Fixed Points of the Restricted Delaunay Triangulation Operator

Marc Khoury[1] **and Jonathan Richard Shewchuk**[2]

1 University of California, Berkeley, USA
2 University of California, Berkeley, USA

Abstract

The restricted Delaunay triangulation can be conceived as an operator that takes as input a k-manifold (typically smooth) embedded in $\mathbb{R}^d$ and a set of points sampled with sufficient density on that manifold, and produces as output a k-dimensional triangulation of the manifold, the input points serving as its vertices. What happens if we feed that triangulation back into the operator, replacing the original manifold, while retaining the same set of input points? If $k = 2$ and the sample points are sufficiently dense, we obtain another triangulation of the manifold. Iterating this process, we soon reach an iteration for which the input and output triangulations are the same. We call this triangulation a *fixed point* of the restricted Delaunay triangulation operator.

With this observation, and a new test for distinguishing "critical points" near the manifold from those near its medial axis, we develop a provably good surface reconstruction algorithm for $\mathbb{R}^3$ with unusually modest sampling requirements. We develop a similar algorithm for constructing a simplicial complex that models a 2-manifold embedded in a high-dimensional space $\mathbb{R}^d$, also with modest sampling requirements (especially compared to algorithms that depend on sliver exudation). The latter algorithm builds a non-manifold representation similar to the flow complex, but made solely of Delaunay simplices. The algorithm avoids the curse of dimensionality: its running time is polynomial, not exponential, in d.

1998 ACM Subject Classification F.2.2 Analysis of Algorithms and Problem Complexity: Nonnumerical Algorithms and Problems

Keywords and phrases restricted Delaunay triangulation, fixed point, manifold reconstruction, surface reconstruction, computational geometry

Digital Object Identifier 10.4230/LIPIcs.SoCG.2016.47

1 Introduction

Manifold reconstruction is the problem of discovering the structure of a k-dimensional manifold embedded in $\mathbb{R}^d$, given only a set of points sampled from the manifold. The classic application is *surface reconstruction*, the problem of discovering a surface (2-manifold) in three-dimensional space. Manifold reconstruction also has applications in higher-dimensional spaces, such as discovering relationships in data and studying algebraic varieties.

A large vein of research in surface and manifold reconstruction develops algorithms that are *provably good* [1, 2, 13, 5]: if the points sampled from a manifold are sufficiently dense, these algorithms are guaranteed to produce a geometrically accurate representation of the unknown manifold with the correct topology. Most of these algorithms are based on Delaunay triangulations, and their output is a set of Delaunay simplices that approximates the unknown manifold. In particular, they generate *restricted Delaunay triangulations*. In this paper, we study restricted Delaunay triangulations that model 2-manifolds embedded in ambient spaces of three or more dimensions. Our algorithm's guarantees hold for relatively sparse point samples, especially in four or more dimensions.

The restricted Delaunay triangulation (RDT) is a geometric structure that formally extends the Delaunay triangulation to curved manifolds. Think of the RDT as a mathematical operator that takes two inputs—a manifold and a finite set of points that lie on the manifold—and produces as output a triangulation of the manifold whose vertices are the input points. This result is conditional: if the points are not sampled densely enough, the RDT will not be a triangulation of the manifold. The RDT is (unconditionally) a subcomplex of the full-dimensional Delaunay triangulation in $\mathbb{R}^d$. It has proven itself as a mathematically powerful tool for surface meshing and surface reconstruction.

What happens if we feed an RDT of a 2-manifold back into the operator, replacing the original manifold, while retaining the same set of input points? If the sample points are sufficiently dense, and the triangulation is a sufficiently good approximation of the manifold from which they were sampled, we obtain another triangulation of the manifold. We prove that if we iterate this process, we eventually reach an iteration for which the input triangulation and output triangulation are the same. We call this triangulation a *fixed point* of the restricted Delaunay triangulation operator for the specified point set. Such *fixed-point triangulations* appear to be particularly accurate representations of the unknown manifold.

Fixed-point triangulations bifurcate into those that stay everywhere close to the unknown manifold and those that somewhere come very close to its medial axis. We call the set of simplices that can appear in the former the *manifold fixed-point complex* (MF complex), and we discuss how to compute it. The MF complex is a fixed point of the RDT operator; moreover, it is the union of all the fixed points that stay far from the medial axis.

The MF complex closely resembles the well-known *flow complex* [17], but it better approximates (in Hausdorff distance and tangents) the unknown manifold. Unlike the flow complex, the MF complex is a subcomplex of the Delaunay triangulation, so it offers greater simplicity. In three dimensions, it is easy to extract a subcomplex of the MF complex that is itself a manifold, homeomorphic to the unknown manifold. In higher dimensions, we prove that such a subcomplex exists. We believe that heuristic algorithms will have little difficulty extracting a homeomorphic manifold, but the problem seems difficult in theory.

Manifold reconstruction algorithms in high-dimensional ambient spaces face the *curse of dimensionality*: a Delaunay triangulation of n vertices in the ambient space $\mathbb{R}^d$ can have up to $\Theta(n^{\lceil d/2 \rceil})$ d-simplices. Even if the complexity is close to linear, as often happens, it is not feasible to triangulate large point clouds of dimension much higher than 10. Our algorithm's running time is polynomial in d, thanks to a method for computing low-dimensional cross-sections of high-dimensional Voronoi diagrams that we borrow from Flötotto [16].

An important ingredient for computing either the MF complex or the flow complex is a test that examines "critical points" (where Delaunay simplices intersect their Voronoi duals) and distinguishes those near the unknown manifold from those near its medial axis. We propose a test that is correct for sparser samples than the test proposed by Dey, Giesen, Ramos, and Sadri [12] in their paper on computing flow complexes. More importantly, we extend it to higher dimensions, without constructing the Delaunay triangulation in $\mathbb{R}^d$. Our new test can be used to construct flow complexes as well.

A benefit of our approach is that our analysis guarantees good reconstructions for coarser point samples than traditional Delaunay reconstruction algorithms such as the Crust [1] and Cocone [2, 13] algorithms or the Dey et al. flow complex algorithm. In $\mathbb{R}^3$, we compute a homeomorphic reconstruction from a 0.143-sample of a surface (see the definition in Section 2); by comparison, the analysis of the Cocone algorithm guarantees a homeomorphic reconstruction only for a 0.05-sample. We attribute the improved constant to our new critical point classifier, a new method for proving homeomorphisms, and the exploitation of

critical points (like the flow complex, but with Delaunay triangles). In higher dimensions, a 0.143-sample also suffices, though we do not guarantee a homeomorphic reconstruction. This coarse sampling requirement sharply distinguishes our algorithm from previous algorithms for manifold reconstruction in codimension two or higher [5, 9]. (See Section 3).

In $\mathbb{R}^3$, our algorithm produces triangles that are similar but not identical to those produced by Edelsbrunner's Wrap algorithm [14] (which is based on flows); and similar but not identical to those used by Dey et al. to help construct the flow complex. The differences are subtle but theoretically important, as they permit us to benefit from the forthcoming Lemma 5.

Unfortunately, our algorithm reconstructs only 2-manifolds. The problem is that the Delaunay 3-simplices used to triangulate a 3-manifold can have extremely unstable affine hulls and normals; they can be perpendicular to the original manifold, and their normals parallel. If we feed such a defective triangulation back into the RDT operator, it will likely produce a triangulation with nasty folds and self-intersections, not homeomorphic to Σ. Nevertheless, we suspect that RDTs will someday be adapted to reconstructing general k-manifolds. For that reason, we discuss some parts of our algorithm in a general-dimensional framework.

2 The Restricted Delaunay Triangulation

Let Σ be a k-manifold embedded in $\mathbb{R}^d$. Let $V \subset \Sigma$ (for "Vertices") be a finite set of n points sampled from Σ. Let $\operatorname{Del} V$ and $\operatorname{Vor} V$ denote, respectively, the d-dimensional Delaunay triangulation and Voronoi diagram of V. Our goal is to find a subset of the k-faces in $\operatorname{Del} V$ that forms a triangulation that geometrically approximates Σ and is topologically equivalent to Σ. For $k = 2$, this is always possible if Σ is sufficiently smooth and V is sufficiently dense.

Throughout this paper, we use pq to denote the line segment connecting p to q, and $|pq|$ to denote its length; i.e., the distance from p to q. The *codimension* of Σ is $d-k$. We presume familiarity with the duality between Delaunay triangulations and Voronoi diagrams: every simplex τ of dimension j in $\operatorname{Del} V$ has a dual face of dimension $d-j$ in $\operatorname{Vor} V$, denoted τ^*. In particular, the Voronoi cell of a site $u \in V$ is denoted $u^* = \{p \in \mathbb{R}^d : \forall w \in V, |up| \leq |wp|\}$.

