and well defined proper Delaunay triangulation. For pwf surfaces, [Bobenko and Springborn 2005] have established that there always exists a unique Delaunay tessellation of the vertex set. However, the tessellation does not yield a proper Delaunay triangulation in general.

The lack of a proper Delaunay triangulation is an indication of a sparse set of vertices. This is easier to see in terms of Voronoi diagrams.

Theorem 2. *Let M be a pwf surface whose vertex set P is in general position. The Delaunay triangulation of P on M is proper if and only if the Voronoi diagram of P on M is well formed.*

Proof. The criteria for a proper triangulation have been enumerated to correspond with the enumeration of criteria for a well formed Voronoi diagram when the triangulation in question is a Delaunay triangulation. For each criterion for a well formed Voronoi diagram we establish that its violation leads to a violation of the corresponding criteria for a well formed Delaunay triangulation and vise versa. The proof for each case makes implicit reference to Theorem 1 from which the correspondence follows immediately.

1. If $\mathcal{V}(p)$ is not topologically a disk then it must contain an internal Voronoi edge and the corresponding Delaunay edge e will begin and end at p. Then a Delaunay triangle adjacent to e will not have three distinct vertices. Conversely, if a Delaunay triangle does not have three distinct vertices, then it has an edge e that terminates at the same point p at both ends. This edge corresponds to an internal edge in the Voronoi diagram and thus $\mathcal{V}(p)$ is not topologically a disk.

2. The number of neighbours of $\mathcal{V}(p)$ is equal to the valence of p in the Delaunay tessellation. Since P is in general position, the Delaunay tessellation is itself *the* Delaunay triangulation.

3. If $\mathcal{V}(p)$ and $\mathcal{V}(q)$ meet at two distinct edges, then they give rise to two distinct Delaunay edges $e = [p, q]$ and $e' = [p, q]$ that share the same endpoints. Conversely, if two distinct Delaunay edges meet at the same endpoints p and q, then, since P is in general position, they will each correspond to distinct Voronoi edges between $\mathcal{V}(p)$ and $\mathcal{V}(q)$.

Proof completed. □

This theorem allows us to use Voronoi diagrams to characterize the meshes that will admit a proper Delaunay triangulation. The dual picture does not yield additional information, but it gives another way to view the situation. For example, by considering the Voronoi diagram it may be easier to see how a refinement algorithm should proceed so as to produce a proper Delaunay triangulation.

4 Delaunay Meshes

In this section, we consider the construction of Delaunay meshes. Following our discussion in the previous section, a natural approach would be to "straighten" the geodesic arcs of an intrinsic proper Delaunay triangulation. In Section 4.1 we show using a simple example that this does not work in general. The reminder of the section is devoted to our edge swapping and refinement algorithm.

Much of our discussion will focus on closed meshes, where we recall that according to Lemma 2, an edge e in the interior of a manifold triangle mesh is locally Delaunay if and only if the sum of its opposite angles does not exceed π. If the angle sum exceeds π, we say that e is not locally Delaunay or NLD, for brevity. A closed manifold triangle mesh is Delaunay if all of its edges are locally Delaunay. Meshes with boundaries are handled specifically in Section 4.2.5.

4.1 Geodesic Triangulations to Meshes

We can construct a geodesic Delaunay triangulation of some (relatively) smooth surface as a preliminary step to producing a mesh: the vertices of the geodesic triangulation will be the vertices of the mesh, and the mesh connectivity is defined by the geodesic edges of the Delaunay triangulation. This is how the remeshing algorithm of [Peyré and Cohen 2003] works for example. It can be expected that most of the edges in a mesh produced in this way will be locally Delaunay if the sampling is adequate. However, no matter how dense the sampling is, the final mesh need not be a Delaunay mesh. The transformation from a smooth surface S to a piecewise linear mesh M comes at the cost of a geometric approximation error. This distortion can cause a geodesic Delaunay edge on the original surface to become an NLD edge when it is realized as an edge in M.

To construct an example of this, consider a planar quadrilateral $puqv$ such that all four sides are of equal length and the opposite angles are equal, i.e., $\angle upv = \angle uqv$ and $\angle puq = \angle pvq$. Suppose further that one of the diagonals is slightly shorter than the other. Specifically, let $|e| = |[p,q]| = l$ and $|e'| = |[u,v]| = l + \epsilon$. For the symmetric quad $puqv$, the longer diagonal edge e' is NLD since it is subtended by larger angles, i.e., $\angle upv + \angle uqv > \angle puq + \angle pvq$.

Consider a cylinder S of radius r. Allow the quad to hinge on the diagonal e' and place its four vertices on the cylinder so that e' is parallel to the axis of the cylinder (Figure 1). In the geodesic realization of the quad, the geodesic diagonal corresponding to e, drawn as the short circular arc between p and q in Figure 1, will have length $s = 4r \arcsin(\frac{l}{4r})$. Thus its length will be longer than that of the other diagonal in the surface of S, $|e'| = l + \epsilon$, if $\frac{l}{4r} > \sin(\frac{l+\epsilon}{4r})$, which is easily realizable. In this case, e' will be the locally Delaunay edge on the surface of the cylinder S, and could be present in the Delaunay triangulation of S, but it is NLD in its mesh realization M. As a result, mesh M would not be Delaunay.

