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

We define a *Delaunay mesh* to be a manifold triangle mesh whose edges form an *intrinsic Delaunay triangulation* or *iDT* of its vertices, where the triangulated domain is the piecewise flat mesh surface. We show that meshes constructed from a smooth surface by taking an iDT or a restricted Delaunay triangulation, do not in general yield a Delaunay mesh.

We establish a precise dual relationship between the iDT and the Voronoi tessellation of the vertices of a piecewise flat (pwf) surface and exploit this duality to demonstrate criteria which ensure the existence of a proper Delaunay triangulation.

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

Keywords: Delaunay triangulation (DT), intrinsic DT, restricted DT, Voronoi diagram, Delaunay mesh

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 to 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 (see [Shewchuk 1997]). Under certain conditions it can also be extended to non-Euclidean geometries. In particular the intrinsic Delaunay triangulation or iDT 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].

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

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. In general, the dual of the Voronoi diagram will not be a proper triangulation (simplicial complex). However, we observe that a proper Delaunay triangulation is guaranteed if and only if the Voronoi diagram satisfies the so-called closed ball property defined by [Edelsbrunner and Shah 1994] (see Section 3.3).

A manifold triangle mesh is an example of a pwf surface of practical interest. In general, the Delaunay triangulation of the vertex set of a triangle mesh does not coincide with the physical triangulation inherent in the mesh. We define a *Delaunay mesh* to be a manifold triangle mesh whose physical triangulation is also its Delaunay triangulation. We show that a mesh obtained from a smooth surface via an iDT or as a restricted Delaunay triangulation is not in general a Delaunay mesh. However, we argue that applying an edge swapping algorithm to a mesh so produced is an effective means of obtaining a Delaunay mesh.

1.1 Motivation

Intrinsic Delaunay triangulations of surfaces have made recent appearances in the geometry processing literature. Most notably, Bobenko and Springborn [Bobenko and Springborn 2005] 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. for constructing an injective parameterization [Floater 1998]. They advocate constructing an iDT of the 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 can be significantly improved in most cases. This improvement is beneficial to 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 sparked our interest in the notion of Delaunay meshes. As explained in [Fisher et al. 2006], constructing an iDT 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 were itself an iDT of its vertices: a Delaunay mesh.

We show that many modern surface reconstruction and remeshing algorithms produce meshes that are quite close to being Delaunay, thus only small adjustments need be made in order to obtain a Delaunay mesh as a final product.

1.2 Contributions

In this paper,

- we establish the duality between the Voronoi diagram and the Delaunay tessellation of a pwf surface (Section 3.2);
- using this duality, we characterize, in terms of the Voronoi diagram, meshes that admit a proper iDT (Section 3.3);
- we present a natural definition of a Delaunay mesh and show that it differs from meshes produced from a smooth surface via the iDT or as a restricted Delaunay triangulation (Section 4).

2 Related Work

In addition to the work of [Bobenko and Springborn 2005] on iDTs 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.

Although many works exist that employ Delaunay concepts and describe algorithms which produce meshes that are close to being Delaunay, to the best of our knowledge, no previous work has investigated Delaunay meshes as defined here. There are two well-known mesh constructions that have been referred to as Delaunay meshes: restricted Delaunay triangulations as studied by [Edelsbrunner and Shah 1994] and meshes obtained from an iDT of a surface, e.g., in [Peyré and Cohen 2003].

A remeshing algorithm from the first category was presented by [Chew 1993] where a Delaunay refinement technique was adapted to curved surfaces. This algorithm produced an approximate geodesic triangulation of the surface with a guaranteed angle bound of $[30°, 120°]$; intersections of the surface with a sphere are used in place of geodesic disks. The resulting mesh turns out to be a restricted Delaunay triangulation of the surface [Edelsbrunner and Shah 1994]. Restricted Delaunay triangulations have played an important role in surface reconstruction. A prominent example was presented by [Amenta and Bern 1998].

