Discrete Comput Geom 22:481–504 (1999)

Discrete & Computational
Geometry

Surface Reconstruction by Voronoi Filtering*

N. Amenta[1] and M. Bern[2]

[1]Computer Sciences, University of Texas,
Austin, TX 78712, USA
amenta@cs.utexas.edu

[2]Xerox Palo Alto Research Center,
3333 Coyote Hill Rd., Palo Alto, CA 94304, USA
bern@parc.xerox.com

Abstract. We give a simple combinatorial algorithm that computes a piecewise-linear approximation of a smooth surface from a finite set of sample points. The algorithm uses Voronoi vertices to remove triangles from the Delaunay triangulation. We prove the algorithm correct by showing that for densely sampled surfaces, where density depends on a local feature size function, the output is topologically valid and convergent (both pointwise and in surface normals) to the original surface. We briefly describe an implementation of the algorithm and show example outputs.

1. Introduction

The problem of reconstructing a surface from scattered sample points arises in many applications such as computer graphics, medical imaging, and cartography. In this paper we consider the specific reconstruction problem in which the input is a set of sample points S drawn from a smooth two-dimensional manifold W embedded in three dimensions, and the desired output is a triangular mesh with vertex set equal to S that faithfully represents W. We give a provably correct combinatorial algorithm for this problem. That is, we give a condition on the input sample points, such that if the condition is met the algorithm gives guaranteed results: a triangular mesh with position and surface normals within a small error tolerance of W. The algorithm relies on the well-known constructions of the Delaunay triangulation and the Voronoi diagram.

This paper is an extension of previous work by Amenta et al. [1] on reconstructing curves in two dimensions. Our previous work defined a planar graph on the sample

* The work of the first author was performed in part while at Xerox PARC and was partially supported by NSF Grant CCR-9404113.

points called the *crust*. The crust is the set of edges in the Delaunay triangulation of the sample points that can be enclosed by circles empty not only of sample points, but also of Voronoi vertices. The crust comes with a guarantee: if the curve is well-sampled, then the crust contains exactly the edges between sample points adjacent on the curve. Our notion of well-sampled, which involves the medial axis of the curve, is sensitive to local geometry. Hence our algorithm, unlike other algorithms for this problem, allows highly nonuniform sampling, dense in detailed areas yet sparse in featureless areas. Any provably correct algorithm must impose some sampling density requirement, similar to the Nyquist limit in spectral analysis.

The extension to three dimensions in this paper requires both new algorithmic ideas and new proof techniques. Most notably the algorithm uses only a subset of the Voronoi vertices to remove Delaunay triangles. The algorithm picks only two Voronoi vertices—called *poles*—per sample point: the farthest vertices of the point's cell on each side of the surface. With this modification, the straightforward generalization of our two-dimensional algorithm now works. Delaunay triangles with circumspheres empty of poles give a piecewise-linear surface pointwise convergent to W. The poles also enable further filtering on the basis of triangle normals. Adding this filtering gives a piecewise-linear surface that converges to W both pointwise and in surface normals (and hence in area). We believe that poles may be useful in other algorithms, perhaps whenever one wishes to estimate a surface normal or tangent plane.

This paper is organized as follows. Section 2 describes previous work on surface reconstruction. Section 3 gives our algorithm. Section 4 states our theoretical guarantees, and Section 5 gives their proofs. Section 6 shows some example outputs.

2. Previous Work

Previous work on the reconstruction problem falls into two camps: computer graphics and computational geometry. The algorithms in use in computer graphics typically compute an approximating surface, that is, a surface passing close by, rather than exactly through, the original sample points. The algorithms devised by computational geometers typically compute an interpolating surface, that is, a surface passing through the sample points, usually a carefully chosen subset of the Delaunay triangulation.

The first and most widely known reconstruction algorithm in the computer graphics community is the work of Hoppe et al. [18]–[20]. This algorithm estimates a tangent plane at each sample using the k nearest neighbors, and uses the distance to the plane of the closest sample point as a signed distance function. The zero set of this function is then contoured by a continuous piecewise-linear surface using the marching cubes algorithm. A later algorithm by Curless and Levoy [11] is designed for data samples collected by a laser range scanner. This algorithm sums anisotropically weighted contributions from the samples to compute a signed distance function, which is then discretized on voxels to eliminate the marching cubes step. These algorithms appear to be successful in practice, but have no provable guarantees. Indeed there exist arbitrarily dense sets of samples, for example ones with almost collinear nearest neighbor sets, for which the algorithm of Hoppe et al. would fail.

The most famous computational geometry construction for associating a polyhedral

shape with an unorganized set of points is the α-shape of Edelsbrunner et al. [13], [14]. Like our reconstructed surface, the α-shape is a subcomplex of the Delaunay triangulation. A Delaunay simplex (edge, face, etc.) belongs to the α-shape of S if its circumsphere has radius at most α. The major drawback of using α-shapes for surface reconstruction is that the optimal value of α depends on the sampling density, which often varies over different parts of the surface. For uniformly sampled surfaces, however, α-shapes are workable. Bernardini et al. [5] follow α-shape-based reconstruction with a clean-up phase to resolve sharp dihedral angles. Edelsbrunner and Raindrop Geomagic Inc. report that they have since developed another, proprietary reconstruction algorithm.

An early algorithm due to Boissonnat [7] is related to ours. He proposed a sculpting heuristic for selecting a subset of Delaunay tetrahedra to represent a solid object. The heuristic is motivated by the observation that "typical" Delaunay tetrahedra have circumspheres approximating maximal empty balls centered at points of the medial axis; our algorithm relies on this same observation. Boissonnat's algorithm, however, overlooks the fact that even dense sample sets can give Delaunay tetrahedra with circumspheres that are arbitrarily far from the medial axis; indeed it is this second observation which motivates our definition of poles. Goldak et al. [17] made a similar oversight, asserting incorrectly that the Voronoi diagram vertices asymptotically approach the medial axis as the sampling density goes to infinity.

Along with the relation between Delaunay circumcenters and the medial axis, our algorithm uses another important idea already found in the literature. The *restricted Voronoi diagram* is the intersection of the three-dimensional Voronoi diagram of the sample points with the surface W. We call a Delaunay triangle a *good triangle* if it is dual to a vertex of the restricted Voronoi diagram. We use a theorem of Edelsbrunner and Shah [15] to show that when the sampling is sufficiently dense these good triangles do in fact produce a topologically correct reconstruction. Unfortunately, without knowing W, we can only produce a superset of the good triangles. In the case in which W is known, Chew [9] used the good triangles to produce a "surface Delaunay triangulation."

Finally, for the two-dimensional problem of reconstructing a curve in the plane, there has been a flurry of recent algorithms with provable guarantees. Figueiredo and Miranda Gomes [16] prove that the Euclidean minimum spanning tree can be used to reconstruct uniformly sampled curves in the plane. Bernardini and Bajaj [4] prove that α-shapes also reconstruct uniformly sampled curves in the plane. Attali [3] gives yet another provably correct reconstruction algorithm for uniformly sampled curves in the plane, using a subgraph of the Delaunay triangulation in which each edge is included or excluded according to the angle between the circumcircles on either side. Finally, the paper by ourselves and Eppstein [1] showed that both the crust and the β-skeleton [21] (another empty-region planar graph) correctly reconstruct curves even in the case of nonuniform sampling. Our two-dimensional results are in this way strictly stronger than those of the other authors.

3. Description of the Algorithm

We start by reviewing the algorithm of Amenta et al. [1] for the problem of reconstructing curves in $\mathbb{R}^2$. Let W be a smooth (twice differentiable) curve embedded in $\mathbb{R}^2$, and let S be a set of sample points from W. Let V denote the vertices of the Voronoi diagram of S.

The *crust* of S contains exactly the edges of the Delaunay triangulation of $S \cup V$ with both endpoints from S. Saying this another way, the crust contains exactly those Delaunay edges around which it is possible to draw a circle empty of Voronoi vertices. If S is a sufficiently dense sample, this simple algorithm constructs a polygonal approximation of W (Theorem 1 in Section 4 below).

The straightforward generalization of this algorithm fails for the task of reconstructing a smooth two-dimensional manifold embedded in three dimensions. The problem is that the Voronoi center of four roughly cocircular samples may fall very close to the surface, thereby punching holes in the crust.

The fix is to consider only the *poles*. The poles of a sample point s are the two farthest vertices of its Voronoi cell, one on each side of the surface. Since the algorithm does not know the surface, only the sample points, it chooses the poles by first choosing the farthest Voronoi vertex regardless of direction (or a fictional pole at infinity in the case of an unbounded Voronoi cell), and then choosing the farthest in the opposite half-space. See step 2 in Fig. 1. Lemma 6 in Section 5 shows that this method is indeed correct for well-sampled surfaces. Denoting the poles by P, we define the *crust* of S to be the triangles of the Delaunay triangulation of $S \cup P$, all of whose vertices are members of S.