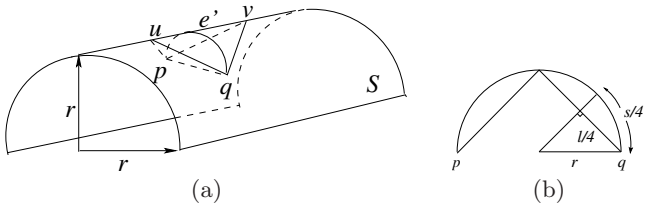

Figure 1: The cylinder example to illustrate the discrepancy between intrinsic Delaunay triangulations and Delaunay meshes. (a) In a quad with opposing angles equal and all sides equal, the longer diagonal edge is NLD. For the quad with planar faces this is edge $e' = [u,v]$, but in its geodesic realization on the cylinder, the other diagonal, the (geodesic) circular arc $[p,q]$, is longer and therefore NLD. (b) A cross-sectional profile of the cylinder at edge $e = [p,q]$. Detailed explanations are given in the text.

Note that a similar example could be constructed if S were a pwf surface. In other words, if we were to take a given mesh M and produce a new mesh M' with the same vertices, but with connectivity defined by the intrinsic Delaunay triangulation of M, then M' will not be a Delaunay mesh in general, even if M has a well formed Voronoi diagram.

Even if the sampling is sufficiently dense with respect to the local feature size r, the above problem can still occur as long as ϵ is sufficiently small. In practice, ϵ can indeed be arbitrarily small while the samples would still technically be in general position. A small ϵ corresponds to Voronoi vertices that are extremely close together.

Rather than placing constraints on the relative distance of Voronoi vertices of a sample set, it is presumably easier to perform a few post-processing edge swaps.

4.2 Edge Swapping and Refinement

The edge swapping algorithm mentioned in the previous section swapped geodesic arcs on a manifold of fixed geometry. In contrast, we now consider swapping the physical edges of a mesh in order to obtain a Delaunay mesh. This changes the nature of the problem considerably, for now the geometry of the underlying domain is changing each time an edge is swapped. In the case of fixed geometry, an edge could be identified as being Delaunay by the empty circumdisk property and its Delaunay status could not be changed by other swaps. Now this is no longer the case; an edge which is Delaunay at one instant may no longer be Delaunay in the mesh that results when a nearby edge is swapped.

4.2.1 Algorithm

Our edge swapping algorithm is superficially similar to the one described by [Bobenko and Springborn 2005], but now we have a different surface after each swap. Also, we have to deal with edges that are unswappable for topological reasons. An outline of our algorithm, which takes any manifold triangle mesh M as input, is given in Algorithm 1.

The order in which the edges are swapped is realized using a priority queue. The priority is related to the extent an edge is locally Delaunay and is set to be the sum of the opposite angles at the edge minus π. We ensure that at every step in the algorithm the current mesh, possibly refined, is a manifold triangle mesh. The edge swapping and refinement steps, along with discussions on termination and topological

```
while M contains an NLD edge do
  while M contains a swappable NLD edge e do
    M ← DelaunaySwap(M, e) (Section 4.2.2)
  end while
  while M contains an unswappable NLD edge e do
    M ← TetRefine(M, e) (Section 4.2.4)
  end while
end while
```

Algorithm 1: Edge swapping and refinement algorithm.

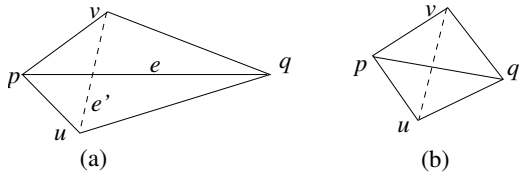

Figure 2: (a) A swap tet $puqv$ with opposing edges e and e'. (b) A regular tet shows that both diagonals $[p, q]$ and $[u, v]$ can be locally Delaunay.

issues leading to unswappable edges, will be described subsequently. Let us focus on closed meshes for now. Meshes with boundaries will be handled in Section 4.2.5.

4.2.2 Swappability and Uniqueness

In the planar case, an edge e is seen to be a diagonal of a quadrilateral formed by the two triangles adjacent to it. An edge swap entails replacing e with the other diagonal of the quadrilateral. In a geodesic triangulation of a piecewise flat surface, the edge e and its two adjacent faces can be isometrically unfolded onto the plane, so an edge swap can be interpreted the same way. However, when swapping a physical edge of a mesh, we cannot unfold the quadrilateral.

Consider an edge e connecting two vertices p and q in a mesh M. Suppose that e is adjacent to triangles $f_1 = [p, q, u]$ and $f_2 = [q, p, v]$. And consider the Euclidean line segment $e' = [u, v]$. The edges of f_1 and f_2 together with e' form a tetrahedron σ, as shown in Figure 2(a). We call σ the *swap tet* associated with e. Performing an edge swap on e involves replacing e with the new edge e' and faces f_1 and f_2 with faces $f_1' = [p, u, v]$ and $f_2'[q, v, u]$. We say that e' is the *opposing* edge to e.

