One-Pass Delaunay Filtering for Homeomorphic 3D Surface Reconstruction

TR99-08

Nina Amenta* Sunghee Choi†

April 12, 1999

Abstract

We give a simple algorithm for surface reconstruction from a set of point samples in R^3, using only one three-dimensional Voronoi diagram computation. We also give a fairly simple proof that the reconstruction is topologically correct when the input is a sufficiently dense sample from a smooth surface.

1 Introduction

We give an algorithm for fitting a surface triangulation to a set S of point samples in three dimensional space. We assume no additional information besides the three-dimensional coordinates of the points. Practical variants of this problem, in which more information might be given, arise in computer graphics, reverse engineering, medical imaging and computer vision.

Like many previous algorithms, our approach is to select some subset of the Delaunay triangles of S as the surface triangulation. This is a natural idea,

*Computer Sciences, University of Texas, Austin, TX 78712. Supported by NSF grant CCR-9731977.

†Computer Sciences, University of Texas, Austin, TX 78712. Supported by NSF grant CCR-9731977.

since Delaunay triangles connect points which are 'close' in a scale-invariant, combinatorial sense. There have been a number of proposals for 'filtering' three-dimensional Delaunay triangles to produce a surface [6],[12],[1].

We define two necessary criteria which a Delaunay triangle must meet in order to belong to a correct surface triangulation, if we assume that the input point set S is a dense enough sample from a smooth surface F. Adopting the definition used in a number of recent papers ([1],[2],[9],[11]), we consider a sample to be "dense enough" when the distance from any point on F to the nearest sample is proportional to the distance to the medial axis, with a small enough constant of proportionality. The *medial axis* of a surface F is defined as the closure of the set of points in R^3 with more than one nearest point on F. Unlike uniform sampling, this definition requires the sampling to be dense near small surface features (where the medial axis is close to the surface) but possibly sparse far away from any feature (where the medial axis is also far away). We consider F to be "smooth" when it is a twice differentiable closed manifold without boundary. Note that this implies that S is finite.

We call triangles meeting our two criteria *surface triangles*. The first criterion is that the normal of a surface triangle must be close to the surface normals (of the original smooth surface F) at its vertices. The second is that a surface triangle must be small, with respect to the distance to the medial axis at its vertices. Of course, given S alone and no other information about F, we cannot test these properties directly. We prove that, under the assumption that S is a dense enough sample from a smooth surface F, we can test the two criteria using the 3D Voronoi diagram of S.

We then explore the conditions under which a set T of surface triangles forms a manifold homeomorphic to F; one such condition is certainly that T must be a piecewise-linear manifold. We show that the function $\mu : T \rightarrow F$ mapping each point on T to the nearest point of F induces a homeomorphism under the additional reasonable condition that the angle between the normals of any two adjacent triangles of T is acute (so that T is 'smooth'). We show that the set of surface triangles always contains such a smooth piecewise-linear manifold T.

Finally, we sketch a simple algorithm for selecting T from the Delaunay triangulation of a sufficiently dense sample S. First, we *filter* the Delaunay triangles using our two criteria, and then we select T from the remaining triangles. We call the resulting triangulated manifold the *short crust* of S, since all of its triangles are small.

2 Previous work

The first author, with Marshall Bern and in part with David Eppstein and Manolis Kamvysellis, have considered this problem in a series of papers [1],[3],[2]. These papers describe a filtering algorithm for which the resulting set of triangles, the *crust* of S, is guaranteed to form a manifold close to the original surface F. The *short crust* algorithm of this paper uses several basic lemmas from [1]. This paper improves on the crust algorithm in two ways. First, the proof of correctness is considerably simpler than that offered for the crust algorithm [1]. Second, the algorithm itself is simpler and faster, since it eliminates a second-pass Delaunay triangulation step.

The idea of selecting a surface reconstruction from the 3D Delaunay triangulation is a venerable one. Boissonnat proposed two such algorithms in an early paper [6], which introduced the key idea of finding triangles with large empty circumsphere.

Edelsbrunner and Mücke [12] proposed the use of α-shapes for selecting Delaunay triangles to form a surface reconstruction. This idea is clearly provably correct when the sampling is uniformly dense, but not in any non-uniform model such as ours. While in many practical reconstruction problems the sampling *is* nearly uniform, none the less in practice finding an appropriate value of α is notoriously difficult.

In computer graphics, a different approach to the problem has predominated. Both Hoppe et al [16] and Curless and Levoy [8] used algorithms which reconstruct the surface as the zero-set of a distance function defined by the input point set. These methods are approximating rather than interpolating, and so far do not have well-defined sampling requirements or performance guarantees. They are, however, very fast and robust and are well-accepted in practice.