More recently, remeshing algorithms have appeared that are based on geodesic distances on a surface. [Peyré and Cohen 2003] described a farthest point sampling method based on geodesic distances. By taking the dual of the Voronoi diagram of the sampled points, a triangulation is produced that is a good approximation to the iDT of the sample points.

3 Delaunay Triangulations and Voronoi Diagrams on pwf Surfaces

In this section we review, in Section 3.1, the Delaunay triangulation on a pwf surface as defined by [Bobenko and Springborn 2005] and then show in Section 3.2 that its duality with the Voronoi diagram can be extended to this setting. In Section 3.3, we discuss the properties of the Voronoi diagrams of meshes that admit a proper Delaunay triangulation.

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 more care than is needed in the planar setting. For example, in the former, there may not be a unique shortest geodesic between two points, or 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 on the surface S.*

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.

For the remainder of Section 3 we confine our attention to a compact pwf surface M without boundary, on which the finite set of sample points, P, includes all of the cone points of M. We refer to the elements of P as *vertices*, emphasizing that the model pwf surface we have in mind is a mesh.

On a pwf surface every point p 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 (if p is a vertex). 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].

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 disk in $\mathbb{R}^2$ and $\bar{D}$ is the closure of D, such that the restriction $\phi|_D$ is an isometric immersion (i.e., every $p \in D$ has a neighbourhood which is mapped isometrically) and $\phi(D) \cap P = \emptyset$ (i.e., $\phi(D)$ is empty of vertices).

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 since 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 whose image is a given empty geodesic 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 $\phi = \phi' \circ T$.

Thus working with immersed empty disks is really equivalent to working with geodesic disks. However, the former allows us to 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 $\tilde{D}$ with $\tilde{D} \cap D \neq \emptyset$, 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 P on M 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 its convex hull, $\text{conv}[\phi^{-1}(P)]$, is a polygon and its image under ϕ defines a face of the tessellation. If $\phi^{-1}(P)$ contains exactly two points, then the image of $\text{conv}[\phi^{-1}(P)]$ under ϕ is an edge. It was established [Bobenko and Springborn 2005] that these faces and edges do 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 P on M is often referred to as the *intrinsic* Delaunay triangulation (*iDT*) to emphasize that it is based on geodesic distances on M 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. Edge 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 only two additional vertices of M are considered. If the edge e is not locally Delaunay, we call it *NLD*, as an abbreviation.

A convenient characterization of a locally Delaunay edge is given in [Bobenko and Springborn 2005]: 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 iDT can be obtained by systematically swapping the geodesic edges that are NLD. Namely, an NLD edge e is replaced by the edge e' that is the other diagonal (guaranteed to be locally Delaunay) 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 [Shewchuk 1997] holds without modification to the case of a fixed pwf surface.

For our purposes, we are primarily concerned with the iDT of the vertex set on a mesh. 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 iDT of the mesh vertices on its surface; we simply call it the iDT for brevity. The iDT 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 iDT 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. Thus the Hopf-Rinow theorem [do Carmo 1976] assures us that a minimal geodesic exists between any two points.

The *Voronoi diagram* of a set of samples P on a pwf surface (M, d_M) 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 minimum 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 some common vertex in P.*

Note that a point q in the interior of $\mathcal{V}(p)$ is characterized by having a unique minimal geodesic connecting it with p.

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.

To see that Voronoi edges are geodesics between Voronoi vertices, let $\phi : \bar{D} \to M$ be an immersed empty disk with $\{p, q\} = \phi^{-1}(P)$ and $c \in D$ the centre. Thus $\phi(c)$ 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 $\tilde{\phi} : \bar{\tilde{D}} \to M$ 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. In other words, $[c, c']$ 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 2. *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 that $\mathcal{V}(p)$ 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 thus 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.

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.

Furthermore, 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 3. *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' while always keeping e in the empty disk. As a result, we must have $e' = e$. □

Lemma 4. *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 centred at any point on the line 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 must lie on the same Voronoi edge. □

The results of this section are summarized in Theorem 1, establishing 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

A triangulation is *proper* if it is the realization of a simplicial complex. In general, the iDT on a pwf surface need not be proper, however we have shown [Dyer et al. 2007] that, if the vertices are in general position, a necessary and sufficient condition for the iDT to be proper is that the Voronoi diagram be *well formed*. That is, each Voronoi cell

1. is a topological disk;
2. meets each neighbour at a single contiguous edge;
3. has at least three distinct Voronoi neighbours.