Steps 1–4 compute the crust (which we shall sometimes call the *raw crust* to distinguish it from the more finished versions). The crust has a relatively weak theoretical guarantee: it is pointwise convergent to W as the sampling density increases. Step 5 produces an output with a stronger guarantee: convergence both pointwise and in surface normals. Step 6 thins the set of output triangles to produce a piecewise-linear manifold.

Step 5 removes triangles based on the directions of their surface normals. Let T be a triangle of the crust and let s be its vertex of maximum angle. Step 5 removes T if

1. Compute the Voronoi diagram of the sample points S.
2. For each sample point s:
   (a) If s does not lie on the convex hull of S, let p^+ be the vertex of $Vor(s)$ farthest from s.
   (b) If s does lie on the convex hull of S, let p^+ be a point at infinite distance outside the convex hull with the direction of sp^+ equal to the average of the outward normals of hull faces meeting at s.
   (c) Among all vertices p of $Vor(s)$ such that $\angle p^+sp$ measures more than $\pi/2$, choose the farthest from s to be p^-.
3. Let P denote all poles p^+ and p^-, except those p^+'s at infinite distance. Compute the Delaunay triangulation of $S \cup P$.
4. (Voronoi Filtering) Keep only those triangles in which all three vertices are sample points.
5. (Filtering by Normal) Remove each triangle T for which the normal to T and the vector to the p^+ pole at a vertex of T form too large an angle (greater than θ for the largest-angle vertex of T, greater than $3\theta/2$ for the other vertices of T).
6. (Trimming) Orient triangles and poles (inside and outside) consistently, and extract a piecewise-linear manifold without boundary.

Fig. 1. The surface reconstruction algorithm.

the angle between the normal to T and the vector from any one of T's vertices to its first-chosen pole is too large. The definition of "too large" depends on which vertex of T is under consideration: for the vertex with the largest angle, too large means greater than an input parameter θ, and for the other two vertices it means greater than $3\theta/2$. Angles are unsigned angles in the range $[0, \pi/2]$. As stated in Theorem 5, the choice of θ is connected with the sampling density. If the user of our algorithm does not have an estimate of the sampling density (the parameter r in Definition 3 below), then the user can slowly decrease θ, backing off when holes start to appear in the surface, similar to choosing a surface from the spectrum of α-shapes [14].

Step 6 ensures that the reconstructed surface is a manifold; before this final step, the computed surface will resemble the original surface geometrically, but may have some extra triangles enclosing small bubbles and pockets. The problem once again is cocircular samples: all four faces of a flat tetrahedron may make it past steps 4 and 5.

Step 6 first orients all triangles. Start with any sample point s on the convex hull of S. Call the direction to p^+ at s the *outside* and the direction to p^- the *inside*. Pick any triangle T incident to s, and define the outside side of T to be the one visible from points on the sp^+ ray. Orient the poles of the other vertices of T to agree with this assignment. Orient each triangle sharing a vertex with T so that they agree on the orientations of their shared poles, and continue by breadth-first search until all poles and triangles have been oriented. Theorem 5 below guarantees that this orientation is consistent.

We can now extract a piecewise-linear two-dimensional manifold from the crust. Define a *sharp* edge to be an edge 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 its 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.

Step 6 trims off pockets by greedily removing triangles with sharp edges. Now the remaining triangles form a "quilted" surface, in which each edge bounds at least two triangles, with consistent orientations. Finally, Step 6 extracts the outside of this quilted surface by a breadth-first search on triangles.

4. Theoretical Guarantees

What sets our algorithm apart from previous algorithms are its theoretical guarantees. We begin with the required sampling density, which is defined with respect to the medial axis.

Definition 1. The *medial axis* of a manifold W embedded in $\mathbb{R}^d$ is the closure of the set of points in $\mathbb{R}^d$ with more than one nearest neighbor on W.

Figure 2 gives an example of the medial axis in $\mathbb{R}^2$; in $\mathbb{R}^3$, the medial axis is generally a two-dimensional surface. Note that we allow the surface W to have more than one connected component.

Definition 2. The *local feature size* $LFS(p)$ at a point p on W is the Euclidean distance from p to (the nearest point of) the medial axis.

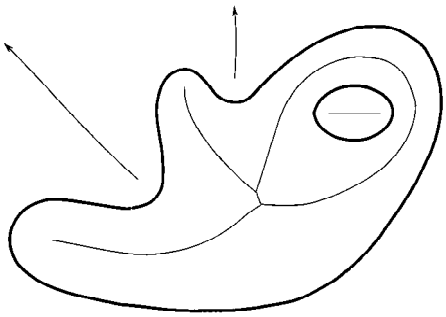

Fig. 2. The medial axis of a smooth curve.

Definition 3. Set $S \subset W$ is an *r-sample* of W if no point p on W is farther than $r \cdot LFS(p)$ from a point of S.

Definition 4. The *poles* of a sample s are the two farthest Voronoi vertices from s, one on either side of W.

Notice that the notion of r-sample does not assume any global—or even local—uniformity. Further notice that to prove an algorithm correct, we must place some condition on the set of sample points S, or else the original surface could be any surface passing through S. Our paper on curve reconstruction [1] proved the following theoretical guarantee.

Theorem 1 [1]. *If S is an r-sample of a curve in $\mathbb{R}^2$ for $r \leq 0.40$, then the crust includes all the edges between pairs of sample points adjacent along W. If S is an r-sample for $r \leq 0.25$, then the crust includes exactly those edges.*

To state our results for the three-dimensional problem, we must define a generalization of adjacency. Consider the Voronoi diagram of the sample points S. This Voronoi diagram induces a cell decomposition on surface W called the *restricted Voronoi diagram*: the boundaries of the cells on W are simply the intersections of W with the three-dimensional Voronoi cell boundaries. We call a triangle with vertices from S a *good triangle* if it is dual to a vertex of the restricted Voronoi diagram; good triangles are necessarily Delaunay triangles. Our first three-dimensional result shows that good triangles deserve their name. To our knowledge, our proof of this result is the first proof that the three-dimensional Delaunay triangulation of a sufficiently dense set of samples contains a piecewise-linear surface homeomorphic to W.

Theorem 2. *If S is an r-sample of W for $r \leq 0.1$, then the good triangles form a polyhedron homeomorphic to W.*

Our next two theorems state the theoretical guarantees for the three-dimensional (raw) crust.

Theorem 3. *If S is an r-sample for $r \leq 0.1$, then the crust includes all the good triangles.*

Theorem 4. *If S is an r-sample for $r \leq 0.06$, then the crust lies within a fattened surface formed by placing a ball of radius $5r\,LFS(q)$ around each point $q \in W$.*

Step 5 adds another guarantee: convergence in surface normals. The raw crust sometimes includes small skinny triangles with normals that deviate significantly from the surface normals. For example, the insides of the sausages shown on the left in Figure 15 have a sort of washboard texture. Convergence in surface normal ensures that the area of the trimmed θ-crust converges to that of the surface.

Theorem 5. *Assume S is an r-sample and set $\theta = 4r$. Let T be a triangle of the θ-crust and let t be a point on T. The angle between the normal to T and the normal to W at the point $p \in W$ closest to t measures $O(\sqrt{r})$ radians.*

5. Proofs

In this section we give the proofs of the theoretical guarantees. We first prove some basic geometric consequences of our definitions and hypotheses. The proofs of Theorems 2 and 3 will then be straightforward, albeit somewhat tedious. The proofs of Theorems 4 and 5 will require some additional ideas and constructions.

We start by defining some terminology. At each point $p \in W$, there are two maximal tangent *medial balls* centered at points of the medial axis. The vectors from p to the centers of its medial balls are normal to W, and W does not intersect the interiors of the medial balls. Since $LFS(p)$ is at most the radius of the smaller medial ball, W is also confined between the two tangent balls of radius $LFS(p)$. We call these the *tangent balls at p*; we shall use the tangent balls to bound the curvature of W in terms of $LFS(p)$. We call a maximal empty ball centered at a Voronoi vertex a *Voronoi ball*, and the Voronoi ball centered at a pole a *polar ball*.

Next we set some notation. We use $d(p, q)$ to denote the Euclidean distance from p to q. We use $|\angle psq|$ to denote the measure of angle psq in radians.

Our first lemma is a Lipschitz condition for the $LFS(p)$ function.