There has a been a lot of closely related work on reconstructing curves in the plane using Delaunay triangulation, much of it recent. See [19], [14], [18],

[4], [5], [9], [15], and [11]. Many of these algorithms come with theoretical guarantees.

3 Good triangles and dense enough sampling

In two dimensions, it is clear that the "right answer" to the reconstruction problem is a piecewise-linear curve connecting points that are adjacent along the original curve from which the samples were taken. It is not immediately obvious how to generalize this idea to three dimensions. We use a definition of the "correct" set of triangles, used by Chew [7] and by Edelsbrunner and Shah [13], which we shall call the set of *surface Delaunay triangles*. Consider the three-dimensional Voronoi diagram of S, and its intersection with F. The Voronoi diagram forms a partition of F into regions; this decomposition is the *restricted Voronoi diagram* of S in F. The surface Delaunay triangles are the triangles dual to the restricted Voronoi diagram. Equivalently,

Definition: A *surface Delaunay triangle* is a Delaunay triangle of S dual to an edge of the three-dimension Voronoi diagram of S intersecting F.

This definition makes sense even when an edge of the Voronoi diagram intersects F in multiple points (although in that case the surface Delaunay triangles may not form a piecewise-linear manifold). Of course, we cannot identify the surface Delaunay triangles from S alone, since the definition depends on F as well. A point set S might be a dense enough sample from two different surfaces F and F', and the set of surface Delaunay triangles with respect to F might differ from the set of surface Delaunay triangles with respect to F'.

We now define a "dense enough" sample. We make this definition with respect to a *Local Feature Size* function $LFS : R^3 \rightarrow R$, the definition of which again depends on the surface F.

Definition: For a point $x \in R^3$, $LFS(x)$ is the Euclidean distance from x to the nearest point on the medial axis of F.

The following lemma shows that the LFS function is Lipschitz.

Lemma 1 (Amenta and Bern [1]) *For any two points p and q on F,*

$|LFS(p) - LFS(q)| \leq d(p,q)$.

Intuitively, LFS will be small where two parts of the surface pass close together, since they will be separated by the medial axis. The medial axis is also close to the surface where the curvature is high, so LFS depends on curvature as well. The following lemma is a Lipschitz condition on the surface normal with respect to LFS.

Lemma 2 (Amenta and Bern [1]) *For any two points p and q on F with $d(p,q) \leq \rho \min\{LFS(p), LFS(q)\}$, for any $\rho < 1/3$, the angle between the normals to F at p and q is at most $\rho/(1-3\rho)$.*

We now define the sampling requirement.

Definition: A sample $S \subset F$ is an r-sample if the distance from any point $x \in F$ to the nearest sample point $s \in S$ is at most $r\ LFS(x)$.

We will see later that S is dense enough for our purposes when $r \leq .1$. We know from the following theorem that a correct output exists for the same value of r.

Theorem 3 (Amenta and Bern [1]) *If S is an r-sample of F for $r \leq .1$, then the surface Delaunay triangles form a polyhedron homeomorphic to F.*

We will be careful to choose filtering criteria which are met by all surface Delaunay triangles, so that the set of triangles which pass the filter is guaranteed to include such a polyhedron.

4 Triangles flat on the surface

In our search for filtering criteria, we consider only properties of the surface Delaunay triangles which can be inferred from S and its Voronoi diagram, without any additional knowledge of F. One such property is that the surface Delaunay triangles are nearly flat on the surface.

Lemma 4 (Amenta and Bern [1]) *Let t be a surface Delaunay triangle and s a vertex of t with angle at least $\pi/3$, and choose $r < 1/7$. (a) The angle* *between the normal to t and the normal to F at s is at most* $\arcsin \frac{\sqrt{3}r}{1-r}$. *(b) The angle between the normal to t and the normal to F at any other vertex of t is at most* $2r/(1-7r) + \arcsin \frac{\sqrt{3}r}{1-r}$.

Although we do not know any of the surface normals, we can approximate the normals at points in S from the Voronoi diagram of S. Informally, the idea is that when S is sufficiently dense, every Voronoi region is long and skinny and roughly perpendicular to the surface. The way we quantify this is to say that, given a sample s and a point v in its Voronoi region, the angle between the vector from s to v and the surface normal at s has to be small (linear in r) when v is far away from s (as a function of LFS).