Lemma 6. *If edge e in a (closed) mesh is not locally Delaunay, then its opposing edge e' is.*

This is easy to see since the sum of the interior angles of a space quad is at most 2π. Refer to Figure 2(a), we have

$$\angle puq + \angle uqv + \angle qvp + \angle vpu \leq 2\pi, \tag{1}$$

with equality holding only when the quad is planar. Note that Lemma 6 is true in this setting, but its converse, which holds for any planar quad in general position [de Berg et al. 1998], is not true in a mesh. The regular tet, as shown in Figure 2(b), gives an example where both edges e and e' are locally Delaunay. Consequently, there can be multiple Delaunay meshes on the same vertex set and defining the same topological surface. Thus in this sense, without demanding further qualifications, we do not have a general uniqueness theorem for Delaunay meshes, contrary to the case of fixed geometry, be it planar 2D or a fixed pwf surface.

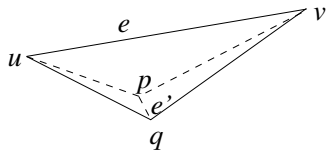

Figure 3: A Delaunay swap may decrease minimum angle and increase Harmonic index: Let e be NLD, then swapping e into e' is a Delaunay swap. However, one can bend $\triangle[u, q, v]$ towards $\triangle[u, p, v]$ so that $\angle puq$ becomes arbitrarily small and the Harmonic index of $\triangle[p, u, q]$ becomes arbitrarily large.

While Lemma 6 ensures that any swappable edge in our algorithm would improve matters locally, it does not quite lead to a termination proof. At the same time, an NLD edge may be unswappable if the resulting mesh would become non-manifold. We address these two issues next.

4.2.3 Delaunay Swaps and Area Minimization

Termination of edge swapping is traditionally shown by defining some functional on a triangulation and proving that it is increased (or decreased) with each swap. However, most of the traditional measures which apply to 2D or intrinsic Delaunay triangulations do not extend to the case of Delaunay swaps on a mesh. For example, the minimal angle in the triangles adjacent to an edge can be decreased after a Delaunay swap. Likewise, the harmonic index, exploited by [Bobenko and Springborn 2005], may increase after a swap. See Figure 3 for an example. It turns out that a measure that is consistently non-increasing with each Delaunay swap is the mesh surface area, which we prove below.

Theorem 3. *If the sum of the two angles opposite to an edge e is greater than the corresponding angle sum for the opposing edge e', then the combined area of the two triangles adjacent to e is greater than or equal to the combined area of the two triangles that would be adjacent to e'. Equality arises only when e and e' lie in the same plane.*

Proof. We are concerned with the area of the quadrilaterals defined by the two triangles adjacent to each edge. Note that the quadrilaterals can be made planar without distorting the area by unfolding them on the edge. The edges e and e' define two different quadrilaterals, but they share the same set of sides. Let a, b, c and d be the lengths of each of the sides. We exploit Bretschneider's formula [Bretschneider 1842] for the area of a quadrilateral $ABCD$:

$$\mathcal{A} = \sqrt{(s-a)(s-b)(s-c)(s-d) - abcd\cos^2\left(\tfrac{A+C}{2}\right)} \tag{2}$$

where $s = (a+b+c+d)/2$ is the semi-perimeter and A and C are angles opposite edge e. Let B' and D' be the angles opposite edge e' in the other quadrilateral.

Noting that $\cos^2\theta$ is monotonically decreasing in the interval $[0, \pi/2]$ and that $\frac{A+C}{2} > \frac{B'+D'}{2}$ by hypothesis, we have $\cos^2(\frac{A+C}{2}) < \cos^2(\frac{B'+D'}{2})$ if $\frac{A+C}{2} < \frac{\pi}{2}$. Thus, by equation (2), the area of the quadrilateral associated with e' is less than the area of that associated with e. On the other hand, if $\frac{A+C}{2} \geq \frac{\pi}{2}$, then by equation (1), we have $\pi/2 \geq \pi - \frac{A+C}{2} \geq \frac{B'+D'}{2}$ with equality in the planar case. Thus $\cos^2(\frac{A+C}{2}) = \cos^2(\pi - \frac{A+C}{2}) \leq \cos^2(\frac{B'+D'}{2})$.

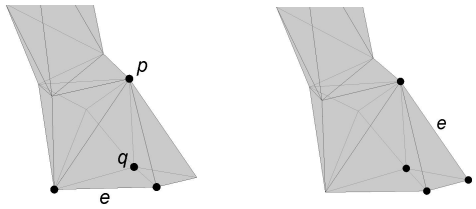

Figure 4: Unswappable swap tets. (a) 2-exposed. (b) 3-exposed. The tet is given by vertices indicated by the black dots.

Again equation (2) gives us a decrease in area except when e and e' lie in the same plane, in which case the area is unchanged. □

Thus, the surface area of the mesh is monotonically non-increasing as we run the edge swapping part of our algorithm. Since the number of possible triangulations is finite, the first inner loop will either terminate, or be stuck in an endless sequence of planar edge swaps. This latter possibility is eliminated by the termination proof of the planar Delaunay edge swapping algorithm [Lawson 1977]. Thus the first inner loop of Algorithm 1 terminates. Next, we discuss the handling of NLD edges that were left since they could not be swapped due to topological constraints; this is dealt with in the second inner loop by `TetRefine`.