Lemma 1. *For any two points p and q in $\mathbb{R}^3$, $|LFS(p) - LFS(q)| \leq d(p, q)$.*

Proof. $LFS(p) \geq LFS(q) - d(p, q)$, since the ball of radius $LFS(q)$ around q contains the ball of radius $LFS(q) - d(p, q)$ around p and contains no point of the medial axis. Similarly, $LFS(q) \geq LFS(p) - d(p, q)$. □

Our second lemma shows that the line segment between two nearby points on the surface must be nearly parallel to the surface.

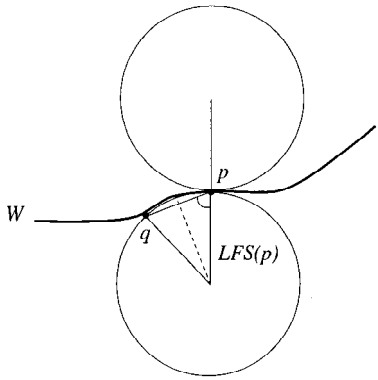

Fig. 3. The angle between pq and the normal to W at p must be nearly $\pi/2$.

Lemma 2. *For any two points p and q on W with $d(p,q) \leq \rho\, LFS(p)$, the smaller angle between the line segment pq and the surface normal at p is at least $\pi/2 - \arcsin(\rho/2)$.*

Proof. Surface W, and hence point q, lies between the tangent balls of radius $LFS(p)$ at p. The angle in question, marked in Fig. 3, is minimized when q lies right on one of the tangent balls as in the figure. This position gives the angle claimed in the lemma. □

Our third lemma is a sort of Lipschitz condition for the direction of surface normals, which can be regarded as a function from W to the two-dimensional sphere.

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

Proof. We parameterize the line segment pq by length. Let $p(t)$ denote the point on pq with parameter value t and let $f(t)$ denote the nearest point to $p(t)$ on the surface W. In other words, $f(t)$ is the point at which an expanding sphere centered at $p(t)$ first touches W. Point $f(t)$ is unique, because otherwise $p(t)$ would be a point of the medial axis, contradicting $d(p,q) \leq \rho\, LFS(p)$.

Let $n(t)$ denote the unit normal to W at $f(t)$, and let $|n'(t)|$ denote the magnitude of the derivative with respect to t, that is, the rate at which the normal turns as t grows. The change in normal between p and q is at most $\int_{pq} |n'(t)|\, dt$, which is at most $d(p,q) \cdot \max_t |n'(t)|$.

The surface W passes between the tangent balls of radius $LFS(f(t))$ at $f(t)$, so the greater of the two principal curvatures at $f(t)$ is no more than the curvature of these tangent balls. The rate at which the normal changes with $f(t)$ is at most the greater principal curvature, and hence $|n'(t)|$ is at most the rate at which the normal turns (as a function of t) on one of these tangent balls. Referring to Fig. 4, we see that

$$dt \geq (LFS(f(t)) - d(f(t), p(t))) \cdot \sin d\theta.$$

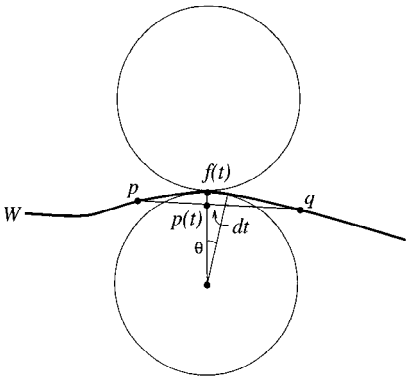

Fig. 4. Bounding $|n'(t)|$ in terms of the radius $LFS(f(t))$ and $d(f(t), p(t))$.

Now $\sin d\theta$ approaches $d\theta$ as θ goes to zero, so

$$|n'(t)| = d\theta/dt \leq 1/(LFS(f(t)) - d(f(t), p(t))).$$

We have

$$d(f(t), p(t)) \leq d(p(t), p) \leq \rho\, LFS(p)$$

and

$$d(f(t), p) \leq d(f(t), p(t)) + d(p(t), p) \leq 2\rho\, LFS(p),$$

so, by Lemma 1, $LFS(f(t)) \geq (1 - 2\rho)\, LFS(p)$. Altogether we obtain $\max_t |n'(t)| \leq 1/((1 - 3\rho)\, LFS(p))$, which yields the lemma. □

The next lemma shows that the cells of the Voronoi diagram of S are long (part a) and skinny (part b). We let $Vor(s)$ denote the closure of the Voronoi cell of s, that is, all points at least as close to s as to any other sample point. We ignore the degenerate case that $Vor(s)$ is unbounded on both sides of W.

Lemma 4. *Let s be a sample point from an r-sample S.*

(a) *On either side of W at s, some point of $Vor(s)$ has distance at least $LFS(s)$ from s.*
(b) *The intersection of $Vor(s)$ and W is contained in a ball of radius $(r/(1-r))\, LFS(s)$ about s.*

Proof. On either side of W at s, the center c of the tangent ball of radius $LFS(s)$ lies within $Vor(s)$, and hence (a) holds. For part (b), let $p \in Vor(s) \cap W$. Since s is the closest sample point to p, $d(p, s) \leq r\, LFS(p) \leq r(LFS(s) + d(p, s))$ by Lemma 1. So $d(p, s) \leq (r/(1 - r))\, LFS(s)$. □

The next lemma shows that these long skinny Voronoi cells are nearly perpendicular to the surface.

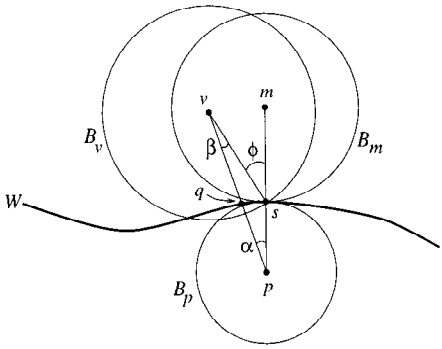

Fig. 5. The vector from s to a distant point in $Vor(s)$, such as a pole, must be nearly normal to the surface.

Lemma 5. *Let s be a sample point from an r-sample S. Let v be any point in $Vor(s)$ such that $d(v, s) \geq \nu\, LFS(s)$ for $\nu > 0$. The angle at s between the vector to v and the normal to the surface* (*oriented in the same direction*) *is at most* $\arcsin(r/\nu(1-r)) + \arcsin(r/(1-r))$.

Proof. Let B_v be the Voronoi ball centered on v. Let B_m be the medial ball touching s on the same side of the surface W, and let m be its center. Let φ be the angle between the segments sv and sm, that is, the angle referred to in the lemma. Let B_p be the ball of radius $LFS(s)$, tangent to W at s, but lying on the opposite side of W from B_m; let p be the center of B_p. The surface W passes between B_m and B_p at s, and does not intersect the interior of either of them, as shown in Fig. 5.

Since p and v lie on opposite sides of W, line segment pv must intersect W at least once. Let q be the intersection point closest to p. No sample point can lie in either B_p or B_v, so the nearest sample point to q must be s. Since B_p has radius $LFS(s)$, $d(q, s) \geq \sin(\alpha)\, LFS(s)$, where α is the angle $\angle spq$. We are interested in angle $\angle vsm$, which is $\varphi = \alpha + \beta$. Since B_v has radius at least $\nu\, LFS(s)$, $d(q, s) \geq \nu \sin(\beta)\, LFS(s)$, where β is the angle $\angle svq$. Since S is an r-sample, $d(q, s)$ must be less than $(r/(1-r))\, LFS(s)$. Combining the inequalities, we obtain $\alpha \leq \arcsin(r/(1-r))$ and $\beta \leq \arcsin(r/\nu(1-r))$, which together give the bound on φ. □

Together Lemmas 4(a) and 5 show that the vector from a sample point to its first pole p^+ is a good approximation to the surface normal. This observation may have wider applicability; for example, the Voronoi diagram and the poles could be used to obtain provably reliable estimates of tangent planes in the algorithm of Hoppe et al.

Our next lemma shows that Step 2 of the crust algorithm does indeed correctly identify the second pole p^-. Recall that p^- is defined to be the farthest Voronoi vertex from s on the opposite side of the unknown surface from p^+.

Lemma 6. *Let s be a sample point from an r-sample S with $r \leq \frac{1}{4}$. The second pole p^- of s is the farthest Voronoi vertex v of s such that the vector sv has negative dot product with sp^+.*

Proof. By Lemma 4(a), $d(s, p^-) \geq LFS(s)$, so by Lemma 5 the angle between sp^+ and sp^- is at least $\pi - 4\arcsin(r/(1-r)) \geq 1.78 > \pi/2$, hence $sp^- \cdot sp^+ < 0$. Lemma 5 also shows that for any Voronoi vertex v on the same side of W as p^+, with $d(s, v) \geq LFS(s)$, the angle between sv and sp^+ is at most $4\arcsin(r/(1-r)) \leq \pi/4$. Hence any v farther from s than p^- must have $sv \cdot sp^+ > 0$. □

Our next lemma bounds the angle between the normal to a good triangle and the surface normals at its vertices.