Lemma 5 (Amenta and Bern [1]) *Let s be a sample point from an r-sample S. Let v be any point in $Vor(s)$ such that $d(v, s) \geq \rho LFS(s)$ for $\rho > \frac{r}{1-r}$. Let $\angle nv$ be the angle at s between the vector $\vec{v}$ to v and the surface normal $\vec{n}$ at s. Then* $\angle nv \leq \arcsin \frac{r}{\rho(1-r)} + \arcsin \frac{r}{1-r}$.

Conversely, if the angle is large, then point v has to be close to s. Specifically, if $\angle nv \geq \arcsin \frac{r}{\rho(1-r)} + \arcsin \frac{r}{1-r}$, then $d(v, s) \leq \rho LFS(s)$. Rearranging things, we get:

Corollary 6 *For any v such that $\angle nv = \alpha \geq \frac{r}{1-r}$, we have $d(v, s) \leq \rho LFS(s)$ with*

$$\rho = \frac{r}{(1-r)\sin(\alpha - \arcsin \frac{r}{1-r})}$$

Lemma 5 tells us that if we can find a point v in the Voronoi region which is sufficiently far away from s, $\vec{v}$ will be a good approximation of $\vec{n}$. In fact the farthest point in the Voronoi region is always sufficiently far away for this purpose. Consider extending a line segment perpendicularly in both directions from the surface at s, until it hits the medial axis in two points m^+, m^- (if it goes off to infinity, we consider that a medial axis point at infinity). These medial axis points are the centers of balls tangent to the surface at s with interiors empty of points of F. The points m^+, m^- are at least as close to s as to any other point on F, including of course all other points in S, and so must be contained in the Voronoi region of s. And since m^+, m^- are medial axis points, they are both at distance at least $LFS(s)$

from s. The farthest vertex p of the Voronoi region of s must then be at least that far away as well. We call p the *pole* of s [1].

Definition: If s does not lie on the convex hull of S, let the pole p be the vertex of $Vor(s)$ farthest from s, and let $\vec{p}$ be the vector from s to p. If s lies on the boundary of the convex hull of S, let $\vec{p}$ be the direction of any ray extending from s to infinity within the Voronoi region of s.

One way to choose $\vec{p}$ when s is on the convex hull is to average the outward-facing normals of the adjacent convex hull facets

Observation 7 *Let $\angle np$ be the angle between $\vec{n}$ and $\vec{p}$. Since $d(s,p) \geq LFS(s)$, Corollary 6 implies that $\angle np \leq 2 \arcsin \frac{r}{1-r}$.*

In summary, the normals of surface Delaunay triangles are close to the surface normals at their vertices, and those surface normals are in turn close to the vectors from the vertices to their poles. We therefore select triangles with normals close to the vectors to the poles.

Criterion 1 *Let s be the vertex of triangle t with largest angle, and let $\vec{t}$ be the normal vector to t. The angle $\angle tp$ between $\vec{t}$ and $\vec{p}$ may be at most $2 \arcsin \frac{r}{1-r} + \arcsin \frac{\sqrt{3}r}{1-r}$.*

Proceeding from Observation 7, we make the following observation about Criterion 1.

Observation 8 *The angle $\angle tn$ for a triangle meeting Criterion 1 is $O(r)$; in particular, $\angle tn \leq 4 \arcsin \frac{r}{1-r} + \arcsin \frac{\sqrt{3}r}{1-r}$.*

The angle between the normals given in Criterion 1 is reasonable; for instance, when $r \leq .1$, angle $\angle tn \leq .64$ radians.

[1] In the crust papers [1],[3] the *two* poles of s are defined to be the two farthest Voronoi vertices on either side of the surface. One pole suffices, however, for estimating the surface normal.

5 Small triangles

The other property of surface Delaunay triangles which we will use in the filtering process is that they are small with respect to LFS.

Lemma 9 *The radius of the circumcircle of a surface Delaunay triangle t is at most $\rho\, LFS(s)$, where s is any vertex of t, and $\rho = r/(1-r)$.*

Proof: Let v be the surface Voronoi vertex dual to t. The distance from v to s is at most $rLFS(v)$, which, by Lemma 1, is at most $r/(1-r)LFS(s)$. $\square$

Again, since we don't know F, we want to infer that a triangle has this property by examining the Voronoi region of a vertex s. One's first thought might be to use the fact that $d(s,p) \geq LFS(s)$, where p is the pole of a sample s. Unfortunately this is just an upper bound, and it is quite possible that $d(s,p)$ is *much* greater than $LFS(s)$, for instance when s is a point on the convex hull. Overestimating $LFS(s)$ of course would lead to accepting too many triangles as surface triangles.