In the study of restricted Delaunay triangulations, [Edelsbrunner and Shah 1994, lemma 4.1] showed that a manifold restricted Delaunay simplicial complex (see Section 4.2 for definitions) would result if the restricted Voronoi diagram satisfies the *closed ball property*. On a 2-manifold without boundary this yields three conditions [Amenta et al. 2000]: (1) Each Voronoi cell is a closed topological 2-ball, (2) The intersection of two Voronoi cells is either empty or a closed 1-ball and (3) The intersection of three or more Voronoi cells is either empty or a closed 0-ball (a single point).

The first two conditions are identical to the first two conditions of well formedness. In fact, a Voronoi diagram that does not satisfy the closed ball property is not well formed and vise versa: the criteria are equivalent. Indeed if $\mathcal{V}(p)$ has only two neighbours, $\mathcal{V}(q)$ and $\mathcal{V}(r)$, then the intersection of these three cells will consist of two distinct points: a violation of the third closed ball condition. If a cell has only one neighbour then its neighbour is not a topological disk: properties 1 and/or 2 are violated.

Conversely, if a Voronoi diagram does not satisfy the third closed ball condition, then there are three cells intersecting at more than one distinct point. Assuming the first two conditions are satisfied, this means that the intersection of any two of the cells is a topological 1-ball. So the intersection of $\mathcal{V}(p) \cap \mathcal{V}(q)$ with $\mathcal{V}(q) \cap \mathcal{V}(r)$ is two points; they must have common endpoints. If $\mathcal{V}(q)$ is a topological 2-ball, it can only have two neighbours: $\mathcal{V}(p)$ and $\mathcal{V}(r)$.

Note that the closed ball property also implies that the samples are in general position since if more than three

Voronoi cells are incident to a Voronoi vertex we would have two cells whose intersection is a single point.

The iDT of a pwf surface admits edges which loop around to terminate at a single point as well as multiple edges between two vertices. Such pathologies cannot appear in restricted Delaunay triangulations where the edges are strictly Euclidean line segments. Nonetheless, the closed ball property is necessary and sufficient to tame them both:

Theorem 2. *Let M be a pwf surface with vertex set P. The intrinsic Delaunay triangulation of P on M is proper if and only if the intrinsic Voronoi diagram of P on M satisfies the closed ball property.*

4 Delaunay Meshes

The restricted Delaunay triangulation and the iDT of a surface are two well-known means of producing triangle meshes. In this section, we investigate the difference between these constructs and Delaunay meshes as defined in our paper. Towards the end, we briefly discuss how their similarities can be exploited so as to produce Delaunay meshes.

4.1 Relation to Intrinsic Delaunay Triangulations

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 implied from the edges of the geodesic 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 M produced in this way will be locally Delaunay if the sampling is adequate. However, no matter how dense the sampling is, the final mesh M need not be a Delaunay mesh. The transformation from a smooth surface S to a piecewise linear mesh 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 $\angle puq = \angle pvq$. Thus $puqv$ is a diamond. Suppose further that one of the diagonals is slightly shorter than the other. Specifically, let $|e| = |[p,q]| = \ell$ and $|e'| = |[u,v]| = \ell+\epsilon$. 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 quadrilateral 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(a)). In the geodesic realization of the quadrilateral, the geodesic diagonal corresponding to e, drawn as the short circular arc between p and q in Figure 1(a), will have length $s = 4r\arcsin(\frac{\ell}{4r})$ (see Figure 1(b)). Thus its length will be longer than that of the other diagonal e' on the surface of S, where $|e'| = \ell+\epsilon$, as long as $\frac{\ell}{4r} > \sin(\frac{\ell+\epsilon}{4r})$. This is realizable for a sufficiently small ϵ, even though sampling density requirements may constrain the size of $\frac{\ell}{4r}$.