The *restricted Voronoi diagram* of V with respect to Σ, denoted $\operatorname{Vor}|_\Sigma V$, is a cell complex much like $\operatorname{Vor} V$, but the Voronoi cells contain only points on the manifold Σ. Thus $\operatorname{Vor}|_\Sigma V = \{c \cap \Sigma : c \in \operatorname{Vor} V\}$, where c ranges over all the cells of all dimensions in the Voronoi diagram. In other words, we replace each Voronoi face c with its *restriction* $c|_\Sigma = c \cap \Sigma$. For each site $u \in V$, $\operatorname{Vor}|_\Sigma V$ contains the *restricted Voronoi cell* of u, namely, $u^*|_\Sigma = \{p \in \Sigma : \forall w \in V, |up| \leq |wp|\}$.

RDTs are defined *not* by restricting Delaunay simplices to a manifold, but by dualizing the restricted Voronoi diagram. The *restricted Delaunay triangulation* of V with respect to Σ, denoted $\operatorname{Del}|_\Sigma V$, is the subcomplex of $\operatorname{Del} V$ that contains every simplex in $\operatorname{Del} V$ whose Voronoi dual face intersects Σ. (I.e., $\operatorname{Del}|_\Sigma V$ contains every simplex dual to a face of $\operatorname{Vor} V$ that contributes a nonempty face to $\operatorname{Vor}|_\Sigma V$.)

We find it useful to define $\operatorname{Del}|_P V$ not just for manifolds, but for any arbitrary point set $P \subseteq \mathbb{R}^d$. Although the RDT operator is usually applied to smooth manifolds, we will apply it to piecewise linear manifolds and even individual simplices.

Let us characterize the restricted Delaunay simplices. A *circumsphere* of a j-simplex is any hypersphere in $\mathbb{R}^d$ that passes through all $j+1$ vertices of the simplex. A d-simplex has one unique circumsphere, whereas a lower-dimensional simplex has infinitely many. The *diametric sphere* of a simplex τ is the unique circumsphere with least radius. The center of τ's diametric sphere, called the *circumcenter* of τ, lies on the affine hull of τ. The *affine circumsphere* of τ is the unique $(j-1)$-sphere that passes through all $j+1$ vertices; it lies on τ's affine hull and it is a cross-section of every circumsphere of τ. The *circumradius* of τ is the radius of its affine circumsphere, which is also the radius of its diametric sphere.

A circumsphere S is *empty* if no vertex in V is inside S. (Vertices precisely on S don't count.) Simplices in an ordinary Delaunay triangulation $\operatorname{Del} V$ are characterized by the well-known *empty circumsphere condition*: τ is *Delaunay* if its vertices are in V and it has at least one empty circumsphere. For τ to be *restricted Delaunay*, it must satisfy these conditions plus one more: it has at least one empty circumsphere whose center lies on the manifold Σ.

If Σ is a k-manifold, $\operatorname{Del}|_\Sigma V$ typically contains only simplices of dimensions k and lower. If simplices of higher dimension appear, we perturb Σ infinitesimally to make them go away. For any cell complex $\mathcal{T}$, the *underlying space* $|\mathcal{T}|$ of $\mathcal{T}$ is the union of simplices in $\mathcal{T}$; i.e., $|\mathcal{T}| = \bigcup_{\tau\in\mathcal{T}} \tau$. (The operator $|\cdot|$ collapses a set of sets of points into a set of points.) Under favorable conditions, $|\operatorname{Del}|_\Sigma V|$ can be a k-manifold and a topologically and geometrically good approximation of Σ.

RDTs have become a standard tool in provably good surface reconstruction algorithms such as Crust [1] and Cocone [2, 13] and in many guaranteed-quality surface meshing algorithms [6, 10, 15]. In the mesh generation problem, Σ is known and we choose the mesh vertices $V \in \Sigma$ so that $\operatorname{Del}|_\Sigma V$ is a high-quality mesh. In the manifold reconstruction problem, V is known but Σ is not, and we try to infer Σ from V by guessing the subcomplex of $\operatorname{Del} V$ most likely to be $\operatorname{Del}|_\Sigma V$ or a good approximation thereof. Without knowing Σ, however, we cannot necessarily determine $\operatorname{Del}|_\Sigma V$, even if V is sampled extremely densely.

A major research victory in surface reconstruction was the characterization of how dense a point sample $V \subset \Sigma$ needs to be to guarantee a correct reconstruction, taking into account that the requirements vary through space. Let M be Σ's medial axis. The *local feature size* function is $\operatorname{lfs} : \Sigma \to \mathbb{R}, p \mapsto d(p, M)$, where $d(p, M)$ denotes the distance from p to M. We require that Σ is smooth enough that $\inf_{p\in\Sigma} \operatorname{lfs}(p) > 0$. A finite point set $V \subset \Sigma$ is an *ϵ-sample* of Σ if for every point $p \in \Sigma$, $d(p, V) \le \epsilon \operatorname{lfs}(p)$. That is, there is some sample $v \in V$ in the ball of radius $\epsilon \operatorname{lfs}(p)$ centered at p. The function $\operatorname{lfs}(\cdot)$ has proven itself as an appropriate, space-varying measure of the sampling density necessary so that algorithms for surface reconstruction and surface meshing can guarantee the correctness of their outputs.

For a 2-manifold $\Sigma \subset \mathbb{R}^d$ and a 0.2695-sample $V \subset \Sigma$, we can show that $|\operatorname{Del}|_\Sigma V|$ is homeomorphic to Σ. (See the full-length version of this paper.) There is a homeomorphism that doesn't move any point too far, so the Hausdorff distance between Σ and $|\operatorname{Del}|_\Sigma V|$ is small. The triangles in $\operatorname{Del}|_\Sigma V$ are approximately parallel to the surface Σ nearby.

To beat the curse of dimensionality, we borrow from Flötotto [16] and Boissonnat and Ghosh [5] the idea of the *tangential Delaunay complex*, which is simply the restricted Delaunay triangulation $\operatorname{Del}|_\Pi V$ where Π is a k-flat (i.e., a k-dimensional affine subspace). Whereas a d-dimensional Voronoi diagram $\operatorname{Vor} V$ with n vertices may have complexity as high as $\Theta(n^{\lceil d/2\rceil})$, a k-dimensional cross-section of $\operatorname{Vor} V$ has complexity only $O(n^{\lceil k/2\rceil})$ and can be computed much faster. The trick is to observe that a k-dimensional cross-section $\operatorname{Vor}|_\Pi V$ of $\operatorname{Vor} V$ is a *power diagram*. If we project each point $v \in V$ orthogonally onto Π and assign each projected point a weight of $-d(v, \Pi)^2$, where $d(v, \Pi)$ is the distance from v to Π, then the k-dimensional power diagram of the weighted, projected points is the desired cross-section of the d-dimensional Voronoi diagram. The dual of $\operatorname{Vor}|_\Pi V$ is a *weighted Delaunay triangulation* $\operatorname{Del}|_\Pi V$ and can be computed by standard algorithms for computing $(k+1)$-dimensional convex hulls such as the randomized incremental algorithm of Clarkson and Shor [11], which takes expected $O(n^{\lceil k/2\rceil} + n \log n)$ time.

For more details about Delaunay triangulations, restricted Delaunay triangulations, medial axes, the local feature size, and ϵ-samples, see Dey [13] or Cheng et al. [10].

3 Slivers and Ambiguities in Manifold Reconstruction

All Delaunay-based surface reconstruction algorithms face a problem that doesn't sound like it should be a problem: $\mathrm{Del}\, V$ typically contains many different subcomplexes that would serve as good triangulations of Σ. Even in three dimensions, four nearly-cocircular vertices on Σ may form a flat, kite-shaped tetrahedron called a *sliver*, which is usually aligned roughly parallel to the surface. Slivers create ambiguities for surface reconstruction, because the top two faces or the bottom two faces of a sliver might equally well represent the surface. The choice might not matter, but to obtain a manifold a choice must be made.

Extracting a manifold from $\mathrm{Del}\, V$ typically requires many such choices. Unfortunately, these choices are not independent. Ultra-thin slivers can stack up in interlocking layers. In higher dimensions, slivers and their higher-dimensional analogs layer along multiple axes.

Sections 4 through 6 analyze the effects of taking the output of the RDT operator and feeding it back into the RDT operator. Slivers are the reason why the RDTs produced by successive iterations may differ from each other. We will see that successive iterates tend to be better than their predecessors until they arrive at a fixed point. However, fixed points are not unique. For example, if a sliver in $\mathbb{R}^3$ contains its own circumcenter, there may be one fixed-point triangulation that chooses the sliver's bottom two faces, and another that chooses the top two. Making the point sample V very dense does nothing to prevent this.