We use, instead, an idea suggested by Tamal Dey [10] and used in two dimensional curve reconstruction by Dey, Mehlhorn and Ramos [11]. We require a triangle t to have an empty circumsphere B whose radius does not exceed that of the circumcircle of t by more than a small multiplicative factor; hence, t lies in a plane nearly bisecting B. We prove that any triangle which meets *both* this criterion and Criterion 1 must be small with respect to the LFS function at each vertex s.

Criterion 2 *Triangle t has a point v on its dual Voronoi edge such that the radius of the circumcircle of t is at least* $\cos(\arcsin \frac{r}{2(1-r)} + \arcsin \frac{\sqrt{3}r}{1-r})$ *times the radius of the circumsphere centered at v.*

Lemma 10 *Every surface Delaunay triangle meets Criterion 2.*

Proof: We show that the property holds for the circumsphere centered at the surface Voronoi vertex v dual to t. Let s be the vertex of t of largest angle. The distance from s to v is at most $\frac{r}{1-r}LFS(s)$. Since v is a point on the surface, it must lie outside the two balls of radius $LFS(s)$ tangent

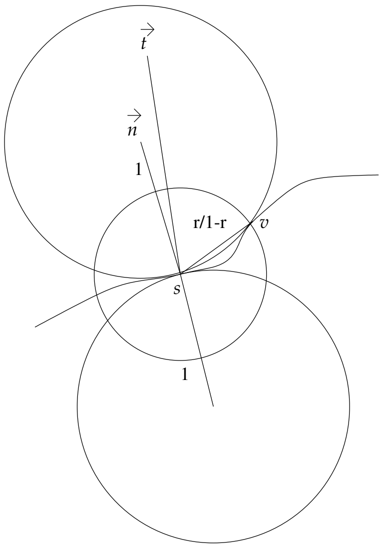

Figure 1: Proof of Lemma 10.

to the surface at s (since the tangent balls centered at the m_i are empty of surface points have radius at least $LFS(s)$). Assuming that v is pessimally positioned (see Figure 1) at the intersection of one of these large tangent balls with the ball of radius $\frac{r}{1-r}LFS(s)$ around s, the angle $\angle nv$ between the surface normal $\vec{n}$ and the vector $\vec{v}$ from s to v must be at least $\pi/2 - \arcsin(\frac{r}{2(1-r)})$. The angle $\angle nt$ is at most $\arcsin\frac{\sqrt{3}r}{1-r}$ by Lemma 4. Thus, $\angle tv \geq \pi/2 - (\arcsin\frac{r}{2(1-r)} + \arcsin\frac{\sqrt{3}r}{1-r})$. The lemma follows since the radius of the circumcircle of t is $\sin(\angle tv)$ times the radius of the circumcircle at v. $\square$

It is possible for a triangle which is quite large with respect to the LFS function at all of its vertices to meet Criterion 2. But any triangle which *also* meets Criterion 1 must be small.

Theorem 11 *The circumcircle of any triangle that meets Criteria 1 and 2 is at most $\rho LFS(s)$, where s is the vertex of t of largest angle, and $\rho = O(r)$.*

Proof: We let v be the center of the smallest circumsphere of t. We bound the distance from v to s, and hence the circumradius of t, using the bounds on the angles between the vectors $\vec{n}$, the surface normal at s, $\vec{v}$, the vector from s to v, $\vec{t}$, the normal to t, and $\vec{p}$, the vector from s to its pole p. Note that $\angle nv \geq \angle tv - \angle nt$, and $\angle nt \leq \angle np + \angle pt$.

We first show that $\angle nv$ must be large - $\pi/2 - O(r)$. For any triangle meeting Criterion 2,

$$\angle tv \geq \pi/2 - [\arcsin \frac{r}{2(1-r)} + \arcsin \frac{\sqrt{3}r}{1-r}]$$

For any triangle meeting Criterion 1,

$$\angle pt \leq 2 \arcsin \frac{r}{1-r} + \arcsin \frac{\sqrt{3}r}{(1-r)}$$

Observation 7 is that

$$\angle np \leq 2 \arcsin \frac{r}{1-r}$$

We put these together to find that

$$\angle nv \geq \pi/2 - [\arcsin \frac{r}{2(1-r)} + 2 \arcsin \frac{\sqrt{3}r}{1-r} + 4 \arcsin \frac{r}{1-r}]$$

Now we show that the distance $d(s, v)$ must be small, using Corollary 6 with $\alpha = \angle nv$, so that that $d(s, v) \leq \rho LFS(s)$ with

$$\rho = \frac{r}{(1-r)\cos(\arcsin \frac{r}{2(1-r)} + 2 \arcsin \frac{\sqrt{3}r}{1-r} + 5 \arcsin \frac{r}{1-r})}$$

Finally, the radius of the circumcircle of t might actually be somewhat smaller, $d(v, s)\cos(\pi/2 - \angle tv)$.