Indeed, as long as $\frac{\ell}{4r} > 0$ and in range, we can select ϵ such that $0 < \epsilon < \arcsin(\frac{\ell}{4r}) - \frac{\ell}{4r}$. In practice, ϵ can be arbitrarily small while the samples would still technically be in general position. A small ϵ corresponds to Voronoi vertices that are very close together. In this case, e' is the locally Delaunay edge on the surface of the cylinder S and

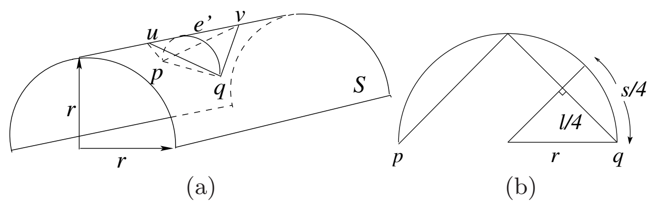

Figure 1: The cylinder example to illustrate the discrepancy between iDTs and Delaunay meshes. (a) In a quadrilateral with opposing angles equal and all sides equal, the longer diagonal edge is NLD. When the quadrilateral lies in the plane, this is edge $e' = [u,v]$, but in its geodesic realization on 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 $e = [p,q]$.

present in the Delaunay triangulation of S. But it is NLD in its mesh realization M, which would not be Delaunay.

Note that a similar example could be constructed if S were a pwf surface [Dyer et al. 2007]. 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 iDT of M, then M' would not be a Delaunay mesh in general, even if M has a well formed Voronoi diagram.

4.2 Relation to Restricted Delaunay Triangulations

The restricted Delaunay triangulation of a surface on a sample set is a subset of the 3D Delaunay tetrahedralization of the sample points and is a fundamental construct underlying many surface reconstruction algorithms. Using similar arguments as for iDTs we can show that the restricted Delaunay triangulation is not a Delaunay mesh in general.

The restricted Delaunay triangulation is defined in terms of the restricted Voronoi diagram. Let P be a set of samples on a surface S in $\mathbb{R}^3$. For $p \in P$, let $V(p) = \{x \in \mathbb{R}^3 |\, \|x-p\| \le \|x-q\|, \forall q \in P\}$ denote the 3D Voronoi cell of p. The *restricted Voronoi diagram* of P with respect to S is the partition of S into the closed cells $\mathcal{V}(p) = V(p) \cap S$, which are the restrictions of the 3D Voronoi cells to S.

The *restricted Delaunay triangulation* of P with respect to S is the dual of the restricted Voronoi diagram. This dual is made up of Euclidean simplices. Specifically, an edge $e = [p,q]$ belongs to the restricted Delaunay triangulation if and only if $\mathcal{V}(p) \cap \mathcal{V}(q) \ne \emptyset$. Note that e is the straight line segment connecting $p \in P$ and $q \in P$ in $\mathbb{R}^3$.

In order to construct an example where the restricted Delaunay triangulation yields an edge that is NLD in the resulting mesh, we again make use of the diamond-shaped quadrilateral $puqv$ as defined in Section 4.1. Let the quadrilateral be bent at the NLD edge $e' = [u,v]$ by an angle of 2α and inscribe the vertices u,v,p,q on an ellipsoid S such that e' is parallel to the principle axis of S, as shown in Figure 2(a). It is not hard to show that the circumcentre (centre of the circumsphere) of the tetrahedron $puqv$, which we denote by c, lies outside of the tetrahedron. Furthermore, c sits right above the midpoint of e' and let it be at a distance z away. Simple trigonometry yields $z = (\epsilon + \frac{\epsilon^2}{2\ell})/(2\sin\alpha)$.