4.2.4 Refining Unswappable Swap Tets

An edge e in M is *unswappable* if its opposing edge e' already exists in M. Swapping e into e' would result in a non-manifold edge. To see when this can happen, we examine the swap tet σ associated with e. All six edges of σ belong to M. At least two of the faces of σ, those adjacent to e, also belong to M. There are three possible cases:

1. **No other faces of σ belong to M:** We say that the swap tet is *2-exposed*, reflecting the number of faces shared by the mesh and the swap tet, as shown in Figure 4(a). If edge e is NLD, then it corresponds to sparse sampling around a thin structure, which can be expected to arise if there is a violation of the third criterion for a well formed mesh. Namely, the Voronoi cells $\mathcal{V}(p)$ and $\mathcal{V}(q)$ would connect at two disjoint edges, shown in Figure 5 in dark red.

2. **Only one of the faces adjacent to e' in σ belongs to M:** This would be the *3-exposed* case, where an edge is adjacent to a valence-three vertex, as shown in Figure 4(b). Let edge e be NLD with two opposite vertices p and q and let v be the valence-three vertex. This situation is conveniently depicted in Figure 5 as well, where we can observe that v has only two Voronoi neighbours, p and q, violating the second criterion of well-formedness defined in Section 3.3. Note that in the plane, a valence-three vertex cannot be incident to an NLD edge. On a surface, this situation indicates an inadequately sampled local feature, or a spike of noise.

3. **All four faces of σ belong to M:** This *4-exposed* case only happens when σ is the entire mesh M, assuming that M is a single connected mesh.

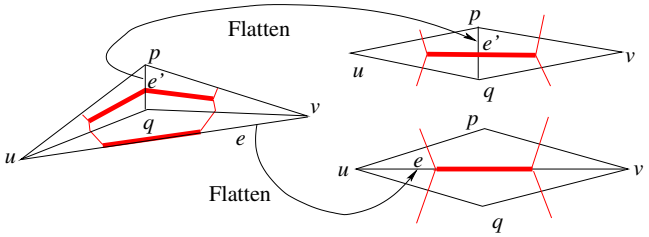

Figure 5: A 2- or 3-exposed swap tet $puqv$ with an NLD edge e. Lines in red are those from the Voronoi diagram. One way to view the situation is to suppose that the tet is 2-exposed. Then Voronoi cells $\mathcal{V}(p)$ and $\mathcal{V}(q)$ are incident at two Voronoi edges, marked in dark red. Alternatively, let the tet be 3-exposed with v being the valence-three vertex. Then v has only two Voronoi neighbors, p and q.

As we have just observed, unswappable NLD edges can be expected to arise when criteria of well-formedness are violated, implying that they are a manifestation of poor sampling. Naturally, our remedy would involve adding (or possibly removing) vertices. It is not hard to show that an NLD edge e can always be split into two locally Delaunay edges by inserting a vertex where e is crossed by the geodesic connecting the two vertices opposite to e; see Figure 6. We use this split to repair a 2-exposed NLD edge. Those bounding edges of the swap tet which become NLD after the split may be handled by subsequent swaps.

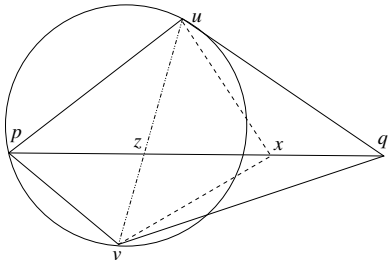

Figure 6: Splitting an NLD edge $e = [p, q]$. Let us fold the faces adjacent to e onto the plane. Clearly, edge $[p, x]$ will be locally Delaunay if and only if it lies inside the circumcircle of $\triangle[p, v, u]$. Likewise, x must lie inside the circumcircle of $\triangle[v, q, u]$ in order for $[x, q]$ to be locally Delaunay. Thus any point in the intersection of these two circumcircles will cut e into two locally Delaunay edges. A natural point to choose is z, the intersection of $[u, v]$ and $[p, q]$.

NLD edges that are 3-exposed edges must be dealt with more carefully. Suppose e is a 3-exposed edge adjacent to a valence-three vertex v. Note that e cannot be adjacent to two valence-three vertices, as that would imply that the swap tet is 4-exposed. If we split e at x as we would a 2-exposed edge, then $[x, p]$ may become a 3-exposed NLD edge again. Instead, we split all three edges adjacent to v at a distance s from v. We choose s to be half the length of the shortest edge adjacent to v. In this way the three new equal length edges adjacent to v are guaranteed to be Delaunay; see Figure 7. However, at least one of the new edges that are not adjacent to v can become a 2-exposed NLD edge, which may be repaired in the next iteration.

Both of these repair operations fix the problem edge without altering the geometry. An alternate means of dealing

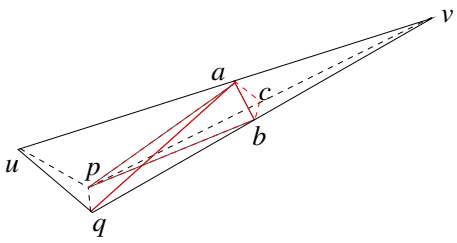

Figure 7: Edge $[u, v]$ is a 3-exposed NLD edge adjacent to the valence-three vertex v. The red lines indicate how it is repaired; a, b and c are all equidistant from v. Note that this repair has now created a 2-exposed NLD edge $[u, a]$.

with a 3-exposed NLD edge would be to simply remove the valence-three vertex. This may be appropriate if the vertex was considered to be a misplaced sample on a smooth surface and its removal would not alter topology. The latter would be the case for the 4-exposed case, where the entire mesh is a swap tet. In this case we would treat it in exactly the same way as the 3-exposed case, as illustrated in Figure 7.