Lemma 7. *Let T be a good triangle and let s be a vertex of T with angle at least $\pi/3$, and choose $r < \frac{1}{7}$.*

(a) *The angle between the normal to T and the normal to W at s is at most* $\arcsin(\sqrt{3}r/(1-r))$.
(b) *The angle between the normal to T and the normal to W at any other vertex of T is at most* $2r/(1-7r) + \arcsin(\sqrt{3}r/(1-r))$.

Proof. For part (a), let C be the circumcircle of T and let ρ_C be its radius. Consider the tangent balls of radius $LFS(s)$ tangent to W at s on either side of W. These balls intersect the plane of T in twin disks of common radius ρ_B, tangent at point s, as shown in Fig. 6. Our first aim is to bound ρ_B in terms of ρ_C.

Since the balls of radius $LFS(s)$ are empty of sample points, the twin disks cannot contain vertices of T. In order to maximize ρ_B relative to ρ_C, we assume that the twin disks pass through the vertices of T and that the angle at s measures exactly $\pi/3$. Now it is not hard to show that ρ_B is maximized exactly when T is equilateral: if we move s away from the midpoint of the arc covered by the twin disks, keeping the twin disks passing through the vertices of T, the radius ρ_B decreases, until s reaches one of the other vertices of T and $\rho_B = \rho_C$. Since the worst-case configuration is equilateral T, we can conclude that $\rho_B \leq \sqrt{3}\rho_C$.

We can bound these radii in terms of $LFS(s)$. Let u denote the restricted Voronoi diagram vertex dual to T. Since u lies on the line through the center of C normal to the plane of C, $\rho_C \leq d(u, s)$. By Lemma 4(b), $d(u, s) \leq (r/(1-r))\, LFS(s)$, so altogether $\rho_B \leq \sqrt{3}r\, LFS(s)/(1-r)$.

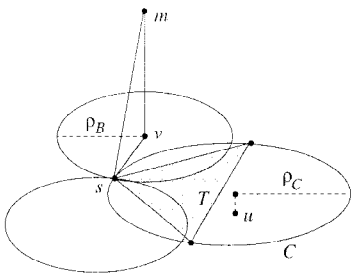

Fig. 6. Bounding the angle between the normal to the triangle and the normal to the surface at s.

Now to find the angle between the normal to T and the normal to W at s, we consider one of the tangent balls B at s. Let m denote the center of B and let v denote the center of the twin disk of radius ρ_B that is the intersection of B with the plane of T, as shown in Fig. 6. The segment sm is normal to W at s and the segment mv is normal to T, so the angle we would like to bound is $\angle smv$. The triangle smv is right, with hypotenuse of length $LFS(s)$ and leg opposite $\angle smv$ of length $\rho_B \leq \sqrt{3}r\, LFS(s)/(1-r)$. Hence $|\angle smv| \leq \arcsin(\sqrt{3}r/(1-r))$.

For part (b), let s' be one of the other vertices of T. Since T is a good triangle, s and s' are neighbors in the restricted Voronoi diagram. Let p be a point on the boundary of both restricted Voronoi diagram cells. Then

$$d(p,s) \leq r\, LFS(p) \leq \frac{r}{1-r} \min\left\{LFS(s), LFS(s')\right\}.$$

So $d(s,s') \leq (2r/(1-r)) \min\{LFS(s), LFS(s')\}$. Applying Lemma 3 with $\rho = 2r/(1-r)$, the angle between the normals to W at s and s' is at most $2r/(1-7r)$ for $r < \frac{1}{7}$. □

We need one more lemma for the proof of Theorem 2. This lemma is a topological result concerning the medial axis that should be independently useful.

Lemma 8. *If a ball B intersects surface W in more than one connected component, then B contains a point of the medial axis of W.*

Proof. Assume $B \cap W$ has more than one connected component. Let c be the center of B and let p be the nearest point on W to c. If p is not unique, then c is a point of the medial axis and we are done. Let q be the nearest point to c in a connected component of $B \cap W$ that does not contain p. Imagine a point c' moving from c toward q along segment cq. Throughout this journey, c' is closer to q than to any point outside B, so the closest point on W to c' must be some point of $B \cap W$. At the beginning of the journey, the closest point to c' is p and at the end it is q, so at some critical c' the closest point must change connected components. Such a c' is a point of the medial axis. □

We now give the proof of Theorem 2: the good triangles form a polyhedron homeomorphic to W. The proof relies on the lemmas above along with a result of Edelsbrunner and Shah [15].

Proof of Theorem 2. The theorem of Edelsbrunner and Shah tells us that it suffices to show that S has the following *closed-ball property*: the closure of each k-dimensional face, $1 \leq k \leq 3$, of the Voronoi diagram of S intersects W in either the empty set or in a closed $(k-1)$-dimensional topological ball.

Let s be a sample point and let $Vor(s)$ be its Voronoi cell. Let the direction of the normal to W at s be vertical.

We begin by showing that in the vicinity of s, the surface W is nearly horizontal; this fact will be useful in establishing the closed-ball property for each value of k. Lemma 4(b) shows that $Vor(s) \cap W$ is small, fitting inside a ball B around s of radius $(r/(1-r))\, LFS(s)$. Since such a small ball cannot intersect the medial axis, Lemma 8 implies that $W \cap B$ has a single connected component. Lemma 3 with $\nu = r/(1-r)$

shows that $W \cap B$ is nearly horizontal, more precisely, the normal to $W \cap B$ is nowhere farther than $r/(1-4r) \leq \frac{1}{6}$ radians from vertical.

Now we consider an edge e of $Vor(s)$, that is, the case $k = 1$. If e has nonempty intersection with W, then e is normal to the good triangle T dual to its intersection point. By Lemma 7(b), e must be within $2r/(1-7r) + \arcsin(\sqrt{3}r/(1-r))$ radians from the normal to W at s. For $r \leq 0.1$, this expression is less than 0.9, so e is within 0.9 radians from vertical, and consequently can intersect W only once within ball B.

Next consider a face f of $Vor(s)$, that is, the case $k = 2$. Face f is contained in plane H, the perpendicular bisector of s and another sample point s', where ss' is an edge of a good triangle T. Plane H must contain a vector parallel to the normal of T, so again Lemma 7(b) shows that the angle between h and the surface normal at s, and hence between f and the surface normal at s is at most 0.9 radians. So we call f *nearly vertical*.

Each connected component of $f \cap W$ is an arc of a curve, with endpoints on edges of f that are dual to good triangles and hence nearly vertical. Assume for the sake of contradiction that there are at least two such connected components, and consider any line segment connecting a point on one component with a point on another. Since $Vor(s) \cap W$ is small, Lemma 2 implies that each of these line segments is *nearly horizontal*, specifically within 0.2 radians of horizontal for $r \leq 0.1$. Hence we can sort the arcs of $f \cap W$ left to right across f, as shown in Fig. 7(a). Let pq be a line segment connecting the right endpoint of one arc with the left endpoint of the next arc. Segment pq is nearly horizontal so it must leave f as it crosses the nearly vertical edge of $Vor(s)$ at p and reenter f at q, a contradiction to the fact that f is convex.

Finally consider $Vor(s)$ itself, the case $k = 3$. Let C be $Vor(s) \cap W$. Now C cannot have a handle because W is nearly horizontal everywhere within ball B. We assume, however, that C is not a topological disk, again aiming for a contradiction. If C has no handles and is not a topological disk, then either it has more than one connected component or it is a topological disk with holes.

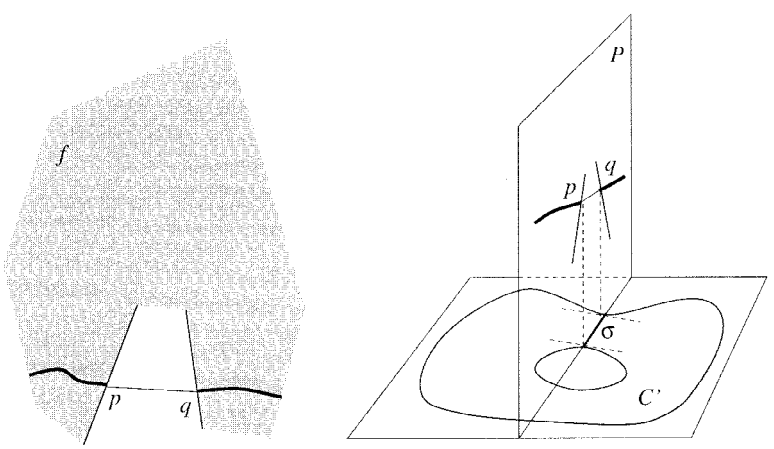

Fig. 7. Two impossible situations: (a) Face f of $Vor(s)$ intersects W in two arcs. (b) $W \cap Vor(s)$ is a disk with a hole.

