

Medial Axis Approximation and Unstable Flow Complex

Joachim Giesen* Edgar A. Ramos † Bardia Sadri ‡

Abstract

The medial axis of a shape is known to carry a lot of information about it. In particular a recent result of Lieutier establishes that every bounded open subset of \mathbb{R}^n has the same homotopy type as its medial axis. In this paper we provide an algorithm that, given a sufficiently dense but not necessarily uniform sample from the surface of a shape with smooth boundary, computes a *core* for its medial axis approximation, in form of a piecewise linear cell complex, that captures the topology of the medial axis of the shape. We also provide a natural method to freely augment this core in order to enhance it geometrically all the while maintaining its topological guarantees. The definition of the core and its extension method are based on the steepest ascent flow induced by the distance function to the sample. We also provide a geometric guarantee on the closeness of the core and the actual medial axis.

1 Introduction

The medial axis of a bounded open set S in \mathbb{R}^n is the set of points in S with at least two closest points in the boundary of S . In the following we sometimes refer to bounded open subsets of \mathbb{R}^n as shapes. A recent result of Lieutier [Lie04] establishes that any shape and its medial axis have the same topological type, or more precisely, are homotopy equivalent. Consequently, the medial axis can be used to answer topological queries about the shape. Therefore, it can be crucial for a medial axis approximation algorithm to capture the topology of the medial axis when its output is used in applications that make such queries including but not limited to shape analysis, motion planning, and mesh partitioning.

As a geometric object, the medial axis is unstable since small changes in the shape can cause comparatively large changes in its medial axis. This instability of the medial axis bears two consequences. First, it makes the medial axis hard to compute exactly because of numerical instabilities; consequently, exact computation of medial axis has only been attempted for a few limited classes of shapes (see for example [CKM04]). Second, the complete medial axis may be less interesting in practice than an approximation of it which carries the same topological type but is more stable under small perturbations of the shape. Approximations of the medial axis of a shape are often sought with a sample of the boundary of the shape provided as input. Chazal and Lieutier [CL05] defined the λ -medial axis, a subset of the medial axis, which has the desired stability, and is guaranteed to have the same homotopy type as the medial axis for suitably small values of the parameter λ . The largest topologically safe λ depends on the shape and can be very small and hard to determine [ABE04]. Finally, the existing algorithm for computing the λ -medial axis requires a very dense *uniform* sample of the boundary of the shape [CL05].

The ε -sampling theory of Amenta and Bern [AB99] provides a framework for analysis of algorithms that either reconstruct the boundary of a shape or approximate its medial axis using a sample whose density varies and is proportional to the local size of the features of the boundary of the shape. Several results were obtained in this framework, some of the earliest of which pertaining the medial axis are due to Amenta, Bern, and Eppstein [ABE98], and Boissonnat and Cazals [BC02], each establishing that a subset of the Voronoi vertices in the Voronoi diagram of the sample points lie close to the medial axis. Later, Amenta, Choi, and Kolluri [ACK01] designed an algorithm that computes an approximation of the medial axis, in shape of a cell complex, that is provably homotopy equivalent to the medial axis of the shape, provided that that shape boundary is sampled densely enough. The output of this algorithm tends to be noisy and

*Theoretische Informatik, ETH Zürich, CH-8092 Zürich, giesen@inf.ethz.ch. Research partially supported by the Swiss National Foundation under the project “Non-linear manifold learning”

†Department of Computer Science, University of Illinois, Urbana, IL 61801, eramosn@cs.uiuc.edu

‡Department of Computer Science, University of Illinois, Urbana, IL 61801, sadri@cs.uiuc.edu

needs to be cleaned up using heuristics for practical use. The cell complex produced by this algorithm is *not* contained in the Voronoi 2-skeleton of the input sample, a property sought and considered plausible by many medial axis approximation algorithms due to the fact that the Voronoi 2-skeleton is indeed the medial axis of the sample and that effective filtering of it can often produce good results in practice. Dey and Zhao [DZ03] addressed these shortcomings by designing an algorithm that outputs a sub-complex of the Voronoi diagram as a medial axis approximation. This often geometrically pleasing output is guaranteed to converge to the true medial axis when the sample grows infinitely dense, i.e., when $\varepsilon \rightarrow \infty$, but suffers from lacking any topological guarantees. In fact, despite its celebrated geometric quality, with a poor choice of filtering parameters, the provided output can be highly flawed topologically. The results of the present paper can be used to mend this deficiency.

Our Contribution. Our assumptions and the approach we pursue in this paper also fall into the Amenta-Bern framework. We compute a piece-wise linear cell complex, which we call the *core* (of the medial axis approximation). This core is extracted from a refinement of the Voronoi complex of the input sample called the *unstable flow complex*. The idea and the algorithm to compute this complex and the core are derived from the critical point theory of the distance function to the sample points and their induced steepest ascent flow, which we refer to as the *discrete flow*.