Although our refinement schemes can be expected to remedy the problem of unswappable NLD edges, it remains to be shown that the algorithm outlined in Section 4.2 terminates. We do not have a proof at this point, but will pursue it in our future work. In our experiments on a few dozen distinct meshes, which we discuss in Section 5, we never encountered more than five iterations of the outer loop.

4.2.5 Boundaries and Features

In [Bobenko and Springborn 2005; Fisher et al. 2006] meshes with boundaries are dealt with by using a constrained Delaunay triangulation; there was not really another choice available in that context. However, we are by necessity allowing ourselves to add Steiner vertices. Therefore, we can also ensure that all the boundary edges are Delaunay.

Definition 4. *A boundary edge e is Delaunay if and only if it subtends a non-obtuse angle.*

If the angle is obtuse, the Voronoi cells of the endpoints will be separated by the Voronoi cell of the other vertex; see Figure 5 for an example. If this were the case, we would bisect the angle subtended by e, and split e where the bisector intersects. Similarly, feature edges can be identified and if a feature edge is not locally Delaunay, it can be split at the crossing point of the other geodesic diagonal, as described in Section 4.2.4. This provides a method for minimizing geometric distortion when creating Delaunay meshes from coarse meshes or meshes with ridge features.

5 Experimental Results

We have performed our edge swapping and refinement algorithm on a few dozen mesh models and statistics collected on some representative tests are reported in Table 1. The first group of models are well-known meshes which serve as examples of "typical" test datasets. In the second group, we choose several low-resolution meshes some of which, e.g., the knot, horse, and skull, have thin structures. These models have a greater percentage of 3- or 2-exposed tets compared with those from the first group and would provide a more rigorous test for the refinement component of our algorithm. The coarse knot, horse, and isis models were produced via the QSlim mesh simplification software of [Garland 1998]. The final group contains a few meshes that were produced by the remeshing algorithm of [Peyré and Cohen 2003]; these meshes are generally quite close to being Delaunay.

As can be observed from the results reported, our algorithm has consistently terminated after swapping only a small percentage of the total number of edges in the mesh, with only a small number of the outer loops of Algorithm 1 iterated. Thus the edge swapping and refinement scheme is seen to work quite efficiently in practice. We have not found a single case where Algorithm 1 failed to terminate. It is worth noting that although the Peyré hand (with a coarse sampling) and horse models approximate Delaunay geodesic triangulations of the original fine mesh surfaces, the meshes produced do have edges, though only a fraction, that are not locally Delaunay.

We have also collected statistics concerning the angle quality of the original and the resulting Delaunay meshes. As we have mentioned, the smallest angle is not required by theory to increase, but Table 1 shows that in practice it generally does with few exceptions. In all cases the size of the maximum angle, as well as the percentage of both small angles (less than 30°) and large angles (greater than 120°) decreased in the Delaunay mesh.

In the last column of Table 1, we report the approximation error of the Delaunay meshes produced, measured using the well-known Metro tool [Cignoni et al. 1998]. As we can see, the geometric approximation error tends to be large for coarse models. The coarse knot model has an abundance of two-exposed tets, the edge swapping algorithm is expectedly not doing a good job of respecting the geometry; Figure 8(a) and (b) show the original and the Delaunay coarse knot models. In contrast, the geometric error associated with the Delaunay versions of the corresponding full-resolution models is much smaller; see results for models from the first group.

We can reduce the geometric error by adding Steiner points rather than swapping edges when the dihedral angle is too large; this in effect enables "feature preservation". We threshold on the scalar product of the normals adjacent to the edge that is to be swapped. If this product is less than the threshold, we will split the edge rather than swapping it. In Figure 8(c) and Figure 9(c) we show the results when the threshold is 0.9. The edges are split where the other diagonal of the associated quadrilateral would cross it, as described in Section 4.2.4. As can be seen in the images, and the reported Metro error, this does a good job of reducing the geometric approximation error at an expense of increased vertex count. We can also observe that this refinement method is not altogether satisfactory. When the threshold is sufficiently close to 1, implying a severe restriction placed on allowable edge swaps, the algorithm often does not terminate. When the threshold is quite high we see regions of dense sampling indicating that this simplistic refinement scheme is working at its limit (the dark regions in Figure 8(c)).

6 Discussion and Future Work

We have established a Voronoi-Delaunay duality on piecewise flat (pwf) surfaces utilizing the empty circumdisk property. The notion of well-formed Voronoi diagrams was proposed and shown to lead to proper Delaunay triangulations on pwf surfaces. We developed an edge swapping and refinement algorithm which produces Delaunay meshes, a particularly useful geometric surface representation. Experiments