Consider the projection C' of C onto a horizontal plane. Since each pair of points in C are connected by a nearly horizontal segment, this projection is one-to-one, and C' is a planar shape homeomorphic to C. If C' has more than one connected component, let σ be the shortest segment connecting two different components, and let q_1', q_2' be the endpoints of σ. Otherwise, C' has a hole. Select some point x in a hole, let σ' be the line segment connecting x to the nearest point q_1' in C'. The line supporting σ' intersects C' again on the other side of x in a point q_2'. We let σ be the segment q_1', q_2'.

Segment σ is perpendicular to the boundary of C' at q_1'. Let P be the vertical plane through σ, and let q_1, q_2 be the points in $P \cap C$ whose vertical projections are q_1', q_2', respectively. The point q_1 is either a vertex of C or an edge point, and thus lies in either an edge or a facet of $Vor(s)$, respectively.

Consider the case in which q_1 lies in a facet f of $Vor(s)$. Facet f is nearly vertical, and since P is perpendicular to C' at q_1', it is nearly perpendicular to C at q_1, and f intersects P in a nearly vertical line at q_1. In the other case, q_1 is contained in a nearly vertical edge e of $Vor(s)$. We consider the plane Q containing both both e and the horizontal line perpendicular to q_1, q_2 at q_1. Q meets P in a nearly vertical line. Note that we had to choose P carefully, since two nearly vertical planes which are *not* nearly perpendicular might meet in a nearly horizontal line.

In either case, examining the situation in P, we find that in the neighborhood of q_1, the interior of q_1, q_2 is separated from the interior of $Vor(s) \cap P$ by a nearly vertical line. However, both q_1 and q_2 belong to $Vor(s) \cap P$, contradicting the the fact that $Vor(s)$ is convex. □

Next we give a proof of Theorem 3: the raw crust contains all the good triangles. The intuition behind this proof is that restricted Voronoi cells are small and poles are far away, so that the ball centered at a vertex u of the restricted Voronoi diagram, passing through the three sample points whose cells meet at u, must be empty of poles.

Proof of Theorem 3. Let T be a triangle dual to a vertex u of the restricted Voronoi diagram. Consider the ball B_u centered on u with boundary passing through the vertices of T. Since T is a Delaunay triangle, B_u contains no point of S in its interior. Since S is an r-sample of W for $r < 1$, the radius of B_u is less than $r\,LFS(u)$. By the definition of LFS, even the larger ball B_u' of radius $LFS(u)$ centered on u cannot contain a point of the medial axis.

Now assume that B_u contains a pole v of a sample point s. We will show that under this assumption, first, that B_v must contain a point of the medial axis, and, second, that the polar ball B_v must be contained in B_u', thereby giving a contradiction. In particular, B_v must contain the center m of the medial ball B_m at s that is on the same side of W as v. Notice that m necessarily lies in $Vor(s)$ and ball B_m has radius at least $LFS(s)$, while the radius of B_v is at least that of B_m (by Lemma 4(a)). By Lemma 5, $|\angle msv| \leq 2\arcsin = (r/(1-r))$, which is less than 0.23 for $r \leq 0.1$. A calculation shows that B_v must contain the medial axis point m.

Since v lies in B_u, the radius of B_v is no greater than the distance from v to the nearest vertex of T, which is at most $2r\,LFS(u)$ since S is an r-sample. Since $d(u, v) \leq r\,LFS(u)$, ball B_v lies entirely within B_u' since $3r\,LFS(u) \leq LFS(u)$. □

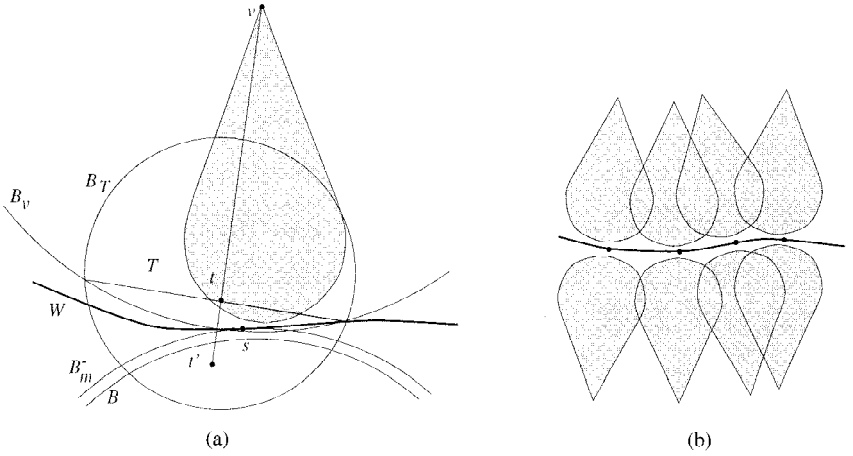

Fig. 8. (a) The Delaunay ball B_T of a triangle intersecting the spindle must contain a big patch of surface W. (b) Spindles of sample points fuse so that all triangles must lie close to W.

We now move on to the proof of Theorem 4, which says that all crust triangles lie within a fattened surface surrounding W. Let s be a sample point and let v be a pole of s. We shall define a forbidden region inside polar ball B_v, which cannot be penetrated by large crust triangles.

Let B_m^+ be the tangent ball of radius $LFS(s)$ at s on the same side of W as v, and let B_m^- be the tangent ball on the other side, with W passing between them. Let B be the ball concentric with B_m^- with radius $(1-r)\,LFS(s)$, as shown in Fig. 8(a). Notice that Lemma 4(a) shows that the radius of B_v is at least that of B.

Definition 5. The *reflection* of a point t through B_v is the point t' along ray vt such that line segment tt' is divided into equal halves by the boundary of B_v. The *spindle* of s is $\{\, t \in B_v \mid \text{segment } tt' \text{ intersects } B \,\}$, that is, all points in B_v whose reflection lies in or beyond B.

The spindle is shaded in Fig. 8(a). Our plan is to show that large crust triangles are confined between the union of spindles on each side of W as shown in Fig. 8(b). (Small crust triangles lie within the fattened surface simply due to their size.) We start by proving two lemmas about spindles: they are indeed forbidden regions, and they have relatively flat bottoms, meaning that their width does not shrink with shrinking r.

Lemma 9. *No crust triangle T whose Delaunay ball B_T has radius greater than $5r\,LFS(s)$ can penetrate the spindle of s.*

Proof. Assume t is a point inside B_v on a crust triangle T with Delaunay ball B_T. We first assert that B_T contains the reflection point t'. Let H be the plane containing the intersection of the boundaries of B_v and B_T. Since the vertices of T lie on B_T outside B_v, T must be contained in the closed half-space bounded by H not containing v. It

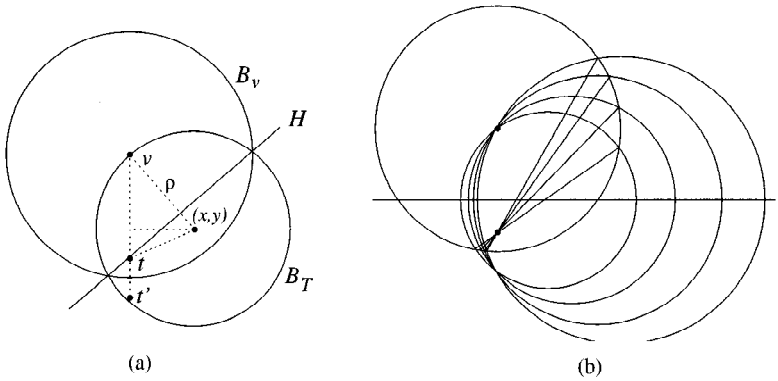

Fig. 9. (a) B_T must contain reflection point t'. (b) The family of possible B_T circles.

suffices to prove this assertion (that B_T contains t') for the case in which t lies right on H, since for any t in the interior of the half-space, segment tt' is a subsegment of the segment connecting a point on H and its reflection.

We may also assume that ball B_T passes through v, since if we replace B_T with the ball that touches v and has the same intersection with H, the part of B_T outside B_v shrinks (making things harder for our lemma).

Now consider any plane containing line vt. Balls B_v and B_T intersect this plane in circles and plane H intersects in a line containing the mutual chord of these circles. See Fig. 9(a).