Dey et al. [DGRS05] show that the critical points of the distance function to an ε -sampling of the boundary of a shape are naturally separated; each such critical point is either very close to the surface, or oppositely, it is very close to the medial axis of the shape. The two classes can be separated algorithmically provided that the sample is sufficiently dense.

In short, the core we compute is the union of the *unstable manifolds* of the critical points close to the medial axis. The unstable manifold of a critical point consists of all points that can be reached from an infinitesimally small neighborhood of that critical point following the discrete flow. We characterize the structure of the unstable manifolds and show how to compute them from the Voronoi complex of the sample points in \mathbb{R}^3 . Our main result states that the core and the medial axis are homotopy equivalent at every dimension for dense enough samples.

Capturing the homotopy type of the medial axis is not the only advantage of the core. Once the core is computed, it can be *extended* “freely” to better capture the geometry of the medial axis while maintaining its homotopy type. To extend the core one picks, using any algorithm at hand, a set of points that approximates a subset of the medial axis of the target shape (see for example [DZ03] for one such algorithm) and adds this set to the computed core along with its *flow closure* under the discrete flow. We show that as long as the chosen approximating set of points does not get too close to the surface of the shape, adding the closure “recaptures” the topological guarantee of the core. Because of this property, computing the core can be used to augment virtually any medial axis approximation algorithm into a topologically accurate one.

Finally we provide an upper-bound on the rate at which the distance of a point to the medial axis can grow as the point moves along its flow line in the discrete flow. This shows that discrete flow closure of points near the medial axis does not escape the medial axis rapidly. A consequence of such a geometric guarantee is that the core gets closer and closer to the medial axis and is contained in it at the limit as $\varepsilon \rightarrow 0$. Moreover, if the core is extended using a set that converges to the real medial axis when $\varepsilon \rightarrow 0$, such as the Voronoi facets filtered by the algorithm of [DZ03], then the extended core also converges to the real medial axis in Hausdorff distance all the while staying homotopy equivalent to it.

The structure of the paper is as follows. In Section 2 we introduce basic definitions and state some known results. Section 3 establishes several homotopy equivalences, most importantly the homotopy equivalence of the medial axis, the core, and the extended core. Section 4 describes how to provide geometric guarantees for the core and flow closures. Concluding remarks and experimental results are provided in Section 5. Appendix A covers the omitted proofs. Finally, Appendix B explains in detail the algorithms for computing the unstable flow complex, the core, and flow closures in \mathbb{R}^3 .

2 Preliminaries

In this section we review some basic definitions as well as some required background material. It is important to notice that although we present all the definitions for \mathbb{R}^3 , all the statements and results of the present

section as well as those in sections 3 and 4 generalize to higher dimensional Euclidean spaces. The algorithms presented in Appendix B, however, are designed with \mathbb{R}^3 considered.

Basic Notions. For a point $x \in \mathbb{R}^3$ and $r > 0$, we denote by $B(x, r)$ the open ball of radius r centered at x , i.e., $B(x, r) = \{y \in \mathbb{R}^3 : \|x - y\| < r\}$, and by $\overline{B}(x, r)$ the closure of $B(x, r)$, i.e., $\overline{B}(x, r) = \{y \in \mathbb{R}^3 : \|x - y\| \leq r\}$. We refer to $B(x, r)$ as the r -neighborhood of x . For $x \in \mathbb{R}^3$ and any subset $S \subset \mathbb{R}^3$, we define the distance between x and S as $d(x, S) = \inf_{y \in S} \|x - y\|$. By S^c we denote the complement of S in \mathbb{R}^3 .

Shape, Surface, and Medial Axis. We consider single-component smooth 2-manifolds without boundary embedded in \mathbb{R}^3 . We call such manifolds *surfaces*. A surface Σ is associated to the two *open* components of its complement $\mathbb{R}^3 \setminus \Sigma$ which we refer to as the bounded or *inner* component S , and the unbounded or *outer* component S^* . We will refer to the bounded component as the *shape* enclosed by Σ and denote it throughout by S . Likewise, in the rest of this paper, S^* refers to the unbounded component associated to the considered surface Σ . The *medial axis* $M(S)$ of an open set S is the set of all points in S that have at least two closest points in ∂S , the boundary of S , i.e. $M(S) = \{x \in S : |A_S(x)| > 1\}$, where $A_S(x)$ is the set of closest points to x in $\partial S = \Sigma$. Note that since Σ is compact, $A_S(x)$ is well-defined and non-empty for every $x \in S$. By medial axis M of a surface Σ , we mean the union of the medial axes of the inner and outer components S and S^* associated to Σ . Thus $M = M(S) \cup M(S^*)$. We call $M(S)$ the *inner* medial axis and $M(S^*)$ the *outer* medial axis of Σ . Thus, M consists of all points in \mathbb{R}^3 that have at least two closest points in Σ .