□

The constants in this theorem are again quite reasonable; for instance when $r = .1$, we get $\rho < .206$ [2].

We define a *surface triangle* to be one which meets Criterea 1 and 2. Note that all surface Delaunay triangles are surface triangles. We can infer from the preceeding theorem that all surface triangles are indeed close to the surface, as follows.

[2] Crust triangles were not be shown to be this small, and can in fact be quite large in flat regions.

Corollary 12 *Every point on any surface triangle is within $O(r)LFS(s)$ of some sample s.*

6 Mapping Surface Triangles to the Surface

In the next section, we will show a homeomorphism between F and any piecewise-linear surface T made up of surface triangles. We define the homeomorphism explicitly, using a function. We initially define a map μ on all of R^3, and then use its restriction to T.

Definition: Let $\mu : R^3 \rightarrow F$ map each point $q \in R^3$ to the closest point of F.

Lemma 13 *The restriction of μ to T is a well defined and continuous function $\mu : T \rightarrow F$.*

Proof: The discontinuities of μ as a map on R^3 are exactly the points of the medial axis. If some point q had more than one closest point on the surface, q would be a point of the medial axis; but every point $q \in T$ is within $O(r)LFS(s)$ of a triangle vertex $s \in F$, and hence can be nowhere near the medial axis. Similarly, μ is continuous except at the medial axis of F, and hence, since T is continuous and avoids the medial axis, μ is continuous on T.
$\square$

Observe that the segment connecting p to $\mu(p)$ is normal to F at $\mu(p)$.

The fundtion μ defines a homoemorphism between T and F if it is continuous, one-to-one and onto. Our approach will be first to show that μ is well-behaved on the samples themselves, and then show that this good behavior continues in the interior of each triangle of T. We begin with the following geometric lemma.

Lemma 14 *Let s be a sample and let m be the center of a medial ball B tangent to the surface at s. No surface triangle intersects the interior of the segment (s, m).*

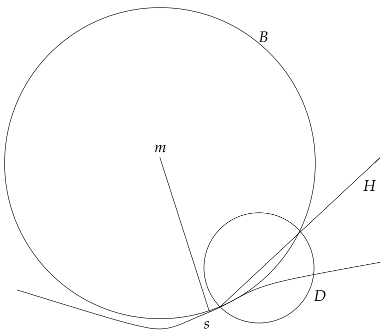

Figure 2: Proof of Lemma 14.

Proof: In order to intersect segment (s, m), a surface triangle t would have to intersect B, and so would the smallest Delaunay ball D of t. Since the vertices of t lie on F and hence not in the interior of B, the intersection of t and B must lie in the closed cap of B bounded by the plane H containing the intersection of the boundaries of B and D. We will show that H avoids the interior of (s, m).

Since D is Delaunay, s cannot lie in the interior of D. H can only intersect the interior of (s, m), then, if D contains m (see Figure 2.) But this is impossible because m is a point of the medial axis, so that the radius of D would be at least $1/2$ $LFS(s')$ for any vertex s' of t, contradicting, by Theorem 11, the assertion that t is a surface triangle.
□

Since any point q such that $\mu(q) = s$ lies on such an open segment (s, m), we have the following.

Corollary 15 *The function μ is one-to-one from T to every sample s.*

In the following section, we will show that μ is indeed one-to-one on all of T. One more geometric preliminary. We already know that the normal of a surface triangle t is close to the surface normals at its vertices (Observation 8). To complete the proof of homeomorphism, we need to show something a little stronger: that the triangle normal agrees with the surface normal at $\mu(q)$ for every $q \in t$.

Lemma 16 *Let q be a point on triangle $t \in T$. The angle between the surface normal $\vec{n}_q$ at $\mu(q)$ and the triangle normal $\vec{t}$ measures at most $O(r)$ radians.*

Proof: The circumcircle of t is small; the distance from q to the vertex s of t with largest angle is $\rho\, LFS(s)$, with $\rho = O(r)$, by Theorem 11. Choosing $r \leq .1$ gives $\rho \leq .206 \leq 1/3$. Substituting ρ into Lemma 2 gives the result. □

7 Homeomorphism

In this section, we show that building a manifold out of surface triangles is sufficient for reconstruction. Let T be a piecewise-linear manifold made up of surface triangles. Since all surface triangles are small, T is everywhere close to F. Under an additional mild assumption on T, we show that μ induces a homeomorphism between T and F.