Assume without loss of generality that the cross section of B_v is the unit circle with center $v = (0, 1)$. Let $t = (0, y_t)$. Denote the center and radius of B_T's cross section by (x, y) and ρ. Since t lies along the mutual chord, it has equal *power distance* to $(0, 1)$ and (x, y):

$$(1 - y_t)^2 - 1 = x^2 + (y - y_t)^2 - \rho^2.$$

Substituting $(1 - y)^2$ for $\rho^2 - x^2$, we obtain

$$y_t{}^2 - 2y_t = (y - y_t)^2 - (1 - y)^2,$$

which simplifies to $y = (1 - 2y_t)/(2 - 2y_t)$. Thus the centers of all possible B_T circles lie on the same horizontal line, as shown in Fig. 9(b).

Any B_T passes through the reflection of $(0, 1)$ across the horizontal line, the point $(0, (1 - 2y_t)/(1 - y_t) - 1)$. For any value of $y_t < 1$, $(1 - 2y_t)/(1 - y_t) - 1 < -y_t$, so B_T contains $t' = (0, -y_t)$.

Thus if the original point t lies within the spindle of s, then B_T must intersect B, the ball concentric with B_m^-. Aiming for a contradiction, assume that t does indeed lie within the spindle of s. Then B_T penetrates each of B_v and B_m^- deeply, at least $r\,LFS(s)$ into each of these balls. Consider the disk D_m bounded by the circle that is the intersection of the boundaries of B_T and B_m^-, as shown in Fig. 10. Since the radius of B_T is at least $5r\,LFS(s)$, the radius of B_m^- at least $LFS(s) \geq 15r\,LFS(s)$, and B_T cuts at least $r\,LFS(s)$ into B_m^-, we can calculate that D_m has radius at least $2.5r\,LFS(s)$. There is an analogous

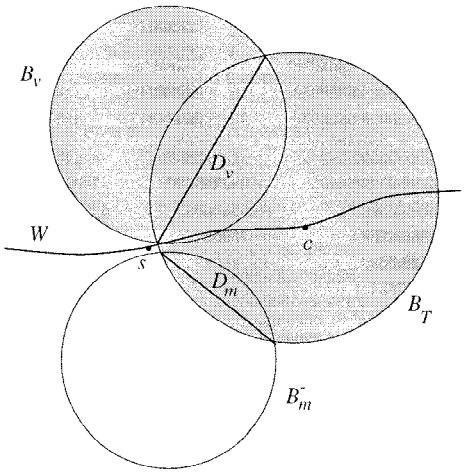

Fig. 10. If B_T does penetrate the spindle of s, then there must be a point c that has no sample points near it.

disk D_v, bounded by the intersection of the boundaries of B_v and B_T, with radius at least $2.5r\ LFS(s)$.

We now assert that there exists a point $c \in W \cap (B_T \cup B_v)$, with $d(c, s) \leq \sqrt{2}\, LFS(s)$, such that the ball of radius $2.5r\ LFS(s)$ around c contains no sample points. Surface W is confined between B_m^- and B_m^+, and hence must cross $B_T \cup B_v$. So some point c of W inside $B_T \cup B_v$ must be at least distance $2.5r\ LFS(s)$ from the boundary of $B_T \cup B_v$. It remains to determine how far c can lie from s; note that B_T need not pass through s. B_T must, however, intersect both the spindle of s and some point on the northern hemisphere of B (taking s as the north pole), so W must cross $B_T \cup B_v$ within $\sqrt{2}\, LFS(s)$ of s (since the radius of B is less than $LFS(s)$).

Now since $d(c, s) \leq LFS(s)$, $LFS(c) \leq (1 + \sqrt{2})\, LFS(s)$. However, c is at least $2.5r\ LFS(s)$ from the nearest sample, so we have obtained a contradiction to W being r-sampled. □

The next lemma shows that spindles have flat bottoms. Recall that a spindle is induced in Voronoi ball B_v by the ball B of radius $(1 - r)\, LFS(s)$ concentric with the tangent ball B_m^- on the opposite side of W. In this lemma we assume that B and B_v have equal radius, although in fact B_v is always at least a little larger than B. It is not hard to confirm that this assumption gives an upper bound: a larger B_v just gives a larger, flatter spindle. Referring to Fig. 11, recall that v is the center of B_v, m is the center of B, and define o to be the point at which segment vm (the axis of symmetry of the spindle) intersects the boundary of B_v.

Lemma 10. *Assume that B and B_v are unit balls, and that the distance between them is at most $\delta \leq 0.06$. Let t be a point outside B and outside the spindle induced by B in B_v. Let p be the closest point on B to t. If $|\angle omp| \leq 0.20$, then $d(t, p) \leq \delta + |\angle omp|$.*

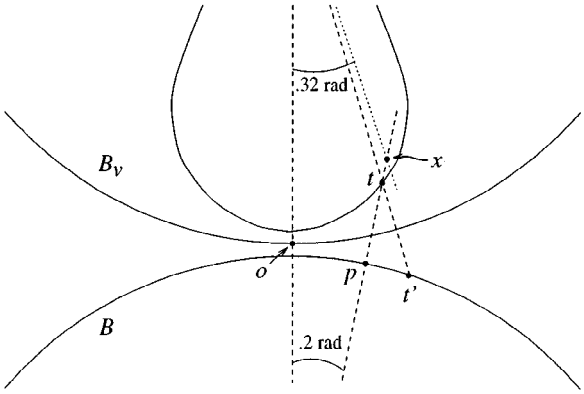

Fig. 11. The spindle curves gradually, so t must be close to B.

Proof. Assume v has coordinates $(0, 1)$. The worst case for the lemma occurs when δ assumes its maximum value, as larger δ means a higher and narrower spindle, thereby maximizing $d(t, p)$ relative to $\delta + |\angle omp|$. So assume m has coordinates $(0, -1.06)$.

Draw the 0.20-radian ray with origin m and the 0.32-radian ray with origin v as shown in Fig. 11. The rays intersect at a point x with coordinates about $(0.259, 0.218)$. By computing the distances to the boundaries of B_v and B along ray vx, we can confirm that x lies inside the spindle. Thus the boundary of the spindle lies below x on the 0.20-radian ray with origin m. Assume t and p are at the extremal positions allowed by the lemma, so that t is on the boundary of the spindle and $|\angle omp| = 0.20$. The distance from x to m is less than 1.252, so $d(t, p) - \delta \leq 0.192 \leq |\angle omp|$. Since $d(t, p)$ decreases with $|\angle omp|$, this inequality also applies to points t and p such that $|\angle omp| < 0.20$ as well. □

We are now in a position to finish the proof of the theorem: all crust triangles lie within the fattened surface formed by placing a ball of radius $5r\,LFS(q)$ around each point $q \in W$.

Proof of Theorem 4. Let B_T be the Delaunay ball of the crust triangle containing point t. Let s be the sample point nearest t. If B_T has radius less than $5r\,LFS(s)$, then there is nothing to prove, since s itself could be the q of the theorem.

So assume B_T has radius at least $5r\,LFS(s)$. Let B_v, B_m^-, and B be respectively the polar ball of s, the tangent ball of radius $LFS(s)$ on the opposite side of W, and the concentric ball with radius reduced by $r\,LFS(s)$ as in Fig. 12. Let o and o' be the points of lune $B_m^- \cap B_v$ closest to the centers of B_m^- and B_v, respectively. Points v and m belong to $Vor(s)$, and hence, so does o'. So by Lemma 4(b), $d(s, o') \leq r\,LFS(s)/(1 - r)$. Since B_v has radius at least that of B_m^-, $d(s, o) \leq d(s, o')$.

Let p and p' be the closest points to t on B and B_m^-, respectively, and let q be the point of W on line pt closest to t. Hence $d(t, q) \leq d(p, t)$. By an argument analogous to that used for o', $d(s, p') \leq r\,LFS(s)/(1 - r)$, and so by the triangle inequality,

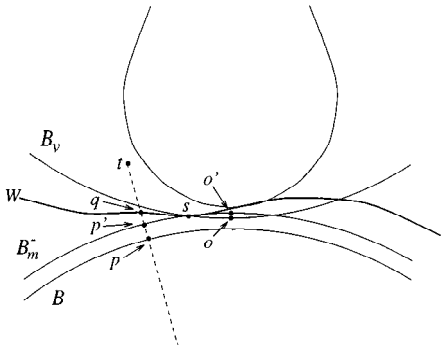

Fig. 12. Crust point t must be near surface point q.

$d(o, p') \leq 2r\, LFS(s)/(1-r)$. So $\angle omp' \leq \arcsin(2r/(1-r))$, which, for $r \leq 0.06$, is less than 0.20 radians. The set-up satisfies the hypotheses of Lemma 10, only with radii scaled by $(1-r)\, LFS(s)$.