Feature Size and Surface Samples. By definition, every point of \mathbb{R}^3 not in M has a unique closest point in Σ . For any point $x \in \mathbb{R}^3 \setminus (\Sigma \cup M)$ we denote by \hat{x} the unique closest surface point to x , i.e., $\hat{x} = \operatorname{argmin}_{y \in \Sigma} \|x - y\|$, and by $\tilde{x} \in M$ we denote the center of the medial ball tangent to Σ at \hat{x} and at the same side of Σ as x . The *medial feature size* is the function $\mu : \mathbb{R}^3 \setminus (\Sigma \cup M) \rightarrow \mathbb{R} \cup \{\infty\}$ defined as $\mu(x) = \|\hat{x} - \tilde{x}\|$. The function $f : \Sigma \rightarrow \mathbb{R}, x \mapsto \inf_{y \in M} \|x - y\|$ which assigns to each point in Σ its distance to the medial axis M , is called the *local feature size*. Notice that for $x \in \mathbb{R}^3 \setminus (\Sigma \cup M)$ it always holds that $f(\hat{x}) \leq \mu(x)$. It can also be easily seen that f is 1-Lipschitz. Throughout this paper, we assume that every point $x \in \Sigma$ has non-zero local feature size and that the infimum of the local feature size function over Σ , $f_0 = \inf_{x \in \Sigma} f(x)$, sometimes called the *reach* of Σ in the literature, is bounded away from zero.

For a constant $\varepsilon > 0$, a finite sample $P \subset \Sigma$ is called a (relative) ε -sample if for all $x \in \Sigma$, there exists a sample point $p \in P$ withing distance $\varepsilon f(x)$ from x .

Distance Functions and Derived Concepts. Given an open set, S we define the *distance function* h_S induced by S as

$$h_S(x) : S \rightarrow \mathbb{R}, \quad x \mapsto d(x, S^c).$$

When S is a shape enclosed by a surface Σ , $h_S(x) = d(x, \Sigma)$ in which we can expand the domain of the function to the entire \mathbb{R}^3 and denote it by h_Σ . Such distance functions have been broadly studied in the literature (see for example [Gro93]) and are known to carry a great deal of information about the set inducing them and its embedding. For example, Σ itself is simply given by $h_\Sigma^{-1}(0)$ and the medial axis M of Σ turns out to consist of all points at which h_Σ is not differentiable. The prevalence of information encoded in h_Σ motivates the use of distance functions induced by discrete samples of Σ , as computationally manageable “approximations” to it, for extraction of the same kind of information.

The function h_S (also h_Σ and more generally any distance function) is 1-Lipschitz and therefore continuous. However, it is not differentiable everywhere in S , i.e., the gradient can be undefined at certain points. As mentioned above, these points constitute $M(S)$. Some of these points are called *critical points*, namely, the points $x \in S$ that are contained in the convex hull of $A_S(x)$. Every other point of S is a *regular* point. Although the gradient is undefined at many of the regular points in S , a unique vector of steepest ascent of h_S exists at every such point (see [Lie04]). Let $d_S(x)$ be the center of the smallest closed ball containing $A_S(x)$ and let $r_S(x)$ be the radius of this ball. We define the *flow vector* at x (with respect to S) as

$$v_S(x) = \frac{x - d_S(x)}{h_S(x)}.$$

The vector $v_S(x)$ agrees with $\nabla h_S(x)$ at every point $x \in S$ for which this gradient is defined and extends the gradient of h_S everywhere else by determining the direction of steepest ascent for h_S [Lie04]. The points x for which $v_S(x) = 0$ are exactly the critical points of h_S . It can be easily verified using Pythagoras's theorem that

$$\|v_S(x)\|^2 = 1 - \left(\frac{r_S(x)}{h_S(x)}\right)^2. \quad (1)$$

By showing the convergence of Euler schemes, Lieutier [Lie04] proved that for any *bounded open* set S , the flow vector field v_S as described above can be integrated, to give a map $\phi_S : \mathbb{R}^+ \times S \rightarrow S$ that is continuous in *both* variables and satisfies for all $x \in S$, (1) $\phi_S(0, x) = x$ and (2) $\phi_S(s, \phi_S(t, x)) = \phi_S(s+t, x)$ for all $s, t \in \mathbb{R}^+$. Starting at any point $x \in S$, the flow map $t \mapsto \phi_S(t, x)$ defines a continuous path in S , called the *orbit* of x and denoted $\phi_S(x)$, when t goes from 0 to $+\infty$. More formally, $\phi_S(x) = \{y \mid \exists t \in \mathbb{R}^+ : \phi_S(t, x) = y\}$. At any point $y \in \phi_S(x)$ the flow vector $v_S(y)$ determines the direction of the flow at y . Lieutier further proved that when S is a bounded open set then for any point $x \in S$ both of the maps $t \mapsto h_S(\phi_S(t, x))$ and $t \mapsto r_S(\phi_S(t, x))$ are increasing and that the former map is continuous and satisfies the following integral equation:

$$h_S(\phi_S(t, x)) = h_S(x) + \int_0^t \|v_S(\phi_S(\tau, x))\|^2 d\tau. \quad (2)$$

It turns out that the critical points of h_S are the fixed points of the flow map, i.e., if c is a critical point of h_S , then $\phi_S(t, c) = c$ for all $t \in \mathbb{R}^+$.