Section 8 presents an algorithm that computes a manifold fixed-point complex (MF complex) that is the union of all fixed-point triangulations that don't get close to the medial axis. The MF complex might contain all four faces of a sliver that contains its own circumcenter (or, if $d > 3$, intersects its own Voronoi dual face). However, typically only a minority of slivers contain their own circumcenters. (The circumcenter of a sliver is very unstable: a small perturbation of a vertex can shoot the circumcenter far away.) The MF complex accepts fewer triangles than algorithms like Crust and Cocone, resolving most, but not all, of the ambiguities, and also obtaining higher quality.

In $\mathbb{R}^3$, there is a fast, simple method for resolving the remaining ambiguities and extracting a manifold, described in Section 8, but it relies on the fact that a codimension-one manifold ($d - k = 1$) partitions space into "inside" and "outside"; it finds the outermost boundary.

For codimension two or higher, Cheng, Dey, and Ramos [9] and Boissonnat and Ghosh [5] solve this problem by replacing $\mathrm{Del}\, V$ with a weighted Delaunay triangulation and using a technique called *sliver exudation* [8] to select vertex weights so that the slivers disappear and the triangulation has only one subcomplex that approximates Σ. At least in theory, sliver exudation also has the virtue that it enables the reconstruction of manifolds of dimension greater than 2, unlike our method. Unfortunately, sliver exudation is so fragile that it is only guaranteed to work with an *extremely* dense sample V—the bounds on ϵ are not derived explicitly, but these algorithms probably require a 10^{-10}-sample or finer to guarantee success. Both groups of authors admit their algorithms are not practical.

In theory, the problem of manifold extraction from an MF complex seems challenging and we wonder if it is NP-hard. (Bern and Eppstein [4] observe that given a collection of triangles in the plane, it is NP-hard to determine whether some subset of them form a triangulation of the triangles' vertex set.) Nevertheless, we suspect that manifold extraction from an MF complex is not hard in practice, as slivers usually appear as isolated singletons or in small groups, especially after the MF complex filters out most slivers. We predict that in practice, 2-manifolds can reliably be extracted from MF complexes by methods based on heuristic search; whereas sliver exudation will never be practical, as its weight perturbations degrade the quality of triangulations that may already be fragile due to sparse sampling.

4 Fixed-Point Triangulations

In the 2-manifold reconstruction problem, we do not know Σ, so we cannot compute $\mathrm{Del}|_\Sigma V$. But suppose we could somehow generate a rough approximation Λ of Σ that is piecewise linear and homeomorphic to Σ, that stays somewhat close to Σ, and that is locally approximately parallel to Σ (in ways we formalize later). We will show that the triangulation $\mathcal{T}_1 = \mathrm{Del}|_\Lambda V$ is also homeomorphic to Σ, and it is likely to be a better approximation to Σ than Λ.

As Λ is piecewise linear, we can compute $\mathrm{Del}|_\Lambda V$ in high dimensions without suffering the curse of dimensionality. We take each flat piece $f \subset \Lambda$ and compute $\mathrm{Del}|_f V$ by computing a Voronoi cross-section as described in Section 2. (Observe that $\mathrm{Del}|_f V$ is a subcomplex of $\mathrm{Del}|_{\mathrm{aff}\, f} V$ where $\mathrm{aff}\, f$ denotes f's affine hull.) We stitch these fragmentary triangulations together to yield $\mathcal{T}_1$. This method requires the computation of many low-dimensional power diagrams, but nonetheless it may be much faster than computing a high-dimensional Voronoi diagram. Moreover, it can more easily take advantage of locality: if the sample points are uniformly distributed on Σ, simple spatial data structures such as spatial hashing might speed up computing $\mathrm{Del}|_f V$ to as fast as constant time. The Voronoi cross-section computations are mutually independent, making them amenable to multicore parallelization.

Next, consider computing $\mathcal{T}_2 = \mathrm{Del}|_{|\mathcal{T}_1|} V$ the same way; now each flat piece is a triangle $\tau \in \mathcal{T}_1$. Likewise compute $\mathcal{T}_3 = \mathrm{Del}|_{|\mathcal{T}_2|} V$, $\mathcal{T}_4 = \mathrm{Del}|_{|\mathcal{T}_3|} V$, etc. We eventually obtain a triangulation $\mathcal{T}^* = \mathrm{Del}|_{|\mathcal{T}^*|} V$ identical to the last input. We call $\mathcal{T}^*$ a *fixed point* of the restricted Delaunay triangulation operator, and we denote it $\mathrm{Del}^*|_\Lambda V$.

Each simplex τ in a triangulation $\mathcal{T}_i$, $i \geq 2$, is there because of a Delaunay simplex $\sigma \in \mathcal{T}_{i-1}$ that intersects τ^*. We say that σ *generates* τ. Next we show that, with one minor exception, a simplex only generates simplices with smaller circumradii—plus it might regenerate itself. This leads us to conclude that the iterations always reach a fixed point. Moreover, each successive iteration tends to shrink the circumradii. The fixed-point triangulation $\mathrm{Del}^*|_\Lambda V$ is composed of small-radius triangles; under the right conditions, this helps us to guarantee that it is close to the true manifold Σ and has similar tangents and normals.

▶ **Lemma 1** (Monotonic Circumradii Lemma)**.** *Consider the restricted Delaunay triangulation $\mathcal{T} = \mathrm{Del}|_{|\mathcal{D}|} V$, where $\mathcal{D} \subseteq \mathrm{Del}\, V$ is a set of Delaunay simplices. Each restricted Delaunay simplex $\tau \in \mathcal{T}$ is generated by a simplex in $\mathcal{D}$ with greater circumradius, or by a simplex with the same circumradius and the same circumcenter as τ (possibly by τ itself).*

Proof. Consider a simplex $\tau \in \mathcal{T}$. By the definition of restricted Delaunay triangulation, the Voronoi face τ^* dual to τ intersects some simplex $\sigma \in \mathcal{D}$, and thus σ generates τ.

Let p be a point in $\sigma \cap \tau^*$. As τ^* is τ's dual Voronoi face, the hypersphere S with center p that passes through every vertex of τ is empty (encloses no vertex in V). Let R be the radius of S, and let r_τ be the circumradius of τ. The affine circumsphere of τ is a cross-section of S, so $r_\tau \leq R$. Moreover, if $r_\tau = R$ then τ's circumcenter is the center p of S.

Let D be the diametric sphere of σ, and let r_σ be the radius of D. If $D = S$, then $r_\sigma = R$ and p is the circumcenter of σ; whereas if D encloses S, then $r_\sigma > R$. As $p \in \sigma$ and S is empty, there is only one other possibility, illustrated in Figure 1: D intersects S, and $D \cap S$ is a $(d-2)$-sphere whose affine hull is a hyperplane Π. As S is empty, σ's vertices are restricted to one side of Π (possibly lying on Π), and therefore so is σ itself. As $p \in \sigma$, at least half of S lies on that side of Π, and that portion of S is enclosed by D. Hence $r_\sigma > R$.

Therefore, $r_\tau \leq r_\sigma$, and equality is possible only if p is the circumcenter of σ and τ. ◀

To better understand fixed-point triangulations and algorithms for computing them, we define a directed almost-acyclic graph called the *restriction DAG* G. There is one node in G for each simplex in $\mathrm{Del}\, V$. G has a directed edge (σ, τ) if σ generates τ.

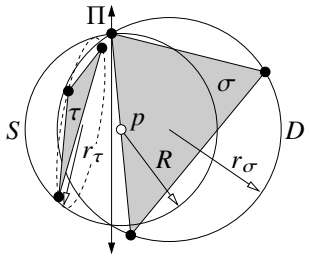

Figure 1 The Monotonic Circumradii Lemma: if σ generates τ, then $r_\tau \leq R \leq r_\sigma$.

Lemma 1 implies that G is almost acyclic. Two kinds of cycles are possible. We say that a Delaunay simplex τ is *anchored* if it intersects its own dual τ^*. An anchored simplex has a self-edge (τ, τ) in G and can generate itself. The second kind of cycle is a clique of two or more anchored simplices, representing a "degenerate" configuration in which several simplices share the same affine circumsphere, and they all intersect their shared circumcenter (which implies that the circumcenter is on the boundary of each simplex). As the simplices in such a clique are all present or all absent in a fixed-point triangulation, we contract each such clique into a single node of G (with self-edge). Then the only cycles are self-edges.