Mesh	#E	Swaps(# | %)		#P	# 3/2-exp	Min angle		Max angle		% Small		% Large		Error
knot	15000	1321	(8.8%)	1	0/0	20.5	20.5	95.9	95.9	7.4	7.3	0.0	0.0	0.0033%
horse	59547	2950	(5.0%)	5	4/4	1.7	1.6	171.3	164.3	8.7	7.9	1.3	1.2	0.474%
igea	165000	19349	(11.7%)	1	0/0	0.1	6.1	179.8	148.3	10.0	5.2	2.2	0.3	0.104%
bone	150000	9906	(6.7%)	4	2/3	0.1	5.8	179.7	150.5	8.2	6.1	1.3	0.4	0.102%
bigfish	41994	5744	(13.7%)	1	0/0	0.0	0.0	179.4	155.2	15.8	14.9	2.7	0.9	0.1%
coarse knot	1200	830	(69.2%)	5	0/238	0.3	12.0	178.7	123.1	42.7	7.6	17.4	0.2	2.37%
coarse horse	1050	216	(20.6%)	2	1/5	3.5	11.6	171.3	133.5	17.9	6.3	5.5	0.3	1.65%
coarse isis	1800	411	(22.8%)	1	0/0	2.8	8.9	173.3	137.1	20.3	9.2	5.8	0.5	1.63%
skull	33156	1952	(5.9%)	4	10/60	2.2	4.3	162.9	150.7	4.8	3.2	0.7	0.1	0.633%
Peyré hand	894	22	(2.5%)	1	0/0	24.7	31.0	129.9	112.5	0.6	0.0	0.2	0.0	1.4%
Peyré horse	1944	17	(0.9%)	3	1/1	12.0	18.1	136.4	113.8	0.2	0.2	0.1	0.0	0.87%

Table 1: Output statistics for our edge swapping and refinement algorithm on several mesh models. Mesh sizes are measured by $\#E$, the number of edges. We report the number of edge swaps performed until termination both in absolute number and as a percentage of the mesh size. The number of passes ($\#P$) refers to the number of times through the outer loop of the Algorithm 1. The fifth column indicates the number of 3-exposed and 2-exposed tets that were repaired, respectively. The next four double columns have before, on the left, and after, on the right, of the minimum (Min angle) and maximum (Max angle) face angles in the meshes, where angles are measured in degrees, as well as the percentage of angles that are less than 30° (% Small) and the percentage of angles that are more than 120° (% Large). The last column reports the geometric approximation error measured by the Metro tool, given as a percentage of the Hausdorff distance between the original and the output Delaunay mesh against the length of the bounding box diagonal of the former. The longest running time was on the bone model which took 6.3 seconds to process on a 2.4 GHz Opteron processor.

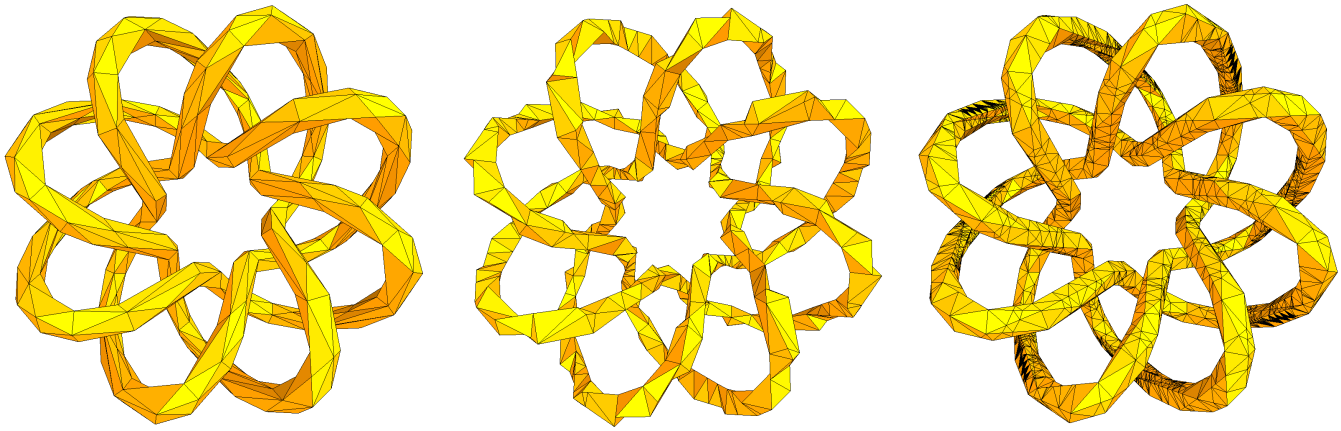

(a) (b) (c)

Figure 8: (a) The original coarse knot model with 400 vertices. (b) Delaunay mesh produced using Algorithm 1. The mesh has 638 vertices vertices and the Metro approximation error is 2.37%. (c) Delaunay mesh produced with "feature preservation" turned on. Namely, no edges with the dot product of their incident face normals exceeding the threshold of 0.9 can be swapped. Metro error is reduced to 0.365%, but it has 5088 vertices.