If the shape S is bounded, then the flow orbit $\phi_S(x)$ stays inside S and converges to a critical point c of h_S as $t \rightarrow +\infty$. Notice that we consider c to be also in the flow orbit of x . For a critical point c of the flow ϕ_S , the set of all points x whose flow orbits converges to c is called the *stable manifold* of c and is denoted by $\mathcal{S}(c)$. In other words,

$$\mathcal{S}(c) = \{x \mid \phi_S(+\infty, x) = c\}.$$

Although there is no flow out of a critical point c of h_S , it is interesting to know where the points very close to c flow. Some of these points flow into c while other flow away from it. We define the *unstable manifold* $\mathcal{U}(c)$ of a critical point c , as the set of all points into which points arbitrarily close to c flow. Formally,

$$\mathcal{U}(c) = \bigcap_{\varepsilon > 0} \bigcup_{y \in B(c, \varepsilon)} \phi_S(y).$$

With an abuse of terminology, we say that c “flows” into the points of $\mathcal{U}(c)$.

As motivated above, given a sample P of $\Sigma = \partial S$, it is natural to try to approximate the distance function h_S by the function $h_P : \mathbb{R}^3 \rightarrow \mathbb{R}, x \mapsto \min_{p \in P} \|x - p\|$, that assigns to each point its distance to the sample P of ∂S . The finite set P is the boundary of the open set $P^c = \mathbb{R}^3 \setminus P$ and we can define the distance function and the flow vector for P^c as we did for S . Since $h_{P^c} = h_P$, in a slight abuse of notation we will denote both of these distance function with h_P and let $A_P(x) = \{y \in P \mid \|x - y\| = h_P(x)\}$. We shall also denote the center of the smallest ball containing $A_P(x)$ as $d_P(x)$ and the associated flow vector field as v_P . Note that $d_P(x)$ is the closest point on the Delaunay face dual to the lowest dimensional Voronoi face that contains x . Sometimes $d_P(x)$ is referred to as the driver of x . Technically speaking, since P^c is unbounded, h_P has a critical point at infinity. The other critical points of h_P can be characterized as the intersection points of Delaunay faces with their dual Voronoi faces [GJ03].

Separation of Critical Points. Dey, et al. [DGRS05] observed that if P is an ε -sample of the smooth boundary Σ of a shape S , then the critical points of the discrete distance function h_P cannot reside everywhere in S . Rather they have to be either very close to Σ or very close to M . Thus critical points of h_P can be classified based on whether they are close to Σ or close to M . We refer to the first class of critical points as *surface critical points* and the second class as *medial axis critical points*. We further subdivide the medial axis critical points h_P into two subgroups: *inner* medial axis critical points are those that are close to $M(S)$ and *outer* medial axis critical points are those close to $M(S^*)$.

Core. The union of the unstable manifolds of the inner medial axis critical points of h_P will play an important role in the present paper we refer to this union as the *core* the approximation. We will show that for a sufficiently dense sample P of Σ this core is homotopy equivalent to the medial axis $M(S)$ of S .

Unstable Flow Complex. In general, unstable manifolds of many critical points may intersect. However, in discrete settings, since there are only a finite number of critical points, we can achieve a cell complex decomposition of space by grouping together, as cells, the points of the space that are flown into from the exactly same “set” of critical points. We define a relation “ \sim ” on the pairs of points in \mathbb{R}^3 under which $x \sim y$ if and only if the set of critical points that flow into x coincides with the set of those that flow into y , or equivalently, if x is in the unstable manifolds of the same set of critical points as y . It is clear that “ \sim ” is an equivalence relation. The *unstable flow complex* induced by a point set P , denoted $\mathbf{U}(P)$ (or just \mathbf{U} when P is understood) is the cell complex whose cells are connected components of the subdivision of space into equivalence classes of the “ \sim ” relation. In appendix B, we study the structure of this complex more closely in three dimensions. It turns out that the full-dimensional cells of this complex coincide with full-dimensional cells of $\text{Vor}(P)$. The lower dimensional cells of \mathbf{U} introduce a subdivision of the 2-skeleton of $\text{Vor}(P)$. We will see that the core is a sub-complex of \mathbf{U} . The importance of these observations is primarily in practice where we desire the output of our algorithms to have a geometric cell-complex structure with no redundancies. Nevertheless, the unstable flow complex of a given set of points can be of independent interest theoretically.

3 Homotopy Equivalences

As in the previous section, we always assume that Σ is a smooth manifold with associated inner and outer components S and S^* . Furthermore, we assume that P is ε -sampling of Σ . Let \mathcal{C} be the *core* of the approximation as defined in Section 2, i.e. \mathcal{C} is the union of unstable manifolds of the inner medial axis critical points of the distance function h_P contained in S . Here we want to show that \mathcal{C} and the medial axis $M(S)$ of S are homotopy equivalent.

Following Lieutier [Lie04] the following criterion is used throughout this paper to prove homotopy equivalence between topological spaces. For the classical definition of homotopy equivalence refer, for example, to [Hat01].