As G (after clique contraction) is a DAG (plus self-edges), we use standard tree terminology: given edges (σ, τ) and (τ, ζ), σ is a *parent* of τ and an *ancestor* of ζ, τ, and σ; ζ is a *child* of τ and a *descendant* of σ, τ, and ζ. The acyclicity of G implies that iterations of the restricted Delaunay operator always reach a fixed point, and allows us to characterize it.

▶ **Theorem 2.** *Let $\mathcal{D} \subseteq \operatorname{Del} V$ be a Delaunay subcomplex. Then the fixed-point triangulation $\operatorname{Del}^*|_{|\mathcal{D}|} V$ exists and contains every anchored simplex that has an ancestor in $\mathcal{D}$ (possibly itself), every descendant of those anchored simplices, and no other simplex.*

▶ **Corollary 3.** *For every arbitrary point set $P \subseteq \mathbb{R}^d$ and every finite point set $V \subset \mathbb{R}^d$, $\operatorname{Del}^*|_P V$ exists and can be computed with a finite number of iterations of the restricted Delaunay operator.*

The corollary is unconditional. For every P and V, we can take $\mathcal{D} = \operatorname{Del}|_P V$ and characterize the simplices in $\operatorname{Del}^*|_P V = \operatorname{Del}^*|_{|\mathcal{D}|} V$ with Theorem 2. Of course, the fact that a fixed point exists doesn't imply that it is anything more than a messy subcomplex of $\operatorname{Del} V$. If the fixed-point triangulation is to be a good manifold, P should be a good approximation of Σ and V should be an ϵ-sample for a small ϵ.

Theorem 2 serves an algorithmic purpose as well as a theoretical one. The fastest way to compute a fixed-point triangulation is not to iterate the restricted Delaunay triangulation operator; rather, we apply Theorem 2 and build part of the restriction DAG. (After we compute which simplices a particular simplex generates, there is no need to compute them again.) It suffices to generate the nodes of G that represent simplices that arise in one of the iterations. Moreover, it suffices to consider just the k-simplices. (We perturb $|\mathcal{D}|$ to prevent higher-dimensional simplices from being generated, and we gain no benefit from generating a lower-dimensional simplex that is not a face of a generated k-simplex.) If $\operatorname{Del} V$ has many slivers, we might need just a small portion of the DAG.

We do not know a useful bound on the number of iterations required to reach a fixed point (we doubt it will ever reach 8 in practice), but we present evidence in the next section that convergence is fast. One can generate a good triangulation by stopping after a few iterations, without reaching a fixed point. Moreover, if many iterations are performed, the

extra computation is likely needed only in a few locations. At any rate, the reconstruction algorithms we present in Section 8 do not actually perform these iterations; rather, they directly find triangles that can appear in a fixed-point triangulation.

We give some intuition why a fixed-point triangulation $\mathcal{T}$ partly resolves some of the ambiguities posed by slivers. In codimension one ($d-k=1$), $|\mathcal{T}|$ partitions $\mathbb{R}^d$ into an "inside" region and an "outside" region. If the circumcenter of a sliver d-simplex τ lies outside $|\mathcal{T}|$, then τ is outside also; if τ's circumcenter is inside, so is τ. (If τ contains its own circumcenter, it can be classified either way, so the fixed-point triangulation is not unique.) In codimension two or higher ($d-k \geq 2$), a fixed-point triangulation maintains an analogous relationship among the manifold k-simplices and the Voronoi $(d-k)$-cells they intersect.

5 Proximity of the Fixed-Point Triangulation to Σ

Lemma 1 tells us that iterations of the RDT operator shrink the circumradii, but says nothing about how much they shrink. Here, we show that if the input manifold does not come too close to the medial axis, the consecutive triangulations fit into a region that shrinks until it is everywhere close to Σ. We show that a parameter governing the distance from an iterate $\mathcal{T}_i$ to Σ converges quickly to a small fixed point. The distance bound yields bounds on the circumradii of the triangles and the accuracy of the tangents for each iterated triangulation, which we use to show that subsequent triangulations are k-manifolds. As a bonus, we considerably tighten the proximity (e.g., Hausdorff distance) bounds that apply to many standard algorithms for Delaunay surface reconstruction and surface mesh generation.

Let M be the medial axis of Σ. Define the *nearest-point map* $\nu : \mathbb{R}^d \setminus M \to \Sigma, q \mapsto \tilde{q}$, a function that maps each point to the nearest point on Σ. For brevity, we write $\tilde{q}$ for $\nu(q)$. If a point q does not have a unique nearest point on Σ, then $q \in M$ by the definition of medial axis, and $\tilde{q}$ is undefined. The homeomorphisms we employ will be the nearest-point map from some reconstruction $|\mathcal{T}|$ to Σ.

For $\omega \in [0,1)$, let the *ω-sausage* be the region $\Sigma_\omega = \{q : d(q,\Sigma) \leq \omega \operatorname{lfs}(\tilde{q})\}$, where $d(q,\Sigma)$ is the distance from q to Σ. Note that $\Sigma_0 = \Sigma$. The ω-sausage is what topologists call a *tubular neighborhood* of Σ. For $k=1$ its shape resembles a traditional sausage (albeit tied in a loop), whereas for $k=2$ it is more of a breakfast patty. Σ_ω is a d-dimensional union of $(d-k)$-balls: at each point $p \in \Sigma$ is centered one $(d-k)$-ball of radius $\omega \operatorname{lfs}(p)$, normal to Σ.

Next, we show that for any point set $P \subseteq \Sigma_{\bar{\omega}}$ with $\bar{\omega}$ a little less than 1, the fixed-point triangulation $\mathrm{Del}^*|_P V$ is included in $\Sigma_{\underline{\omega}}$ with $\underline{\omega}$ close to zero. The following lemma (see the full-length paper) bounds the interpolation error over a restricted Delaunay simplex in Σ_ω.

▶ **Lemma 4** (ϵ-Interpolation Lemma). *Let $\Sigma \subset \mathbb{R}^d$ be a smooth k-manifold without boundary. Let V be an ϵ-sample of Σ for some $\epsilon < 0.5$. Let $\tau \in \mathrm{Del}\, V$ be a simplex (of any dimension) whose dual Voronoi face τ^* intersects the ω-sausage Σ_ω. If $\omega \leq 1-\epsilon$ and $\omega < \frac{1-4\epsilon^2}{1+4\epsilon^2}$ (the latter precondition is equivalent to $\lambda < 0.5$), then for every point $x \in \tau$,*

$$|x\tilde{x}| \leq \Omega(\omega) \operatorname{lfs}(\tilde{x}), \quad \textit{where} \quad \Omega(\omega) = 1 - \sqrt{1 - \frac{\lambda^2}{(1-\lambda)^2}} \quad \textit{and} \quad \lambda = \sqrt{\frac{1+\omega}{1-\omega}}\epsilon.$$

In particular, for $\tau \in \mathrm{Del}|_\Sigma V$ (i.e., τ^* intersects Σ so we can set $\omega = 0$), Lemma 4 gives $|x\tilde{x}| \leq \left(1 - \sqrt{1 - \frac{\epsilon^2}{(1-\epsilon)^2}}\right) \operatorname{lfs}(\tilde{x}) \approx \frac{\epsilon^2}{2} \operatorname{lfs}(\tilde{x})$. Compare with the bound of $15\epsilon^2 \operatorname{lfs}(\tilde{x})$ in Cheng et al. [10, Theorem 13.22]—an improvement by a factor of up to 30.

Lemma 4 implies that if an input to the restricted Delaunay triangulation operator is in the ω-sausage for a middling value of ω, then the output is in the ω-sausage for a smaller ω.

▶ **Lemma 5** (Shrinking Sausage Lemma). *Let $\Sigma \subset \mathbb{R}^d$ be a bounded, smooth k-manifold without boundary. Let V be an ϵ-sample of Σ with $\epsilon \leq 0.3084$. Let $\underline{\omega}$ and $\bar{\omega}$ be two real solutions of $\omega = \Omega(\omega)$, where $\Omega(\omega)$ is defined in Lemma 4 and the solutions of interest are approximately*

$$\underline{\omega} = \frac{1}{2}\epsilon^2 + \epsilon^3 + \frac{17}{8}\epsilon^4 + 5\epsilon^5 + \frac{199}{16}\epsilon^6 + O(\epsilon^7) \quad \text{and} \quad \bar{\omega} = 1 - 8\epsilon^2 + 32\epsilon^4 - 384\epsilon^6 + O(\epsilon^8).$$

Let P be a point set included in the ω-sausage Σ_ω for some $\omega < \min\{1-\epsilon, \bar{\omega}\}$. Let the simplicial complex $\mathcal{T} = \mathrm{Del}^*|_P V$ be the fixed point of the restricted Delaunay triangulation operator given the input P. Then $|\mathcal{T}|$ is a subset of the $\underline{\omega}$-sausage $\Sigma_{\underline{\omega}}$.