By Lemma 9, t must lie between the spindle and B_m. Applying Lemma 10,

$$d(t, p) \leq r\, LFS(s) + |\angle omp|(1-r)\, LFS(s).$$

We now use the fact that $|\angle omp| \leq \arcsin(2r/(1-r)) \leq 3r$, to obtain

$$d(t, p) \leq r\, LFS(s) + 3r(1-r)\, LFS(s) \leq 4r\, LFS(s).$$

Finally, since again $q \in Vor(s)$, we have $d(s, q) \leq r\, LFS(s)/(1-r)$, so, by Lemma 1, $LFS(q) \geq (1-2r)\, LFS(s)/(1-r)$, and hence $5r\, LFS(q) \geq d(t, p) \geq d(t, q)$. □

Theorem 4 establishes that crust triangles are close to the surface, but not necessarily flat on the surface. Step 5 of the algorithm removes triangles whose surface normals differ too much from the vectors to the poles at their vertices, ensuring that the normals of the resulting θ-crust triangles are close to surface normals at the sample points. To establish that the θ-crust normals converge everywhere to the surface normals, we need to extend this guarantee to the interiors of all θ-crust triangles. This extension is immediate for small triangles, but takes some work for large ones. Let T be a triangle of the θ-crust with $\theta = 4r$, let t be a point on T, and let p be the closest point to t on W. Theorem 5 states that the angle between the normal to T and the normal to W at p measures $O(\sqrt{r})$ radians.

The proof shows that a large crust triangle (of size $O(\sqrt{r}LFS(s))$) must have a very large circumsphere (of size $O(LFS(s)/\sqrt{r})$) so that any nearby patch of surface cannot twist away without penetrating the circumsphere deeply, a contradiction since the circumsphere must be empty of samples.

Proof of Theorem 5. First, we establish the easier claim that at each sample point s, the normals to incident θ-crust triangles do not deviate by more than $\psi = O(r)$ radians from the normal to W. This statement follows from the fact that Step 5 of the algorithm

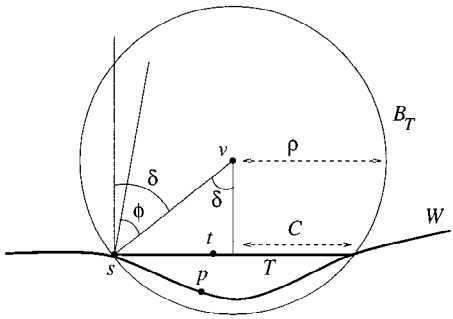

Fig. 13. Repeated use of Lemma 5 shows that if triangle T is large B_T must be enormous.

removes each triangle around s whose normal forms an angle larger than $4r$ with the vector to the pole. By Lemma 5, the pole vector in turn deviates from the normal to W at s by at most $2\arcsin(r/(1-r))$, so that $\psi \leq 7r$ for $r \leq 0.06$.

Now let t be any point on a θ-crust triangle T, and let p be the closest point on W to t. By Theorem 4, $d(t, p) \leq 5r\, LFS(p)$. Let s denote the closest vertex of T to t, let C denote the radius of T's circumcircle, and let ρ be the radius of T's Delaunay ball B_T. If $C \leq \sqrt{r}\, LFS(s)$, then $d(s, p)$ is $O(\sqrt{r})$, and Theorem 5 follows from Lemma 3 and the bound on ψ.

So assume C and hence ρ is at least $\sqrt{r}\, LFS(s)$. Let φ denote the angle between the normal to W at s and the vector from s to the center v of B_T. Lemma 5 with $\nu \geq \sqrt{r}$ implies that $\varphi \leq 2\sqrt{r}/(1-r)$ radians. Next let δ denote the angle between the normal to T at s and the vector from s to v, as shown in Fig. 13. Angle $\delta \leq \varphi + \psi$, where ψ, as above, is the angle between the normal to the surface at s and the normal to T. Since $\psi = O(r)$, we can conclude that $\delta \leq 3\sqrt{r}$.

Now $C = \rho \sin\delta$, so $\rho = C/\sin\delta \geq LFS(s)/3$. Thus the assumption that C is large (at least $\sqrt{r}\, LFS(s)$) shows that ρ must be very large (at least $LFS(s)/3$). However, we can now do better; we return to Lemma 5 with $\nu = \frac{1}{3}$. This time we obtain an upper bound of $O(r)$ on φ and δ, and a lower bound of $\Omega(LFS(s)/\sqrt{r})$ on ρ. (Sadly, we cannot repeat this trick to inflate ρ indefinitely, since ψ remains $O(r)$.)

Notice that since δ is $O(r)$, the plane containing T cuts a small spherical cap on B_T, one subtending a solid angle of only $O(r)$. This means that T itself is small with respect to B_T; the point $t \in T$ can be at most $O(r\rho)$ from the nearest vertex s, bounding (by Lemma 1) $LFS(t) \leq O(r\rho) + LFS(s)$, which is $O(\sqrt{r}\rho)$. Since t is within $5r\, LFS(p)$ of p, $LFS(p)$ is $O(\sqrt{r}\rho)$ as well.

Now assume that the normal to W at p deviates from the normal to T by $\Omega(\sqrt{r})$, and consider the tangent balls of radius $LFS(p)$ at p. The point p is close—within $O(r\, LFS(p))$—to the surface of B_T, while the radius of B_T is much larger—$\rho = O(LFS(p)/\sqrt{r})$—than the radius of the tangent balls at p. For some small enough value of r, the tangent balls intersect B_T in circular patches of radius $\Omega(\sqrt{r})\, LFS(p)$. As in the proof of Lemma 9, W is confined between these two balls, so there must be a similar-size patch of W inside B_T, and hence empty of sample points, which gives a contradiction to S being an r-sample. This contradiction establishes Theorem 5. □

6. Implementation and Examples

Manolis Kamvysselis, an undergraduate from Massachusetts Institute of Technology, implemented steps 1–4 of the crust algorithm during a summer at Xerox PARC. We used Clarkson's *Hull* program [10] for Delaunay triangulation, and *Geomview* [22] to visualize and print the results. We used vertices from pre-existing polyhedral models as inputs. A companion paper [2] reports on our experimental findings.

The only tricky part of the implementation was the handling of degeneracies and near degeneracies. Our test examples, many of which started from approximately gridded sample points, included numerous quadruples of points supporting slivers. Kamvysselis incorporated an explicit tolerance parameter ε; the circumcenter of quadruples within ε of cocircularity was computed by simply computing the circumcenter of a subset of three. This hack did not affect the overall algorithm, as these centers are never poles. Running time was only a little more than the time for two three-dimensional Delaunay triangulations. Notice that the Delaunay triangulation in step 3 involves at most three times the original number of vertices.

Figure 14 shows an especially advantageous example for our algorithm, a well-spaced point set on a smooth surface. Even though our algorithm is not designed for surfaces with boundary, it achieves perfect reconstruction on this example. Of course, the trimming step should not be used in reconstructing a surface with boundary.

Figure 15 shows an effect of undersampling. (We say we have *undersampled* if the sample set is not an r-sample for a sufficiently small r.) In this example, the raw crust contains all the good triangles, along with some extra triangles. The extra triangles turn separated sausages into link sausages, and as mentioned above roughen the inside

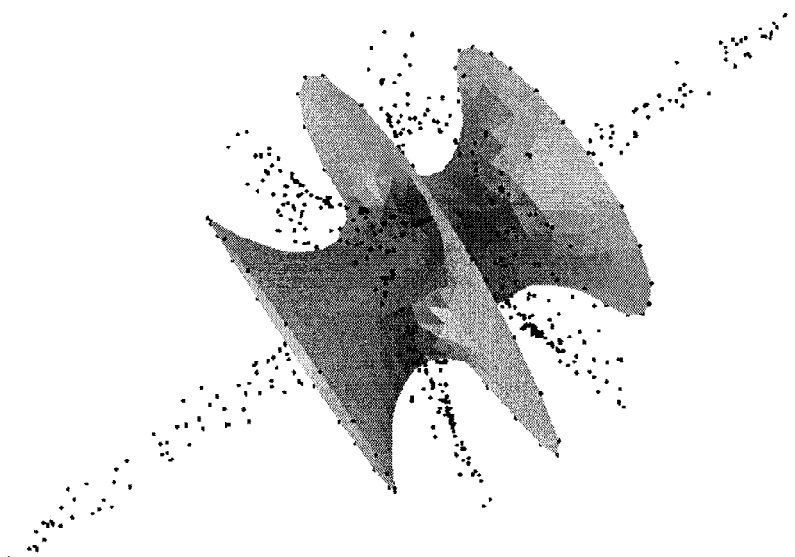

Fig. 14. A reconstructed minimal surface along with the poles of sample points. The crust contains exactly the original triangles. (Sample points courtesy of Hugues Hoppe.)