Proposition 1 *Let X and $Y \subseteq X$ be arbitrary sets and let $H : [0, 1] \times X \rightarrow X$ be a continuous function on both variables satisfying the following three conditions. (1) $\forall x \in X, H(0, x) = x$, (2) $\forall x \in X, H(1, x) \in Y$, and (3) $\forall y \in Y, \forall t \in [0, 1], H(y, t) \in Y$. Then X and Y have the same homotopy type.*

Intuitively, if we interpret the first argument of the map H as time. Using a simple re-parametrization in the first argument, we can replace the interval $[0, 1]$ with any interval $[0, T]$ where $T > 0$ is a real number. It is important that the time interval considered has finite length. The above criterion for homotopy equivalence between X and Y continuously maps points in X to those in Y during the time interval $[0, T]$. At time 0, all points in X are mapped to themselves and at time T , they have all arrived in Y . Notice, the important property that the points in Y stay in Y all the time.

In the following we want to plug in ϕ_S and ϕ_P (the flow resulted from integrating v_P after circumventing the technical difficulty of unboundedness of P^c) for the map H mentioned above. When distance flow maps are used for H , the first condition of Proposition 1 is automatically satisfied since $\phi_S(0, x) = x$ for all x (the same holds for ϕ_P). Satisfying the second condition when using flows corresponds to proving that every point in X flows into Y in *finite* time. Finally, the third condition of Proposition 1 to requiring the set Y to be *closed* for the used flow, meaning that no point of Y flows out of it.

Crucial to the provided homotopy equivalence proofs is the concept of reduced shapes, that we introduce at first.

3.1 Reduced Shapes

Reduced Shapes. Let S be a shape whose boundary is a smooth 2-manifold Σ and let $0 < \delta < 1$. The δ -*tubular neighborhood* Σ_δ of Σ is the set

$$\Sigma_\delta = \Sigma \cup \{x \in \mathbb{R}^3 \setminus (\Sigma \cup M) : \|x - \hat{x}\| < \delta f(\hat{x})\}.$$

and the δ -*reduced shape* S_δ is defined as $S_\delta = S \setminus \Sigma_\delta$. Notice that the definition of Σ_δ , puts the medial axis $M(S)$ of S into S_δ . The following lemma shows that every point of $M(S)$ is in fact an interior point of S_δ . The proof is provided in Appendix A.

Lemma 1 *For every $0 < \delta < 1$, every point of $M(S)$ is an interior point of S_δ .*

Note, that the previous Lemma also implies that every boundary point of S_δ has a unique closest point in Σ . The following Lemma gives a complete characterization of the boundary points of S_δ . The proof is provided in Appendix A.

Lemma 2 *The boundary of S_δ consists of exactly the points $x \in S \setminus M(S)$ satisfying $\|x - \hat{x}\| = \delta f(\hat{x})$.*

3.2 Homotopy Proofs

Our proof that the core \mathcal{C} and the medial axis $M(S)$ of S are homotopy equivalent consists of the following two steps.

- (1) S and S_δ are homotopy equivalent for $0 < \delta < 1$.
- (2) $S_{2\varepsilon^2}$ and \mathcal{C} are homotopy equivalent for $\varepsilon < 0.14$.

The equivalences (1) and (2) together with the homotopy equivalence of S and $M(S)$ [Lie04] establish the homotopy equivalence of $M(S)$ and \mathcal{C} .

As we mentioned earlier an essential part of each of the homotopy equivalence proofs consists of showing, in correspondence to the second condition of Proposition 1, that a considered object is closed under some flow. Trivially, every open set S is closed with respect to the flow ϕ_S it induces. The medial axis $M(S)$ of an open set S is also closed under ϕ_S . This is because $M(S)$ precisely consists of the points $x \in S$ for which $A_S(x) > 1$ and consequently $r_S(x) > 0$. This observation, along with a result of Lieutier [Lie04] that states the map $t \mapsto r_S(\phi_S(t, x))$ is increasing for every $x \in S$, implies that the flow out of every point $x \in M(S)$ stays inside $M(S)$.

Proposition 2 *The medial axis $M(S)$ of any open set S is closed under the flow ϕ_S .*

The following Lemma, due to Lieutier [Lie04], shows that every point $x \in S$ arrives in $M(S)$ under ϕ_S in finite time. In Appendix A, we reproduce the proof of the Lemma for the sake of completeness and because it is a good example for demonstrating the technique we employ in homotopy equivalence proofs in the sequel.

Lemma 3 [Lie04] *Let Δ be an upper bound for the diameter of an open set S . Then for every $x \in S$, $\phi_S(\Delta, x) \in M(S)$.*

The above Lemma, Proposition 2, and the fact that $\phi_S(0, x) = x$ for all $x \in S$, fulfill the three requirements of Proposition 1 and therefore imply the following result of Lieutier [Lie04].