Proof. By Lemma 4, if an input to the restricted Delaunay triangulation operator lies in the ω-sausage for $\omega < \min\left\{1-\epsilon, \frac{1-4\epsilon^2}{1+4\epsilon^2}\right\}$, then the output lies in the $\Omega(\omega)$-sausage. We observe that $\Omega(\underline{\omega}) = \underline{\omega}$; $\Omega(\bar{\omega}) = \bar{\omega}$; for any $\omega \in (\underline{\omega}, \bar{\omega})$, $\Omega(\omega) \in (\underline{\omega}, \omega)$; and for any $\omega \in [0, \underline{\omega})$, $\Omega(\omega) \in (\omega, \underline{\omega})$. (Note that $\underline{\omega} = \bar{\omega}$ for $\epsilon \doteq 0.308406$, hence the requirement that ϵ must be smaller.) Therefore, if we take an input in the ω-sausage with $\omega \in \left[0, \min\left\{1-\epsilon, \frac{1-4\epsilon^2}{1+4\epsilon^2}, \bar{\omega}\right\}\right)$ and iterate the restricted Delaunay triangulation operator until the output is a fixed point, then the output is in the $\underline{\omega}$-sausage. (As $\bar{\omega} \leq \frac{1-4\epsilon^2}{1+4\epsilon^2}$ for all $\epsilon \in [0, 0.308406]$, we omit $\frac{1-4\epsilon^2}{1+4\epsilon^2}$ from the lemma's preconditions.) ◀

We emphasize that convergence is rapid. For example, when $\epsilon = 0.143$, $\underline{\omega} \doteq 0.0145$ and $\bar{\omega} \doteq 0.8471$. If we start with a triangulation in the 0.847-sausage, after five iterations we obtain a triangulation in the 0.0147-sausage. Convergence is even faster for smaller values of ϵ. In the unlikely event that an input takes many iterations to reach its fixed point, one can always stop early; a small constant number of iterations suffices to reach a triangulation nearly as good as the fixed point.

Observe that the conditions of Lemma 5 are mild: P can be an arbitrary point set, so long as it is in the ω-sausage for a suitable ω. We have not yet guaranteed that the fixed-point triangulation is more than a messy subcomplex of $\mathrm{Del}\, V$—but it is a mess close to Σ. The next section describes conditions that ensure that the output is a 2-manifold.

6 Homeomorphism of the Restricted Delaunay Triangulation to Σ

Computational geometers have developed a theory of surface sampling to show that under the right conditions, a restricted Delaunay triangulation $\mathrm{Del}|_\Sigma V$ is homeomorphic to a smooth surface Σ [1, 2, 10, 13]. (More precisely, its underlying space $|\mathrm{Del}|_\Lambda V|$ is homeomorphic.) We present here an alternative theorem that allows us to handle a nonsmooth Λ. Specifically, we show that $|\mathrm{Del}|_\Lambda V|$ is homeomorphic to a smooth 2-manifold Σ without boundary if $V \subset \Sigma$ is a sufficiently dense ϵ-sample and Λ is a piecewise smooth 2-manifold that approximates Σ well (in both location and tangents). We note that Boissonnat and Oudot [7] developed a method for treating restricted Delaunay triangulations when Σ itself is a nonsmooth, Lipschitz surface, but our method meets our needs better.

An advantage of our proof is that it enlarges the value of ϵ for which certain properties hold. For example, we can show that the restricted Delaunay triangulation $\mathrm{Del}|_\Sigma V$ is homeomorphic to a 2-manifold $\Sigma \subset \mathbb{R}^d$ for any 0.2695-sample V, and the fixed-point triangulation $\mathrm{Del}^*|_\Sigma V$ is likewise for any 0.143-sample. (See full-length paper.) Not only does the former claim improve the constant $\epsilon = 0.18$ proven in Dey's book [13] for $\mathbb{R}^3$; it holds in any dimension d.

Let Σ be a smooth k-manifold embedded in $\mathbb{R}^d$. For every point $p \in \Sigma$, let $T_p\Sigma$ denote the k-flat (i.e., a k-dimensional affine subspace) tangent to Σ at p, called the *tangent space* at p, and let $N_p\Sigma$ denote the $(d-k)$-flat normal to Σ at p, called the *normal space* at p. In this paper we are mainly interested in $k = 2$, for which $T_p\Sigma$ is a plane.

Let $\Gamma, \Pi \subseteq \mathbb{R}^d$ be two flats such that $1 \le \dim\Gamma \le \dim\Pi \le d-1$. If they intersect each other at a point x, the angle $\angle(\Gamma,\Pi)$ separating them is $\max_{p\in\Gamma\setminus\{x\}} \min_{q\in\Pi\setminus\{x\}} \angle pxq$. If they do not intersect each other, translate Γ so they intersect, then measure the angle as above. This angle is never obtuse. We sometimes substitute line segments or Voronoi faces for flats, in which case we mean to measure the nonobtuse angles separating their affine hulls. It is well known that for any two points $p, q \in \Sigma$, $\angle(N_p\Sigma, N_q\Sigma) = \angle(T_p\Sigma, T_q\Sigma)$.

Let $\Lambda \subset \mathbb{R}^d$ be a piecewise smooth k-manifold. We say that Λ is *θ-parallel* to Σ if for every smooth piece $\Gamma \subseteq \Lambda$ and every point $p \in \Gamma$, $\angle(T_p\Gamma, T_{\tilde{p}}\Sigma) \le \theta$.

For any Voronoi face $f \in \operatorname{Vor} V$ of dimension 1 or greater, we say that f is *ϕ-normal* to Σ at $p \in f$ if $\angle(f, N_{\tilde{p}}\Sigma) \le \phi$. That is, f is nearly parallel to $\tilde{p}$'s normal space.

To prove that Σ and $|\operatorname{Del}|_\Lambda V|$ are homeomorphic for $k = 2$, we require that $\Lambda \subset \Sigma_\omega$ for a suitable ω, that Λ be everywhere θ-parallel to Σ, and that the faces of $\operatorname{Vor} V$ of dimension $d-2$ or greater be ϕ-normal to Σ everywhere in the ω-sausage for an appropriate ω, where $\theta + \phi < 90°$. The following lemma, proven in the full-length version of this paper, shows that Λ approximates the tangent and normal spaces of Σ with an error linear in ϵ.

▶ **Lemma 6** (Normal-Parallel Lemma). *Let $\Sigma \subset \mathbb{R}^d$ be a bounded, smooth 2-manifold without boundary. Let $V \subset \Sigma$ be an ϵ-sample of Σ. Let $\tau \in \operatorname{Del} V$ be a triangle whose Voronoi dual face τ^* intersects Σ_ω for some $\omega \le 1-\epsilon$. Let N_τ be a flat dual (complementary) to* $\operatorname{aff}\tau$. *Let*

$$\chi = \arcsin\delta + \arcsin\left(\frac{2}{\sqrt{3}}\sin(2\arcsin\delta)\right),\ \ \delta = \frac{\sqrt{\omega^2 + (1+\omega)\epsilon^2}}{1-\lambda},\ \ \lambda = \sqrt{\frac{1+\omega}{1-\omega}}\epsilon,\ \text{and}$$

$$\eta = \arctan\frac{\lambda\sqrt[4]{1-\lambda^2}}{\sqrt{1-\lambda^2\left(1+\sqrt{1-\lambda^2}\right)}} \approx \lambda + \frac{5}{12}\lambda^3 + \frac{57}{160}\lambda^5 + \frac{327}{896}\lambda^7 + O(\lambda^9).$$

Then for every point $p \in \tau$, $\angle(\operatorname{aff}\tau, T_{\tilde{p}}\Sigma) = \angle(N_\tau, N_{\tilde{p}}\Sigma) \le \chi + 2\eta$, and for every point $z \in \tau^* \cap \Sigma_\omega$, $\angle(\operatorname{aff}\tau, T_{\tilde{z}}\Sigma) = \angle(N_\tau, N_{\tilde{z}}\Sigma) \le \chi + \eta$. Therefore, τ is $(\chi+2\eta)$-parallel to Σ, and τ^* is $(\chi+\eta)$-normal to Σ wherever it intersects Σ_ω.