Definition: A pair of triangles $t_i, t_2 \in T$ are *adjacent* if they share at least one common vertex.

Assumption: Two adjacent triangles meet at their common vertex at an angle of greater than $\pi/2$.

This assumption excludes manifolds which contain sharp folds and, for instance, flat tunnels.

Our proof proceeds in three short steps. We show that μ induces a homeomorphism on each triangle, then on each pair of adjacent triangles, and finally on T as a whole.

Lemma 17 *Let U be a region contained within one triangle $t \in T$. The function μ defines a homeomorphism between U and $\mu(U) \subset F$.*

Proof: We know that μ is well-defined and continuous on U, so it only remains to show that it is one-to-one. For a point $q \in t$, the vector $\vec{n_q}$ from $\mu(q)$ to q is perpendicular to the surface at $\mu(q)$; since F is smooth the direction

of $\vec{n_q}$ is unique and well defined. If there was some $y \in t$ with $\mu(y) = \mu(q)$, then q, $\mu(q)$ and y would all be colinear and t itself would have to contain the line segment between q and y, contradicting Lemma 16, which says that the normal $\vec{t}$ of t is nearly parallel to $\vec{n_q}$.
□

Lemma 18 *Let U be a region contained in adjacent triangles of T. The function μ defines a homeomorphism between U and $\mu(U) \subset F$.*

Proof: Let q and y be any two points in U, and let v be the common vertex of the triangles containing U. Lemma 17 implies that if $\mu(q) = \mu(y)$ we can assume that q and y lie in the two distinct triangles t_q and t_y. Let $\vec{n}$ be the surface normal at $\mu(q) = \mu(y)$. Since the ray supported by $\vec{n}$ passes through both t_q and t_y, and the angles $\angle t_q n, \angle t_y n = O(r)$ (Lemma 16), then t_q and t_y must meet at v at an acute angle. This would contradict the Assumption, which is that t_q and t_y meet at v at an obtuse angle. Hence there are no two points in U such that $\mu(q) = \mu(y)$.
□

Finally, in the following theorem, we bring out the topological guns.

Theorem 19 *The mapping μ defines a homeomorphism from the triangulation T to the surface F.*

Proof: Let $F' \subset F$ be $\mu(T)$. We first show that (T, μ) is a covering space of F'. (We relay on the treatment of covering spaces in Massey [17], Chapter 5.) Informally, (T, μ) is a covering space for F' if function μ maps T smoothly onto F', with no folds or other singularities. Showing that (T, μ) is a covering space is weaker than showing that μ defines a homeomorphism, since, for instance, it does not preclude several connected components of T mapping onto the same component of F', or more interesting behaviour, such as a torus under the map wrapping twice around another torus to form a *double covering.*

Formally, the (T, μ) is a covering space of F' if, for every $x \in F'$, there is a path-connected *elemenary neighborhood* V_x around x such that each path-connected component of $\mu^{-1}(V_x)$ is mapped homeomorphically onto V_x by μ.

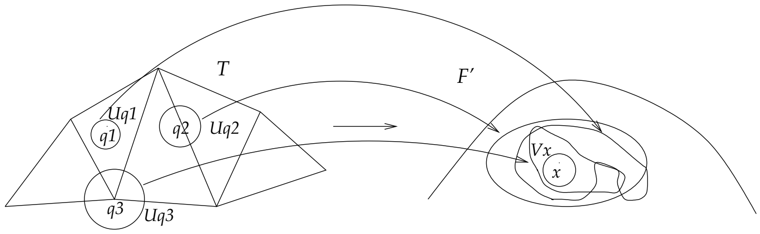

Figure 3: Proof of Theorem 19.

To construct such an elemenary neighborhood, note that the set of points $|\mu^{-1}(x)|$ corresponding to a point $x \in F'$ is non-zero and finite, since μ is one-to-one on each triangle of T and there are only a finite number of triangles. For each point $q \in \mu^{-1}(x)$, we choose an open neighborhood U_q of around q, homeomorphic to a disk and small enough so that U_q is contained only in triangles that contain q.

We claim that μ maps each U_q homeomorphically onto $\mu(U_q)$. This is because μ is continuous, it is onto $\mu(U_q)$ by definition, and, since any two points x and y in U_q are in adjacent triangles, it is one-to-one by Lemma 18.