Corollary 1 [Lie04] *Every bounded open subset S of \mathbb{R}^3 is homotopy equivalent to its medial axis $M(S)$.*

Next we show that the reduced shapes S_δ for $0 < \delta < 1$ are also closed under the flow ϕ_S . By Proposition 1 this implies that S and S_δ are homotopy equivalent for $0 < \delta < 1$. The proof is provided in Appendix A.

Lemma 4 *For any $0 < \delta < 1$, the reduced shape S_δ is closed under the flow ϕ_S .*

Corollary 2 *For any $0 < \delta < 1$, any open set S is homotopy equivalent to its δ -reduced shape S_δ .*

In order to complete our proof, we need to integrate the vector field v_P induced by the same P . However, we face a technical difficulty: the result of Lieutier [Lie04] only applies to *bounded* open sets. Although P^c is open, it is not bounded. This problem can be bypassed using a somewhat standard technique. Let $B_0 = B(c, \Delta/2)$ to be an open ball of center c and diameter Δ such that $S \subset B_0$ and thus Δ is an upper bound for diameter of S and let $B_1 = 3B_0 = B(c, 3\Delta/2)$ be the result of scaling B_0 three times around its center. Finally let Q be the open set obtained by removing the sample points in P from B_1 , i.e., $Q = B_1 \setminus P$. Notice that $\partial Q = P \cup \partial B_1$. For any point x in S , the closest boundary point of Q to x is a point of P . This is because every boundary point of B_1 is at distance strictly greater than Δ or more from x while there is always a point of P within distance strictly less than Δ from x . This in particular means that the flow induced by Q in S , entirely depends on the arrangement of the points in P only and is completely irrelevant

to the boundary of B_1 . In other words, the flow lines of the flow ϕ_Q agree with those of the discrete flow ϕ_P inside S as described by [GJ03]. As such, in the sequel we shall use the notation ϕ_P to denote the *continuous* flow induced by the open set Q inside B_0 . In order to prove an analogue of Lemma 4 for the flow ϕ_P , we first prove following two auxiliary lemmas. The proofs of these two Lemmas are primarily based on the Lipschitzness of the local feature size function. These proofs are also moved to Appendix A.

Lemma 5 *Let x be a point on the boundary of S_δ . Any vector v at x that makes an angle α less than $\arctan\left(\frac{1-\delta}{2\delta}\right)$ with the vector $\tilde{x} - x$ points into S_δ .*

Lemma 6 *Let x be a point on the boundary of S_δ . The angle α that $v_P(x)$ makes with $\tilde{x} - x$ is bounded by*

$$\arccos\left(\frac{2\delta(1-\varepsilon-\delta)-\varepsilon^2}{2(1-\delta)(\delta+\varepsilon)}\right).$$

Now we are ready to prove the analogue of Lemma 4 for the flow ϕ_P .

Lemma 7 *The δ -reduced shape S_δ is closed under the flow ϕ_P , respectively, for any $\varepsilon^2 \leq \delta \leq 10\varepsilon^2$ if $\varepsilon \leq 0.14$.*

Proof. Lemmas 5 and 6 imply that for any point x on the boundary of S_δ the vector $v_P(x)$ points into S_δ if $\delta \geq \varepsilon^2$ and $\varepsilon \leq 0.14$ as can be checked by plugging in the values into the bounds provided by the two lemmas. ■

Lemma 7 now allows us to prove that the core \mathcal{C} and S_δ are homotopy equivalent. To do so we need the following corollary and lemma.

Corollary 3 *For any $1 > \delta \geq \varepsilon^2$ and $\varepsilon \leq 0.14$, every point of the reduced shape S_δ flows under ϕ_P into a inner medial axis critical point of h_P .*

Proof. By Lemma 7 every point of S_δ stays inside S_δ under ϕ_P while all surface critical points of h_P are in Σ_δ . Therefore, flow orbits $\phi_P(x)$ for points $x \in S_\delta$ have to flow into medial axis critical points of h_P . ■

Lemma 8 *For $\varepsilon \leq 0.14$, if P is an ε -sample of Σ , then there is a constant T such that $\phi_P(T, x) \in \mathcal{C}$ for all $x \in S_{2\varepsilon^2}$.*

Proof. Let $\zeta > 0$ be the minimum distance between a Delaunay face D induced by P and its dual Voronoi face V that do not intersect. Consider the reduced set S_δ for $\delta = 2\varepsilon^2$. By Corollary 3, S_δ is closed under the flow ϕ_P . Therefore, every flow line of ϕ_P in S_δ ends in some inner medial axis critical point of h_P in the limit. Consider now any point $x \in S_\delta \setminus \mathcal{C}$ and let $V(x)$ be the cell in the Voronoi complex $\text{Vor}(P)$ that contains x and let $D(x)$ be its dual Delaunay cell in the Delaunay complex $\text{Del}(P)$. There are two cases two consider depending on whether $V(x)$ and $D(x)$ intersect.