Proofs about the homeomorphism of restricted Delaunay triangulations usually rely on the *Topological Ball Theorem* of Edelsbrunner and Shah [15], and ours is no exception. The Topological Ball Theorem states that $|\operatorname{Del}|_\Lambda V|$ is homeomorphic to Λ if the following *topological ball property* holds: for every Voronoi j-face f that intersects Λ, $j \in [d-k, d]$, $f \cap \Lambda$ is a topological $(j+k-d)$-ball and $(\operatorname{int} f) \cap \Lambda = \operatorname{int}(f \cap \Lambda)$. (Here $\operatorname{int} f$ is an "interior" with respect to $\operatorname{aff} f$, and $\operatorname{int}(f \cap \Lambda)$ is an "interior" with respect to the manifold $f \cap \Lambda$, not with respect to the ambient space $\mathbb{R}^d$.) Moreover, Λ intersects no Voronoi j-face for $j < d-k$. The last condition can be simulated by infinitesimal perturbations of Λ. The following theorem is one of our main contributions; see the full-length paper for a proof.

▶ **Theorem 7** (θ-ϕ Homeomorphism Theorem). *Let $\Sigma \subset \mathbb{R}^d$ be a connected, smooth 2-manifold. Let $\Lambda \subset \Sigma_\omega$ be a piecewise smooth 2-manifold where $\omega < (\sqrt{5}-1-2\epsilon^2)/(\sqrt{5}-1+2\epsilon^2)$ (equivalently, $\lambda < \sqrt{\left(\sqrt{5}-1\right)/2} \doteq 0.7861$). Suppose that there is a $\theta \ge 0°$ and a $\phi \ge 0°$ such that $\theta + \phi < 90°$, Λ is θ-parallel to Σ, and every face $f \in \operatorname{Vor} V$ of dimension $d-2$ or $d-1$ is ϕ-normal to Σ at every point on $f \cap \Sigma_\omega$. Suppose that for every $p \in \Sigma$, Λ intersects $N_p\Sigma$ at most once (i.e., $\nu : \Lambda \to \Sigma$ is injective), and that Λ intersects no i-face of $\operatorname{Vor} V$ for $i < d-2$. Then for every $j \in [0,2]$ and every $(d-2+j)$-dimensional face $f \in \operatorname{Vor} V$ such that $f \cap \Lambda \ne \emptyset$, the intersection $f \cap \Lambda$ is a topological j-ball and $(\operatorname{int} f) \cap \Lambda = \operatorname{int}(f \cap \Lambda)$. Thus Λ and $\operatorname{Vor} V$ satisfy the topological ball property and $|\operatorname{Del}|_\Lambda V|$ is homeomorphic to Λ.*

▶ **Corollary 8.** *Under the assumptions of Theorem 7, $|\operatorname{Del}|_\Lambda V|$ is homeomorphic to Σ.*

The obvious use of Corollary 8 is to show that, given a suitable Λ, the fixed-point procedure produces a triangulation of Σ. We have a second use for it: our algorithm in Section 8 does not construct a suitable Λ; instead, it directly computes all the triangles that might be in a fixed-point triangulation. By the following corollary (proven in the full-length paper), we know that some subset of those triangles is a good triangulation of Σ.

▶ **Corollary 9.** *Let $\Sigma \subset \mathbb{R}^3$ be a bounded, smooth 2-manifold without boundary. Let $V \subset \Sigma$ be a finite 0.143-sample of Σ. Then the underlying space of the fixed-point triangulation $\mathrm{Del}^*|_\Sigma V$ is homeomorphic to Σ.*

7 Classifying Critical Points

Recall that a simplex $\tau \in \mathrm{Del}\, V$ is *anchored* if it intersects its own Voronoi dual τ^*. We call the intersection $\tau \cap \tau^*$ a *critical point* (of the distance function $d(p, V) = \min_{u \in V} |pu|$). Dey, Giesen, Ramos, and Sadri [12] show that for a sufficiently dense sample, every critical point is either a *manifold critical point* close to the manifold Σ or a *medial critical point* close to Σ's medial axis M. The following theorem is our tighter version of their result.

▶ **Theorem 10** (Critical Point Separation Theorem)**.** *Let $\Sigma \subset \mathbb{R}^d$ be a smooth k-manifold. Let V be an ϵ-sample of Σ for some $\epsilon \leq 0.5$. Let $\tau \in \mathrm{Del}\, V$ be a simplex (of any dimension) whose dual Voronoi face τ^* intersects τ. Let $z = \tau \cap \tau^*$. Let*

$$\mathring{\omega} = \frac{2 - \sqrt{4 - 12\epsilon^2 + \epsilon^4} - \epsilon^2}{4} = \frac{1}{2}\epsilon^2 + \frac{1}{2}\epsilon^4 + \frac{3}{4}\epsilon^6 + \frac{11}{8}\epsilon^8 + O(\epsilon^{10}) \quad \textit{and}$$

$$\mu = \frac{1 - \sqrt{1 - 4\epsilon^2}}{2} = \epsilon^2 + \epsilon^4 + 2\epsilon^6 + 5\epsilon^8 + O(\epsilon^{10}).$$

Then either $|z\tilde{z}| \leq \mathring{\omega}\, \mathrm{lfs}(\tilde{z})$ (i.e., $z \in \Sigma_{\mathring{\omega}}$) or $|zm| \leq \mu |\tilde{z}m|$ for some $m \in M$, where M is the medial axis of Σ.

Here, we describe a method for distinguishing the two types of critical points that works for larger values of ϵ (in theory) than the well-known *cocone criterion* of Amenta, Choi, Dey, and Leekha [2]. Our method is useful not only for our algorithm, but also for Dey et al.'s algorithm for computing flow complexes. In particular, our test works in any dimension, and enables the computation of fixed-point complexes and flow complexes in higher dimensions. The cocone criterion depends on having accurate estimates of the local tangent planes, but in high dimensions only poor approximations can be obtained cheaply.

Let v be a vertex of an anchored simplex τ. Let $v^* \in \mathrm{Vor}\, V$ denote its Voronoi cell. Consider the set of points at the centers of the medial balls tangent to Σ at v: these centers all lie on the normal space $N_v\Sigma$, and each one lies either on the medial axis M or at infinity. If the codimension is 1 (i.e., $d - k = 1$), there are two such medial points; if the codimension is greater, there are infinitely many. In either case, at least one medial point is not at infinity, and all these medial points, including those at infinity, lie in the Voronoi cell v^*.

Consider first the case of codimension 1. Suppose that we can explicitly construct the Voronoi diagram (i.e., d is small). Following Amenta and Bern [1], we define the *poles of v* to be two vertices of v^*, one on each side of Σ. Let the pole p^+ be the vertex of v^* furthest from v, unless v^* is unbounded, in which case p^+ is a point at infinity whose direction is chosen so the ray vp^+ lies in v^*. Let the pole p^- be the vertex of v^* furthest from v such that $\angle p^+ v p^- > 90°$. Amenta and Bern show that p^- is on the opposite side of Σ from p^+.

Our classifier for codimension 1 is simple. Given a critical point $z = \tau \cap \tau^*$, choose any vertex v of τ and compute v's poles p^+ and p^-. Set $p = p^-$ if v^* is unbounded; otherwise,

select the pole $p \in \{p^+, p^-\}$ that minimizes $\angle pvz$. Then, z is a medial critical point if $|vz|/|vp| > 0.313$; a manifold critical point otherwise. This classifier relies on a lemma of Amenta and Bern, which generalizes to higher dimensions and codimensions without change.

▶ **Lemma 11** (Normal Lemma [1]). *Let $\Sigma \subset \mathbb{R}^d$ be a bounded, smooth k-manifold without boundary. Let $V \subset \Sigma$ be an ϵ-sample of Σ for $\epsilon < 1$. Let p be a point in the Voronoi cell v^* of a sample $v \in V$ such that $|vp| \geq \rho \operatorname{lfs}(v)$ for some $\rho > 0$. Then $\angle(vp, N_v\Sigma) \leq \beta(\rho)$, where $\beta(\rho) = \arcsin \frac{\epsilon}{1-\epsilon} + \arcsin \frac{\epsilon}{\rho(1-\epsilon)}$.*

▶ **Lemma 12** (Classification Lemma for Codimension 1). *Let $\Sigma \subset \mathbb{R}^d$ be a bounded, smooth $(d-1)$-manifold without boundary. Let $V \subset \Sigma$ be an ϵ-sample of Σ for $\epsilon \leq 0.2327$. Let $\tau \in \operatorname{Del} V$ be a simplex such that $\tau \cap \tau^* \neq \emptyset$, and let z be τ's circumcenter $\tau \cap \tau^*$. Let v be any vertex of τ. Choose a pole p of the Voronoi cell v^* as described above. If $|vz|/|vp| > 0.313$, then z is a medial critical point; otherwise, z is a manifold critical point.*

Proof. As v^* contains medial points, each at a distance of at least $\operatorname{lfs}(v)$ from v, $|vp^+| \geq \operatorname{lfs}(v)$. Amenta and Bern [1] show that $|vp^-| \geq \operatorname{lfs}(v)$ too. Hence $|vp| \geq \operatorname{lfs}(v)$.