Now consider the intersection $U'(x) = \cap_{q \in \mu^{-1}(x)} \mu(U_q)$, the intersection of the maps of each of the U_q. $U'(x)$ is the intersection of a finite number of open neighborhoods, each containing x, so we can find an open disk V_x around x. V_x is path connected, and each component of $\mu^{-1}(V_x)$ is a subset of some U_q and hence is mapped homeomorphically onto V_x by μ. Thus (T, μ) is a covering space for F'.

We now show that μ defines a homeomorphism between T and F'. Since T is onto F' by definition, we need only show that μ is one-to-one. Consider one connected component G of F'. A theorem of algebraic topology (see Massey [17], Chapter 5 Lemma 3.4) says that when (T, μ) is a covering space of F', the sets $\mu^{-1}(x)$ for all $x \in G$ have the same cardinality. We now use Corollary 15, that μ is one-to-one at every sample. Since each connected component of F contains some samples, it must be the case that μ is everywhere one-to-one, and T and F' are homeomorphic.

Finally, we show that $F' = F$. F' is closed and compact since T is closed and compact. So F' cannot include part of a connected component of F, and F' must consists of a subset of the connected components of F. Since ev-

ery connected component of F contains a sample s (actually many samples), and $\mu(s) = s$, all components of F belong to F', $F' = F$, and T and F are homeomorphic.
$\square$

8 Algorithm

Finally, we sketch a simple algorithm for selecting a piecewise-linear surface which meets Assumption 7. [3] We note, however, that this is not a practical algorithm; it can fail catastrophically when the input point set is not a dense enough sample from a smooth surface. We include it here only to complete the theoretical proof that we can produce a correct reconstruction given a sufficiently good sample. In practice, other heuristics should be used.

Let T' be the set of surface triangles. T' includes the surface Delaunay triangles, but might well be a superset, since S might be an r-sample for two different surfaces F and F', each inducing a different set of surface Delaunay triangles, both of which are guaranteed to be in T'.

To ensure that our output surface T will obey the Assumption that all dihedral angles are obtuse, we greedily remove all triangles adjacent to *sharp edges*. Define a *sharp edge* to be one which has a dihedral angle greater than $3\pi/2$ between a successive pair of incident triangles in the cyclic order around the edge. In other words, a sharp edge has all of its adjacent triangles within a small wedge. We consider an edge bounding only one triangle to have a dihedral of 2π, so such an edge is necessarily sharp. (Notice that if we greedily remove sharp edges from a set of triangles which does not contain a closed manifold, we might end up removing every triangle; this is the catastrophic failure mode.)

Let T'' be the set of triangles remaining after every triangle adjacent to a sharp edge has been removed. Since T'' has no sharp edge, every edge on the outside of T'' has two neighbors, so the outside of T'' is a piecewise-linear manifold. We let T be the outside surface of T''; we can find T, for example,

[3]This algorithm is essentially the same as the "manifold extraction" step of the crust algorithm.

by depth-first search on the outer triangles of every connectected component of T''.

Lemma 21 below guarantees that T'' still includes the surface Delaunay triangles, and hence that every sample s is still contained in some triangle in T''. Since no surface triangle intersects the line segment from s to its outside medial axis point, (Lemma 14) every sample appears on the outside of T''. So T includes every sample s.

It therefore remains only to prove Lemma 21. We begin with a simple technical lemma, which says that any line which meets F in two points close together must be nearly parallel to the surface.

Lemma 20 *A line intersecting F in two points x, x', such that $d(x, x') \leq O(r)LFS(x)$, must meet the surface normal at x at an angle of at least $\pi/2 - O(r)$.*

Proof Sketch: The point x' must lie outside the two tangent balls of radius $LFS(x)$ at x, and must be near x.
□

Now we prove the lemma.

Lemma 21 *No surface Delaunay triangle has a sharp edge.*

Proof Sketch: Let t and t' be adjacent surface Delaunay triangles, and let e be their shared edge. If t and t' meet at e in an angle of at least $\pi/2$, then e cannot be a sharp edge, even with respect to other triangles adjacent to e.

Since t and t' are surface Delaunay triangles, they have circumspheres B and B', respectively, centered at points v, v' of F. The boundaries of B and B' intersect in a circle C contained in a plane H, with H containing e. H separates t and t', since the third vertex of each triangle must lie on the boundary of its circumsphere, and $B \subseteq B'$ on one side of H, while on the other $B' \subseteq B$.