- (i) $V(x) \cap D(x) = \emptyset$. In this case, the distance $x - d_P(x) \geq \zeta$ and since $h_P(x) \leq \Delta$, we get $\|v_P(x)\| = \|x - d_P(x)\|/h_P(x) \geq \zeta/\Delta$.
- (ii) $V(x) \cap D(x) = \{c\}$, where c is a critical point of d_P . It can be easily observed that if c is a medial axis critical point, then V is entirely contained in the unstable manifold $\mathcal{U}(c)$ of c and this implies that $x \in \mathcal{C}$, contradicting our choice of x . Therefore, c must be a surface critical point and as such $c \in \Sigma_{\varepsilon^2}$ while $x \in S_{2\varepsilon^2}$. With a similar argument as the one used in Lemma 2 the open ball of center c and radius $\frac{1-\delta}{1+\delta}(2\varepsilon^2 - \varepsilon^2)f(\hat{c})$ is entirely contained in $\Sigma_{2\varepsilon^2}$ and therefore $\|x - c\| \geq \xi = \frac{1-\delta}{1+\delta}\varepsilon^2 f_0$. Thus we get $\|v_P(x)\| = \|x - c\|/h_P(x) \geq \xi/\Delta$.

Thus for every point $x \in S_\delta \setminus \mathcal{C}$, $\|v_P(x)\| \geq \vartheta = \min\{\xi/\Delta, \zeta/\Delta\}$. If $\phi_P(t, x) \notin \mathcal{C}$ for all $t \in [0, \Delta/\vartheta^2]$ we get from Eq. (2)

$$h_P(\phi_P(\Delta/\vartheta^2), x) = h_P(x) + \int_0^{\Delta/\vartheta^2} \|v_P(\phi_P(\tau, x))\|^2 d\tau \geq h_P(x) + \int_0^{\Delta/\vartheta^2} \vartheta^2 d\tau = h_P(x) + \Delta.$$

This contradicts the fact that Δ is an upper bound for the diameter of S . Thus if we set $T = \Delta/\vartheta^2$, then $\phi_P(T, x) \in \mathcal{C}$ for all $x \in S_\delta \setminus \mathcal{C}$. ■

Corollary 4 For $\varepsilon \leq 0.14$, if P is an ε -sample of Σ , then \mathcal{C} is homotopy equivalent to $S_{2\varepsilon^2}$.

Proof. By definition, \mathcal{C} is closed under the flow ϕ_P . This along with the result of Lemma 8 satisfy the requirements of Proposition 1, implying that \mathcal{C} and $S_{2\varepsilon^2}$ are homotopy equivalent. ■

Combining the result of the above Corollary with that of Corollary 2 gives us the theorem we set out to prove.

Theorem 1 The shape S , its medial axis $M(S)$, and the core \mathcal{C} consisting of the the union of unstable manifolds of inner medial axis critical points of h_P , are homotopy equivalent.

3.3 Extending the Core

As mentioned in the Introduction section, one of the most pleasing properties of the core is its flexibility in being used with other medial axis approximation algorithms. The following Theorem clarifies this statement.

Theorem 2 Let $W \subset S_{2\varepsilon^2}$ be any set of points and let $\widehat{W} = \{x \mid \exists y \in W : x \in \phi_P(w)\}$. Then $\mathcal{C} \cup \widehat{W}$ is homotopy equivalent to S .

Proof. By Lemma 8 all points in $\Sigma_{2\varepsilon^2}$ flow into \mathcal{C} in finite time. This includes all the points in \widehat{W} . On the other hand, by definition $\mathcal{C} \cup \widehat{W}$ is closed for ϕ_P . These are all the requirements needed to invoke Proposition 2 to establish $\mathcal{C} \cup \widehat{W}$ is homotopy equivalent to $S_{2\varepsilon^2}$ which is itself homotopy equivalent to S . ■

Computing the flow closure of all the points in W can be computationally difficult, depending on the nature of W . If W is a sub-complex of the Voronoi complex $\text{Vor}(P)$ however, the flow closure of all the points in W can be done in bulk by computing the flow closures of a whole face at a time as is explained in section B. This is for example the case in the algorithm of Dey and Zhao [DZ03].

4 Geometric Approximation

Geometric fidelity of the core to the real medial axis $M(S)$ of S can be a concern. Although the separation of critical points [DGRS05] ensures that the medial axis critical points lie very close the medial axis, it provides no guarantee for the paths connecting them on their unstable manifolds to enjoy the same closeness. The same concern is valid when we extend the core with a set of points close the medial axis: to guarantee the topology one must include the flow closures of added points but it is not clear that this closure stays close to the medial axis as well.

For a point x in the core \mathcal{C} or any other approximation of the medial axis, the *relative* approximation error at x can be considered to be the ratio between the distance from x to $M(S)$ and $h_P(x)$. One would like to have that the distance to $M(S)$ along a flow line would grow at most linearly with $h_P(x)$ as this would correspond to maintaining a constant relative error all along the path. However, proving this, if it is true at all, appears elusive. Nevertheless, in this section we show that if we start from a point x close to the medial axis and follow the flow line $\phi_P(x)$, the distance to $M(S)$ along this path grows as a function of h_P at a rate slightly super-linear at worst. More precisely we will show that if we scale the space so that $h_P(x)$ has unit length, then if x has a medial axis point within distance $O(\sqrt{\varepsilon})$, then every point y in the flow line starting at x will have a medial axis point within distance $O(\sqrt{\varepsilon})h_P(y)^{1+O(\sqrt{\varepsilon})}$. The proof of all technical Lemmas are moved to Appendix A due to the lack of space.