If z is a manifold critical point, then z lies in the $\mathring{\omega}$-sausage by Theorem 10. In the full-length paper we show that the circumradius of τ is $|vz| \leq \sqrt{\mathring{\omega}^2 + (1+\mathring{\omega})\epsilon^2} \operatorname{lfs}(\tilde{z})$, and $|\tilde{z}v| \leq \lambda \operatorname{lfs}(\tilde{z})$, where $\lambda = \sqrt{\frac{1+\mathring{\omega}}{1-\mathring{\omega}}}\epsilon$. By the Feature Translation Lemma [13, 10], $|vz| \leq \sqrt{\mathring{\omega}^2 + (1+\mathring{\omega})\epsilon^2} \operatorname{lfs}(v)/(1-\lambda)$. For $\epsilon \leq 0.2327$, $\mathring{\omega} < 0.0287$ and $|vz|/|vp| < 0.3127$.

If z is a medial critical point, then there is a medial axis point $m \in M$ such that $|zm| \leq \mu|\tilde{z}m| \leq \mu(|\tilde{z}z| + |zm|)$; therefore $|zm| \leq |\tilde{z}z|\mu/(1-\mu)$. Hence, $\operatorname{lfs}(v) \leq |vm| \leq |vz| + |zm| \leq |vz| + |\tilde{z}z|\mu/(1-\mu) \leq |vz| + |vz|\mu/(1-\mu) = |vz|/(1-\mu)$ and $|vz| \geq (1-\mu)\operatorname{lfs}(v)$.

By the Normal Lemma, $\angle(vp, N_v\Sigma) \leq \beta(1)$ and $\angle(vz, N_v\Sigma) \leq \beta(1-\mu)$. For codimension 1, $N_v\Sigma$ is a line, so one pole satisfies $\angle pvz \leq \beta(1) + \beta(1-\mu)$ and one satisfies $\angle pvz \geq 180° - \beta(1) - \beta(1-\mu)$. We choose p to be the former.

As v is a vertex of τ and z is the circumcenter of τ, every point q for which $\angle vzq > 90°$ is closer to another vertex of τ and cannot lie in the Voronoi cell v^*. But $p \in v^*$, so $\angle vzp \leq 90°$ and $|vz|/|vp| \geq \cos\angle pvz \geq \cos(\beta(1) + \beta(1-\mu))$. For $\epsilon \leq 0.2327$, $|vz|/|vp| > 0.3134$. ◀

We adapt this test to 2-manifolds in codimension 2 or higher by taking a three-dimensional cross-section of the Voronoi cell v^*. First, we use the following simple algorithm to compute an approximate tangent plane $\hat{T}_v\Sigma$ for v: find v's nearest neighbor $w \in V$, then find the vertex $x \in V$ that minimizes the circumradius of $\triangle vwx$. We show in the full-length paper that $\triangle vwx$ has a small circumradius and its affine hull $\hat{T}_v\Sigma = \operatorname{aff} \triangle vwx$ is an adequate approximation of v's tangent plane $T_v\Sigma$. Let Ξ be the 3-flat that includes $\hat{T}_v\Sigma$ and the critical point z. We compute the cross-section $\Xi \cap v^*$ with the method described in Section 2. As we seek just one cell in a three-dimensional power diagram, this computation can be dualized to a three-dimensional convex hull computation, which can be performed in $O(n \log n)$ time.

If we knew the true tangent plane $T_v\Sigma$, we could use the same test and analysis as Lemma 12; but only an approximation $\hat{T}_v\Sigma$ is available, so we adjust the test to compensate and require a stricter constraint on ϵ. (A better tangent plane approximation method would enable us to loosen this constraint.) Given a critical point $z = \tau \cap \tau^*$, choose any vertex v of τ, compute $\hat{T}_v\Sigma$, let $\Xi = \operatorname{aff}(\hat{T}_v\Sigma \cup \{z\})$, and compute the three-dimensional polyhedron $\Xi \cap v^*$ and its poles p^+ and p^-. Select a pole $p \in \{p^+, p^-\}$ in the same manner as for codimension 1. Then, z is a medial critical point if $|vz|/|vp| > 0.182$; a manifold critical point otherwise. This test works for any $\epsilon \leq 0.1522$. We defer the proof to the full-length paper.

8 Constructing Fixed-Point Complexes

We classify each anchored triangle as a *manifold anchored triangle* or a *medial anchored triangle* according to whether its circumcenter is a manifold critical point or a medial critical point. A *chained triangle* is a triangle that is generated by a manifold anchored triangle or another chained triangle. The *manifold fixed-point complex* (MF complex) is the set of all manifold anchored triangles and chained triangles. Like the flow complex, the MF complex is not a manifold (because of slivers), but it is a useful geometric representation of Σ in other ways: its Hausdorff distance to Σ is small and all its triangles are good estimates of Σ's local tangent space. Unlike the flow complex, the MF complex obtains the latter guarantee without prohibiting sample points from being too close together. Being a subcomplex of $\operatorname{Del} V$, the MF complex is a simpler alternative to the flow complex that offers better approximation bounds. Moreover, for $\epsilon \leq 0.143$, some subcomplex of the MF complex is a 2-manifold that approximates Σ well (namely, $\operatorname{Del}^*|_\Sigma V$). (We conjecture that for small ϵ, every MF triangle participates in at least one such 2-manifold.)

Here, we present an algorithm that constructs the MF complex for a 2-manifold $\Sigma \subset \mathbb{R}^d$ in time polynomial in d. In the special case $d = 3$, standard methods can efficiently extract a 2-manifold from the MF complex; thus we have a surface reconstruction algorithm. In higher dimensions, we believe heuristics can do the same in practice.

We take as input an n-point ϵ-sample $V \subset \Sigma$. If possible, we begin by constructing $\operatorname{Del} V$ and $\operatorname{Vor} V$, then we use $\operatorname{Del} V$ to identify the manifold anchored triangles. If d is too large for the construction of $\operatorname{Del} V$ to be feasible, we identify the anchored triangles by enumerating candidate triangles and checking each one to see if its diametric sphere is empty.

Next, we classify the anchored triangles as manifold or medial as described in Section 7. We find the chained triangles as described in Section 4. This requires us to determine which $(d-2)$-faces of $\operatorname{Vor} V$ intersect anchored or chained triangles. If we have constructed $\operatorname{Del} V$, an efficient way to find these Delaunay-Voronoi intersections is to simultaneously "walk" through both $\operatorname{Del} V$ and $\operatorname{Vor} V$ in synchrony, working outward from each anchored triangle along the chained triangles in a depth-first manner. If constructing $\operatorname{Del} V$ in infeasible, we construct a Voronoi cross-section for each MF triangle as described in Section 2.

Let $\mathcal{T}$ be the MF complex. In $\mathbb{R}^3$, we can extract a 2-manifold from $\mathcal{T}$ with a standard method [2, 13] whose correctness is well known. All MF triangles lie sufficiently close to Σ (in $\Sigma_{\underline{\omega}}$) that we can assign each edge two "sides": two triangles in $\mathcal{T}$ are on the same "side" of a shared edge if they meet at a dihedral angle less than 90°; on opposite sides otherwise. We *prune* from $\mathcal{T}$ every triangle that has an edge with no matching triangle on the other side. Each pruned triangle may render other edges one-sided, prompting further pruning. When no such triangle survives, we label the tetrahedra of $\operatorname{Del} V$ that are not enclosed by the surviving triangles by performing a depth-first search. The search begins at the unbounded outer cell of $\operatorname{Del} V$ and it labels every tetrahedron it visits "outside," but it never crosses a triangle still in $\mathcal{T}$. Finally, we output every triangle that separates an unlabeled tetrahedron from a tetrahedron labeled "outside" or the outer cell of $\operatorname{Del} V$. The method succeeds because by Corollary 9, $\mathcal{T}$ includes the 2-manifold $\operatorname{Del}^*|_\Sigma V$, whose triangles are not pruned and which encloses tetrahedra that will not be labeled "outside." (The output manifold might enclose additional tetrahedra, which do no harm.)