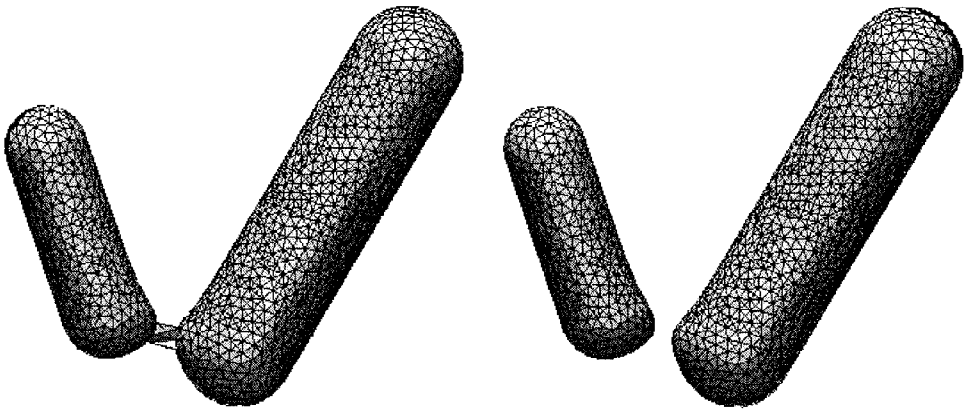

Fig. 15. The raw crust contains some extra triangles linking the sausages; this defect is corrected by step 5. Sample points courtesy of Paul Heckbert.)

surfaces of the sausages. Both of these defects are corrected by step 5, filtering by normals. Figure 16 shows another effect of undersampling: missing triangles around the chest and hooves. Some sample points are not opposed by samples on the other side of these roughly cylindrical surfaces; hence Voronoi cells extend too far and poles filter out some good triangles. An r-sample for a sufficiently small r would be very dense near the hooves, which include some rather sharp corners.

7. Conclusions and Future Work

In this paper we have given an algorithm for reconstructing an interpolating surface from sample points in three dimensions. The algorithm is simple enough to analyze, easy enough to implement, and practical enough for actual use.

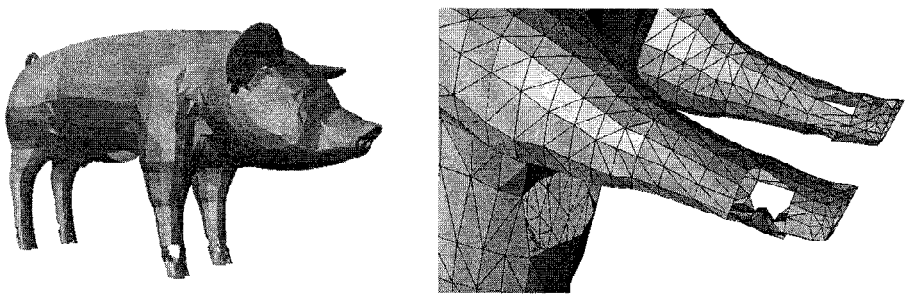

Fig. 16. (a) The pig sample set contains 3511 points. (b) A close-up of the front feet shows an effect of undersampling. (Sample points courtesy of Tim Baker.)

Our previous paper [1] gave two provably good algorithms for reconstructing curves in two dimensions, one using Voronoi filtering as in this paper, and the other using the β-skeleton. It is interesting to ask whether the β-skeleton can be generalized to the problem of surface reconstruction. (We know that the most straightforward generalization of the β-skeleton does not work.)

Another interesting question concerns the generalization of Voronoi filtering to higher dimensions. *Manifold learning* is the problem of reconstructing a smooth k-dimensional manifold embedded in $\mathbb{R}^d$. This problem arises in modeling unknown dynamical systems from experimental observations [8]. Even if Voronoi filtering can be generalized to this problem, its running time for the important case in which $k \ll d$ would not be competitive with algorithms that compute triangulations only in k-dimensional subspaces [8], rather than in $\mathbb{R}^d$.

Along with the two theoretical open questions outlined above, there are several quite practical directions for further research on our algorithms. What is the empirical maximum value of r for which our algorithm gives reliable results? We believe that the value of $r \leq 0.06$ in Theorem 4 is much smaller than necessary. Is the crust useful in simplification and compression of polyhedra? Can the crust be extended to inputs with creases and corners, such as machine parts? Can the crust be modified for the problem of reconstruction from cross sections, in which the input is more structured than scattered points?

Acknowledgments

We would like to thank Alan Cline, Bob Connelly, Tamal Dey, Herbert Edelsbrunner, David Eppstein, David Goldberg, and Manolis Kamvysselis for helpful conversations, and Ken Clarkson and The Geometry Center for making their software available.

References

1. N. Amenta, M. Bern, and D. Eppstein. The crust and the β-skeleton: combinatorial curve reconstruction. *Graphical Models and Image Processing* 60/2:2, 1998, pp. 125–135.
2. N. Amenta, M. Bern, and M. Kamvysselis. A new Voronoi-based surface reconstruction algorithm. *Proc. SIGGRAPH* '98, 1998, pp. 415–421.
3. D. Attali. r-Regular shape reconstruction from unorganized points. *Proc. 13th ACM Symp. Computational Geometry*, 1997, pp. 248–253.
4. F. Bernardini and C. Bajaj. Sampling and reconstructing manifolds using α-shapes. *Proc. 9th Canadian Conf. Computational Geometry*, 1997, pp 193–198.
5. F. Bernardini, C. Bajaj, J. Chen, and D. Schikore. Automatic Reconstruction of 3D CAD Models from Digital Scans. Technical Report CSD-97-012, Purdue University (1997).
6. F. Bernardini, C. Bajaj, J. Chen, and D. Schikore. A triangulation-based object reconstruction method. *Proc. 13th ACM Symp. Computational Geometry*, 1997, pp. 481–484.
7. J.-D. Boissonnat. Geometric structures for three-dimensional shape reconstruction. *ACM Trans. Graphics* **3** (1984), 266–286.
8. C. Bregler and S. M. Omohundro. Nonlinear manifold learning for visual speech recognition. *Proc. 5th International Conf. Computer Vision*, 1995, pp. 494–499.
9. L. P. Chew. Guaranteed-quality mesh generation for curved surfaces. *Proc. 9th ACM Symp. Computational Geometry*, 1993, pp. 274–280.

10. K. Clarkson. *Hull*: a program for convex hulls. http://cm.bell-labs.com/netlib/voronoi/hull.html.
11. B. Curless and M. Levoy. A volumetric method for building complex models from range images. *Proc. SIGGRAPH* '96, 1996, pp. 303–312.
12. H. Edelsbrunner. Surface Reconstruction by Wrapping Finite Sets in Space. Technical Report 96-001, Raindrop Geomagic, Inc., 1996.
13. H. Edelsbrunner, D. G. Kirkpatrick, and R. Seidel. On the shape of a set of points in the plane. *IEEE Trans. Inform. Theory* **29** (1983), 551–559.
14. H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graphics* **13** (1994), 43–72.
15. H. Edelsbrunner and N. Shah. Triangulating topological spaces. *Proc.* 10*th ACM Symp. Computational Geometry*, 1994, pp. 285–292.
16. L. H. de Figueiredo and J. de Miranda Gomes. Computational morphology of curves. *Visual Comput.* **11** (1995), 105–112.
17. J. Goldak, X. Yu, A. Knight, and L. Dong. Constructing discrete medial axis of 3-D objects. *Internat. J. Comput. Geom. Appl.* **1** (1991), 327–339.
18. H. Hoppe. Surface Reconstruction from Unorganized Points. Ph.D. Thesis, Computer Science and Engineering, University of Washington, 1994. http://www.research.microsoft.com/research/graphics/hoppe/thesis/thesis.html.
19. H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. *Proc. SIGGRAPH* '92, 1992, pp. 71–78.
20. H. Hoppe, T. DeRose, T. Duchamp, H. Jin, J. McDonald, and W. Stuetzle. Piecewise smooth surface reconstruction. *Proc. SIGGRAPH* '94, 1994, pp. 19–26.
21. D. G. Kirkpatrick and J. D. Radke. A framework for computational morphology. In *Computational Geometry*, G. Toussaint, ed., Elsevier, Amsterdam, 1988, pp. 217–248.
22. S. Levy, T. Munzner, and M. Phillips. Geomview. http://www.geom.umn.edu/software/download/ geomview.html.

Received August 25, 1998, *and in revised form March* 31, 1999.