Both circumspheres pass through C, so their centers lie on a line perpendicular to H. Since they are the circumcenters of surface Delaunay triangles, the two centers are both within $O(r)LFS(s)$ of s (using the sampling assumption and Lemma 1). Hence $d(v, v') \leq O(r)LFS(v)$, and the surface normal at v

is within $O(r)$ radians of the surface normal at s. So the line l between v and v' must be nearly perpendicular to the surface normal $\vec{n}$ at s - the angle $\angle ln$ is $\pi/2 - O(r)$ (using Lemma 20 and Lemma 2). Hence the angle between H and $\vec{n}$ is at most $O(r)$. Since t and t' are flat to the surface at s, and they lie on opposite sides of H, the angle between them cannot be sharp.
$\square$

9 Conclusions and future work

We have given improved filtering criteria for selecting triangles from a Delaunay triangulation of a dense enough sample from a smooth surface to form a piecewise-linear reconstruction of the surface. We have also given a reasonably simple proof that such a reconstruction is indeed homeomorphic to the original surface.

In practice, the input point set S usually fails to be sufficiently dense near sharp edges and corners, and often it samples a surface F which is a manifold with boundary rather than a closed manifold. Our experience with the crust algorithm leads us to believe that the filtering criterea given here should be fairly robust in these situations. The actual reconstruction algorithm, unfortunately, while technically correct, relays strongly on the assumption that F is a closed manifold. We hope in the future to provide reconstruction algorithms that are more robust and practical with the help of the simpler theoretical framework given here.

Other important goals in this area are to correctly reconstruct surfaces with sharp edges and corners, and to develop reconstruction algorithms that gracefully handle noise and incremental reconstruction algorithms that can avoid examining all of the input data.

Acknowledgements We would like to thank John Havlicek for pointing us to related literature and for reading the paper and giving us insightful comments.

References

[1] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. Submitted to *Discrete and Computational Geometry*. An extended abstract appeared in the *14th ACM Symposium on Computational Geometry*, 1998, 39–48.

[2] N. Amenta, M. Bern, and D. Eppstein. The crust and the β-skeleton: combinatorial curve reconstruction. To appear in *Graphical Models and Image Processing*.

[3] N. Amenta, M. Bern, and M. Kamvysselis. A new Voronoi-based surface reconstruction algorithm. To appear in *Siggraph* 1998.

[4] D. Attali. r-Regular shape reconstruction from unorganized points. In *Proc. 13th ACM Symp. Computational Geometry*, 1997, 248–253.

[5] F. Bernardini and C. Bajaj. Sampling and reconstructing manifolds using α-shapes, *9th Canadian Conference on Computational Geometry*, 1997, 193–198.

[6] J-D. Boissonnat. Geometric structures for three-dimensional shape reconstruction. *ACM Trans. Graphics* 3 (1984) 266–286.

[7] Chew, L.P., Guaranteed-quality mesh generation for curved surfaces, *Proceedings of the ACM Symposium on Computational Geometry*, (1993), pp. 274-280.

[8] B. Curless and M. Levoy. A volumetric method for building complex models from range images. *Proc. SIGGRAPH '96*, 1996, 303–312.

[9] T. Dey and P. Kumar. A simple provable algorithm for curve reconstruction. *Proc. 10th Annual ACM-SIAM Symp. on Discrete Algorithms*, 1999, S893-4.

[10] T. Dey, *personal communication.*

[11] T. Dey, K. Mehlhorn and E. Ramos. Curve reconstruction: connecting the dots with good reason. To appear in *SoCG '99*

[12] H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graphics* 13 (1994) 43–72.

[13] H. Edelsbrunner and N. Shah. Triangulating topological spaces. *Proc. 10th ACM Symp. Computational Geometry*, 1994, 285–292.

[14] L. H. de Figueiredo and J. de Miranda Gomes. Computational morphology of curves. *Visual Computer* 11 (1995) 105–112.

[15] C. Gold. Crust and anti-crust: a one-step boundary and skeleton extraction algorithm. To appear in *SoCG '99*

[16] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. *Proc. SIGGRAPH '92*, 1992, 71–78.

[17] W.S. Massey. *Algebraic Topology: An Introduction*, Springer-Verlag, Graduate texts in Mathematics 56, 1967.

[18] M. Melkemi. A-shapes and their derivatives, *manuscript*, (1997).

[19] O'Rourke, J., Booth, H. and Washington, R., Connect-the-dots: a new heuristic, *Computer Vision, Graphics and Image Processing*, **39**, (1984), pp. 258-266.

[20] R. C. Veltkamp. *Closed object boundaries from scattered points.* LNCS Vol. 885, Springer, 1994.