▶ **Theorem 13.** *Let $\Sigma \subset \mathbb{R}^d$ be a bounded, smooth 2-manifold without boundary. Let $V \subset \Sigma$ be a finite ϵ-sample of Σ for $\epsilon \leq 0.143$. Our fixed-point reconstruction algorithm produces a Delaunay subcomplex $\mathcal{T} \subset \operatorname{Del} V$ such that the map $\nu : p \mapsto \tilde{p}$ sends each point $p \in \mathcal{T}$ a distance of at most $\underline{\omega}\operatorname{lfs}(\tilde{p})$, where $\underline{\omega} \approx \frac{1}{2}\epsilon^2 + \epsilon^3 + \frac{17}{8}\epsilon^4 + 5\epsilon^5 + \frac{199}{16}\epsilon^6$ and $\tilde{p}$ is the point*

nearest p on Σ; and for every triangle $\tau \in \mathcal{T}$ and every point $p \in \tau$, $\angle(\text{aff}\,\tau, T_{\tilde{p}}\Sigma) \leq \chi + 2\eta$, where χ and η are defined in Lemma 6 with $\omega \leftarrow \underline{\omega}$. Moreover, some subset of $\mathcal{T}$ has an underlying space homeomorphic to Σ, and for $d = 3$ we can extract it efficiently.

In almost all practical cases, the MF complex has $O(n)$ triangles, but by placing $\Theta(n)$ sample points on the boundary of an empty sphere, it is possible to create pathological, sliver-packed examples that have $\Theta(n^2)$ triangles in $\mathbb{R}^3$ or $\Theta(n^3)$ triangles in $\mathbb{R}^5$.

If we explicitly construct $\operatorname{Del} V$ and $\operatorname{Vor} V$, the randomized incremental algorithm of Clarkson and Shor [11] constructs them in expected $O(n^{\lceil d/2 \rceil})$ time, but for many point sets that arise in practice it takes expected $O(n \log n)$ time [3]. Then it takes time linear in the size of $\operatorname{Del} V$ to identify the anchored triangles and (in $\mathbb{R}^3$) to extract the final manifold at the end. By walking through $\operatorname{Del} V$, we identify the chained triangles at a worst-case cost of $O(n)$ time per MF triangle, but the typical cost is closer to $O(1)$ per triangle.

If d is large enough that we cannot construct $\operatorname{Del} V$, the brute force algorithm for identifying the anchor triangles (test all $\binom{n}{3}$ potential triangles for an empty diametric sphere) takes $O(n^4)$ time. However, as manifold anchored triangles tend to be small, for many point sets that arise in practice, it may be possible to reduce this time substantially by use of bucketing, spatial hashing, or quadtrees. We can identify the chain triangles by computing a Voronoi cross-section for each MF triangle at a cost of $O(n \log n)$ time per triangle.

In $\mathbb{R}^3$, we identify the poles of every Voronoi cell and classify the critical points as manifold or medial in time linear in the size of $\operatorname{Vor} V$. For $d \geq 4$, we compute a three-dimensional Voronoi cell cross-section for each critical point at a cost of $O(n \log n)$ time each.

9 Conclusions

The main open problem we would like to solve is to efficiently find a fixed-point triangulation that is homeomorphic to Σ for $d \geq 4$. There are two very different approaches: one could find a "starter" triangulation Λ that approximates Σ well enough to apply fixed-point iterations; or one could find an efficient algorithm for extracting a 2-manifold from the MF complex. Both problems have stumped us in theory despite long efforts. As we noted in Section 3, we suspect that manifold extraction is hard in theory, but easy in practice.

It seems promising to consider what the MF complex looks like if we permit simplices of dimension greater than 2, even when Σ is a 2-manifold. In three dimensions, for instance, the MF complex would include sliver tetrahedra that contain their own circumcenters and perhaps other, chained, sliver tetrahedra. We conjecture that for $\epsilon \leq 0.143$, this full-dimensional MF complex is homotopy equivalent to Σ.

We are oddly optimistic that our ideas can be modified to find good representations of higher-dimensional manifolds, despite the problems we discuss at the end of Section 1. We observe that the simplices with the most terrible normals and tangents have Voronoi duals that lie close to Σ. Counterintuitively, we could perhaps better model Σ by taking a restricted Delaunay triangulation $\operatorname{Del}|_\Lambda V$ for a manifold Λ that approximates Σ but is displaced far enough away from Σ to dodge the Voronoi duals of the most troublesome simplices. Such a triangulation would not be a fixed point of the restricted Delaunay triangulation operator, but it might be found by similar methods.

Acknowledgments. The author J. Shewchuk began this work during a sabbatical at INRIA Sophia-Antipolis during the spring of 2010. He would like to thank the Geometrica Group for their kind reception, and he particularly thanks Jean-Daniel Boissonnat and Arijit Ghosh for extended discussions of the problem of manifold reconstruction.

This work was supported in part by INRIA Sophia-Antipolis, in part by the National Science Foundation under Awards CCF-0635381, IIS-0915462, and CCF-1423560, in part by the University of California Lab Fees Research Program under Grant 09-LR-01-118889-OBRJ, in part by a National Science Foundation Graduate Research Fellowship, and in part by an Alfred P. Sloan Research Fellowship. We thank our sponsors for their support.

References

1 Nina Amenta and Marshall Bern. Surface Reconstruction by Voronoi Filtering. *Discrete & Computational Geometry*, 22(4):481–504, June 1999.

2 Nina Amenta, Sunghee Choi, Tamal Krishna Dey, and Naveen Leekha. A Simple Algorithm for Homeomorphic Surface Reconstruction. *International Journal of Computational Geometry and Applications*, 12(1–2):125–141, 2002.

3 Nina Amenta, Sunghee Choi, and Günter Rote. Incremental Constructions con BRIO. In *Proceedings of the Nineteenth Annual Symposium on Computational Geometry*, pages 211–219, San Diego, California, June 2003. Association for Computing Machinery.

4 Marshall Bern and David Eppstein. Mesh Generation and Optimal Triangulation. In *Computing in Euclidean Geometry*, volume 1 of *Lecture Notes Series on Computing*, pages 23–90. World Scientific, Singapore, 1992.

5 Jean-Daniel Boissonnat and Arijit Ghosh. Manifold Reconstruction Using Tangential Delaunay Complexes. *Discrete & Computational Geometry*, 51(1):221–267, January 2014.

6 Jean-Daniel Boissonnat and Steve Oudot. Provably Good Sampling and Meshing of Surfaces. *Graphical Models*, 67(5):405–451, September 2005.

7 Jean-Daniel Boissonnat and Steve Oudot. Provably Good Sampling and Meshing of Lipschitz Surfaces. In *Proceedings of the Twenty-Second Annual Symposium on Computational Geometry*, pages 337–346, Sedona, Arizona, June 2006.

8 Siu-Wing Cheng, Tamal Krishna Dey, Herbert Edelsbrunner, Michael A. Facello, and Shang-Hua Teng. Sliver Exudation. *Journal of the ACM*, 47(5):883–904, September 2000.

9 Siu-Wing Cheng, Tamal Krishna Dey, and Edgar A. Ramos. Manifold Reconstruction from Point Samples. In *Proceedings of the Sixteenth Annual Symposium on Discrete Algorithms*, pages 1018–1027, Vancouver, British Columbia, Canada, January 2005. ACM–SIAM.

10 Siu-Wing Cheng, Tamal Krishna Dey, and Jonathan Richard Shewchuk. *Delaunay Mesh Generation*. CRC Press, Boca Raton, Florida, December 2012.

11 Kenneth L. Clarkson and Peter W. Shor. Applications of Random Sampling in Computational Geometry, II. *Discrete & Computational Geometry*, 4(1):387–421, December 1989.

12 Tamal K. Dey, Joachim Giesen, Edgar A. Ramos, and Bardia Sadri. Critical Points of Distance to an ϵ-Sampling of a Surface and Flow-Complex-Based Surface Reconstruction. *International Journal of Computational Geometry and Applications*, 18(1–2):29–62, 2008.

13 Tamal Krishna Dey. *Curve and Surface Reconstruction: Algorithms with Mathematical Analysis*. Cambridge University Press, New York, 2007.

14 Herbert Edelsbrunner. Surface Reconstruction by Wrapping Finite Sets in Space. In Boris Aronov, Saugata Basu, János Pach, and Micha Sharir, editors, *Discrete and Computational Geometry: The Goodman–Pollack Festschrift*, pages 379–404. Springer-Verlag, Berlin, 2003.

15 Herbert Edelsbrunner and Nimish R. Shah. Triangulating Topological Spaces. *International Journal of Computational Geometry and Applications*, 7(4):365–378, August 1997.

16 Julia Flötotto. *A Coordinate System Associated to a Point Cloud Issued from a Manifold: Definition, Properties and Applications*. PhD thesis, Université Nice Sophia Antipolis, 2003.

17 Joachim Giesen and Matthias John. The Flow Complex: A Data Structure for Geometric Modeling. *Computational Geometry: Theory and Applications*, 39:178–190, 2008.