Although sampling density constraints, i.e., those which respect Amenta's local feature size criteria [Amenta and Bern 1998], can force $\alpha > 0$ to be small, we can still make $z > 0$ arbitrarily small by choosing a sufficiently small ϵ (for fixed ℓ, α remains positive while ϵ goes to zero). Since there is a gap between the ellipsoid surface S and edge e', we

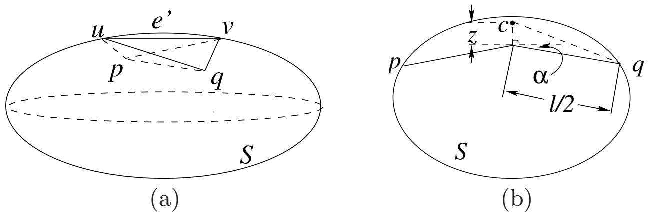

Figure 2: The ellipsoid example to illustrate the discrepancy between restricted Delaunay triangulations and Delaunay meshes. (a) The geometry of the tetrahedron $puqv$ inscribed in an ellipsoid S. (b) A cross-sectional profile of the ellipsoid perpendicular to edge e', where c is the circumcentre of the tetrahedron and it lies above e' but inside the ellipsoid.

can make c lie inside the ellipsoid, as shown in Figure 2(b). Consequently, S would pass above the circumcentre c.

Now consider the 3D Voronoi cells $V(u)$ and $V(v)$. The face f which separates them lies on a plane which is a perpendicular bisector of edge e'. This face extends upwards to infinity and must pass through the circumcentre c, as c is equi-distant to u, v, p and q. Since c is inside the ellipsoid S, the surface of S must intersect face f. It follows that the restricted Voronoi cells $\mathcal{V}(u)$ and $\mathcal{V}(v)$ are neighbours and the NLD edge $e' = [u, v]$ would appear in the restricted Delaunay triangulation; this completes our construction.

4.3 Producing Delaunay Meshes

The examples described in Sections 4.1 and 4.2 showed that the intrinsic and restricted Delaunay triangulations may yield edges that are NLD when four sample points lie close to a geodesic circle. This occurs when two Voronoi vertices are very close together, i.e., the samples are "almost not in general position". A Delaunay mesh could be guaranteed if a sampling criteria were developed that gave a lower bound on the distance between Voronoi vertices.

Alternatively, a Delaunay mesh may be produced with a post-processing step of edge swapping on one of these meshes. If the sample density is constrained by the local feature size [Amenta and Bern 1998] then the tetrahedra involved in the examples of Sections 4.1 and 4.2 will be slivers and the geometric distortion introduced by an edge swap will be small. Such edge swaps reduce the surface area of the mesh [Dyer et al. 2007] thus termination is assured. It is possible however that an edge is unswappable since its counterpart already exists in the mesh; such a situation occurs only in the presence of a Voronoi diagram that is not well formed, an indication of poor sampling [Dyer et al. 2007].

5 Discussion and Future Work

In studying the relationship between the intrinsic Delaunay triangulations, the restricted Delaunay triangulations, and Delaunay meshes we have shown that for a given surface S and a set of samples P on S, these objects can differ regardless of sampling density. However, suppose that S is a Delaunay mesh and P, its vertex set, is sufficiently dense so that the restricted Voronoi diagram is well formed. Then in this case, the restricted Delaunay triangulation and the intrinsic Delaunay triangulation both coincide with the triangulation defined by the physical edges of the mesh.

In future work we will investigate quantitative sampling criteria that would eliminate topological obstructions to Delaunay meshes. When such criteria are met, a Delaunay mesh can be obtained from an initial mesh by edge swapping [Dyer et al. 2007]. We are also working to develop a robust algorithm to produce Delaunay meshes from initial meshes which are not sufficiently well sampled.

References

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.

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

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.

do Carmo, M. P. 1976. *Differential Geometry of Curves and Surfaces*. Prentice-Hall.

Dyer, R., Zhang, H., and Möller, T. 2007. On Voronoi-Delaunay duality and Delaunay meshes. Tech. Rep. TR 2007-04, Simon Fraser University.

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.

Floater, M. S. 1998. Parametric tilings and scattered data approximation. *Int. J. of Shape Modeling 4*, 165–182.

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.

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.

Peyré, G., and Cohen, L. 2003. Geodesic Remeshing Using Front Propagation. In *Proceedings VLSM*, 33–40.

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

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