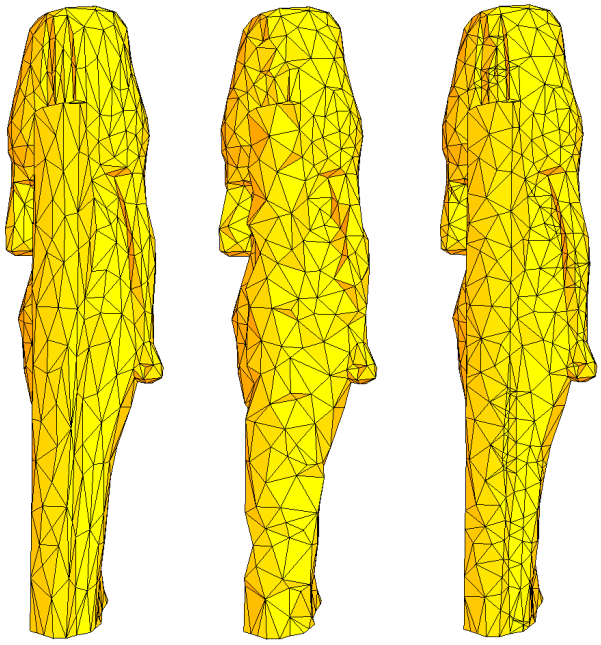

(a) (b) (c)

Figure 9: (a) The original coarse isis model. (b) Delaunay mesh produced using Algorithm 1: 602 vertices with Metro error 1.63%. (c) Delaunay mesh produced with "feature preservation" turned on (threshold = 0.9). Metro error is reduced to 0.65%, but with an addition 230 vertices produced.

show that our algorithm works efficiently and the resulting meshes possess desirable angle quality.

On geometric approximation errors: For a large class of modern meshes, edge swapping is a simple and fast way of turning a given mesh into a Delaunay mesh. For densely isotropically sampled smooth surfaces, the transformation to a Delaunay mesh comes with a small cost in geometric approximation error. From the angle statistics Delaunay meshes are well suited as a domain for the numerical solution of PDEs. It is exactly these well sampled smooth structures that are suited to most numerical computations anyway.

However, it is worthwhile to emphasize, that Delaunay meshes are not meant to improve the visual appearance of coarsely sampled meshes. This has become clear in the examples in Figure 8 as well as Figure 9. Furthermore, the example used in Section 4.1 (see Figure 1) also serves to emphasize that Delaunay meshes by no means minimize the geometric approximation error. The cylinder represents a ridge feature that could appear in a mesh and generally it is better to have the edges align with the direction of minimal absolute principal curvature [Dyn et al. 1990; Dyn et al. 2001]. Thus, in terms of geometry approximation the NLD edge turns out to be the better choice while its opposing edge takes a bite out of the ridge feature.

For efficient geometry approximation on a sparse vertex budget we would want anisotropic sampling – lower sampling along directions of low principle curvature – and a triangulation that aligns long skinny triangles accordingly. However, such a sampling is not appropriate for a numerical evaluation of differential operators, for example. For this task a denser isotropic sampling and a Delaunay mesh would be more appropriate.

Area minimization: The fact that Delaunay swaps are area reducing is intriguing. It seems likely that there are more interesting implications to be drawn. One avenue to explore is the possibility that a sequence of increasingly refined Delaunay meshes that converge point-wise to a given smooth surface will also have converging surface normals. Hildebrandt and Polthier [Hildebrandt et al. 2005] have shown that this is equivalent to surface area convergence.

The notion of well-formed Voronoi diagrams: The Voronoi conditions for the existence of a proper Delaunay triangulation could, for example, be useful in the generation of coarse base meshes for parameterization. While Leibon and Letscher [Leibon and Letscher 2000] have given sufficient conditions for the existence of a well defined proper Delaunay triangulation on a Riemannian surface, we have given *necessary and sufficient* conditions on a pwf surface. In fact, a well formed Voronoi diagram will be a necessary condition on any manifold surface, and this could be used to guide sampling algorithms such as the one presented by [Peyré and Cohen 2003].

Termination and complexity of edge swapping: There are a couple of gaps in our analysis of Algorithm 1. The most important one is the lack of a solid proof of termination. In preliminary tests we have found, that – by itself, without changing the geometry – the edge split refinement that we use to repair two-exposed tets to create a Delaunay mesh does not always converge if we allow no edge swaps at all. It would be good to employ a more sophisticated refinement technique, perhaps an adaptation of the one presented by Rivara and Inostroza [Rivara and Inostroza 1997]. Such an algorithm, with a termination guarantee would allow for the production of Delaunay meshes without geometry distortion. This could be substituted for the second inner loop in Algorithm 1. It would do away with the need for an outer loop.

Having a better understanding of the termination criteria for our algorithm will also aid a complexity analysis. Although the algorithm invariably terminates after swapping only a fraction of the total number of edges in practice, the first inner loop is presumably $\mathcal{O}(n^2)$ in theory. It would be good to have a proof of this for completeness.

References

Alboul, L., Brink, W., and Rodrigues, M. 2006. Mesh optimisation based on Willmore energy. In *22nd European Workshop on Computational Geometry*, 133–136.

Aleksandrov, A. D., and Zalgaller, V. A. 1967. *Intrinsic Geometry of Surfaces*, vol. 15 of *Transactions of mathematical monographs*. AMS.

Amenta, N., and Bern, M. W. 1998. Surface Reconstruction by Voronoi Filtering. In *Symp. Comp. Geom.*, 39–48.

Amenta, N., Choi, S., Dey, T. K., and Leekha, N. 2000. A Simple Algorithm for Homeomorphic Surface Reconstruction. In *Symp. Comp. Geom.*, 213–222.

Bobenko, A. I., and Springborn, B. A. 2005. A discrete Laplace-Beltrami operator for simplicial surfaces. arXiv:math.DG/0503219 v1.

Bretschneider, C. A. 1842. Untersuchung der trigonometrischen Relationen des geradlinigen Viereckes. *Archiv der Math. 2*, 225–261.

CHEN, H., AND BISHOP, J. 1997. Delaunay Triangulation for Curved Surfaces. In *6 Intl. Meshing Roundtable Proceedings, 1997*, 115–127.

CHEW, L. P. 1993. Guaranteed-quality mesh generation for curved surfaces. In *Symp. Comp. Geom*, 274–280.

CIGNONI, P., ROCCHINI, C., AND SCOPIGNO, R. 1998. Metro: Measuring error on simplified surfaces. *Computer Graphics Forum 17*, 2, 167–174.

DE BERG, M., VAN KREVELD, M., OVERMARS, M., AND SCHWARZKOPF, O. 1998. *Computational Geometry. Algorithms and Applications*. Springer-Verlag.

DESBRUN, M., MEYER, M., AND ALLIEZ, P. 2002. Intrinsic parameterizations of surface meshes. In *Proceedings of Eurographics*, 209–218.

DEY, T. K., LI, G., AND RAY, T. 2005. Polygonal surface remeshing with delaunay refinement. In *14th International Meshing Roundtable*, 343–361.

DO CARMO, M. P. 1976. *Differential Geometry of Curves and Surfaces*. Prentice-Hall.

DYN, N., LEVIN, D., AND RIPPA, S. 1990. Data Dependent Triangulations for Piecewise Linear Interpolation. *IMA Journal of Numerical Analysis 10*, 137–154.

DYN, N., HORMANN, K., KIM, S.-J., AND LEVIN, D. 2001. *Optimizing 3D triangulations using discrete curvature analysis*. Vanderbilt University, Nashville, TN, USA, 135–146.

EDELSBRUNNER, H., AND SHAH, N. R. 1994. Triangulating topological spaces. In *Symp. Comp. Geom.*, 285–292.

FISHER, M., SPRINGBORN, B., BOBENKO, A. I., AND SCHRÖDER, P. 2006. An algorithm for the construction of intrinsic Delaunay triangulations with applications to digital geometry processing. In *SIGGRAPH '06 Courses*, 69–74.

FLORIANI, L. D., MAGILLO, P., PUPPO, E., AND SOBRERO, D. 2004. A Multi-Resolution Topological Representation for Non-manifold Meshes. *Computer-Aided Design 36*, 2, 141–159.

GARLAND, M., 1998. The QSlim Mesh Simplification Software.

GU, X., AND YAU, S.-T. 2002. Computing conformal structures on surfaces. *Communications in Information and Systems 2*, 2, 121–146.

HILDEBRANDT, K., POLTHIER, K., AND WARDETZKY, M. 2005. On the Convergence of Metric and Geometric Properties of Polyhedral Surfaces. preprint.

KIMMEL, R., AND SETHIAN, J. A. 1998. Computing geodesic paths on manifolds. In *Proceedings of National Academy of Sciences*, vol. 95(15), 8431–8435.

LAWSON, C. L. 1977. Software for C^1 surface interpolation. Academic Press, New York, J. R. Rice, Ed., 161–194.

LEIBON, G., AND LETSCHER, D. 2000. Delaunay Triangulations and Voronoi Diagrams for Riemannian Manifolds. In *Proc. 16th Symposium on Comp. Geom.*, 341–349.

MEYER, M., DESBRUN, M., SCHRÖDER, P., AND BARR, A. H. 2003. Discrete Differential-Geometry Operators for Triangulated 2-Manifolds. In *Visualization and Mathematics III*. Springer-Verlag, 35–57.

MUSIN, O. R. 1997. Properties of the Delaunay triangulation. In *SCG '97: Proceedings of the thirteenth annual symposium on Computational geometry*, ACM Press, New York, NY, USA, 424–426.

ONISHI, K., AND ITOH, J. 2003. Estimation of the necessary number of points in Riemannian Voronoi diagram. In *Proc. 15th Canadian Conf. Comput. Geom.*, 19–24.

O'ROURKE, J. 1981. Polyhedra of minimal area as 3d object models. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 664–666.

PEYRÉ, G., AND COHEN, L. 2003. Geodesic Remeshing Using Front Propagation. In *Proceedings VLSM*, 33–40.

RIVARA, M.-C., AND INOSTROZA, P. 1997. Using Longest-Side Bisection Techniques for the Automatic Refinement of Delaunay Triangulations. *International Journal for Numerical Methods in Engineering 40*, 581–597.

RUPPERT, J. 1995. A delaunay refinement algorithm for quality 2-dimensional mesh generation. *Journal of Algorithms 18*, 3, 548–585.

SHEWCHUK, J. R. 1997. *Delaunay Refinement Mesh Generation*. PhD thesis, School of Computer Science, Carnegie Mellon University. Technical Report CMU-CS-97-137.

SIFRI, O., SHEFFER, A., AND GOTSMAN, C. 2003. Geodesic-Based Surface Remeshing. In *12th International Meshing Roundtable*, 189–199.

TURK, G. 1991. Generating textures on arbitrary surfaces using reaction-diffusion. In *ACM SIGGRAPH*, 289–298.

VAN DAMME, R., AND ALBOUL, L. 1995. Tight triangulations. In *Mathematical Methods for Curves and Surfaces*. Vanderbilt University Press, 517–526.