Lemma 9 For every point $x \in S$ and for every $p \in A_P(x)$, $\|x - p\|^2 \leq h_S(x)^2 + \varepsilon^2 f(\hat{x})^2 + \varepsilon^2 h_S(x) f(\hat{x})$.

The *driving angle* of a point x with respect to the sample P is defined as $\theta_P(x) := \angle(-v_P(x), x - y)$ for any $y \in A_P(x)$. A critical point has driving angle equal to $\pi/2$ while points not on the 2-skeleton of $\text{Vor}(P)$ have driving angles equal to 0. The following Lemma generalizes the separation of critical points (see [DGRS05]). The proof is rather similar to the one given in [DGRS05] as is provided in Appendix A.

Lemma 10 Let x be a point with driving angle $0 < \theta \leq \pi/2$. Then, x is within distance $\frac{2\varepsilon}{1-\cos\theta}\mu(x)$ from one of \hat{x} or \check{x} .

Recall that the driver $d_P(x)$ is the same for all points x on the same face of the the Voronoi complex $\text{Vor}(P)$. In fact the flow path $\phi_P(x)$ turns exactly when the flow moves from one Voronoi face to another. Consider a point $x \in S$ that lies on the 2-skeleton of $\text{Vor}(P)$, i.e. $|A_P(x)| \geq 2$, and consider a line segment L in the flow orbit $\phi_P(x)$. The distance $h_P(x)$ monotonically increases along a flow path. We can therefore parametrize this line segment using the distance to the sample set P . Let f be a non-decreasing real valued C^1 function and consider for each point $x \in L$, the ball $B_f(x) = B(x, f(h_P(x)))$. We define the set $D(x)$ as those points on the boundary of $B_f(x)$ that are left outside $B_f(x)$ when x moves infinitesimally in the direction of $\phi_P(x)$. In other words if we take $\eta > 0$ such that $\phi_P(\eta, x)$ has the same driver as $d_P(x)$ (meaning that $\phi_P(\eta, x)$ is on the same line segment of $\phi_P(x)$ as x), then

$$D(x) = \bigcap_{0 < \varepsilon < \eta} (\partial B_f(x) \setminus \overline{B}_f(\phi_P(\varepsilon, x))).$$

Lemma 11 $D(x)$ consists of those points $y \in \partial B_f(x)$ that satisfy $\angle(y - x, d_P(x) - x) \leq \psi$ where $\cos \psi = df/dh_P$ at x .

Lemma 12 Let $B = B(x, R)$ be a ball empty of sample points with at least one sample point on its boundary and containing at least one medial axis point. Then the ball $B(x, (1 - 4\varepsilon^2)R)$ does not intersect Σ .

Theorem 3 Let $x_0 \in S$ be a point with $|A_P(x_0)| \geq 2$ and let $x_1 = \phi_P(t_1, x_0)$ be such that for all $0 \leq t \leq t_1$, $\cos(\theta_P(\phi_P(t, x_0))) \geq c$. Then, if there is a medial axis point within distance f_0 from x_0 , there is a medial axis point within distance $f(x) = f_0(h_P(x_1)/h_P(x_0))^\xi$ from x' , provided that $\xi \geq \frac{1}{c}(1 + 4\varepsilon^2 h_0^2/f_0^2)$.

Proof. We prove the theorem by showing that there always remains a medial axis point within distance $f_0(h_P(x)/h_P(x_0))^\xi$ as x moves along the flow line. We do this by showing that this proposition is maintained when x moves infinitesimally along the flow line. To this end, we first recall that the ball $B_f(x)$ is an open ball by definition. If $B_f(x)$ contains a medial axis point z , then for any direction vector v , there is a small enough real number $\tau > 0$, such that the translated ball $B_f(x) + \tau v = \{y + \tau v \mid y \in B_f(x)\}$ contains z as well. Since f is increasing, this implies that $z \in B_f(x + \tau v)$. In particular this implies by choosing $v = v_P(x)$ that if $B_f(x)$ contains a medial axis point, so does $B_f(x + \tau v_P(x))$ for τ sufficiently small.

We thus only need to consider the case where $B_f(x)$ contains no medial axis point while its boundary does. Thus, let $z \in \partial B_f(x)$ be a medial axis point. By Lemma 11, if $\cos(\angle(x - z, -v_P(x))) \leq \frac{d}{dh_P} f(h_P(x))$, then z is contained in $B_f(x + \tau v_P(x))$ for a sufficiently small $\tau > 0$ and therefore we have nothing to prove. Thus we only need to consider the case when

$$\cos(\angle(x - z, -v_P(x))) > \frac{d}{dh_P} (f(h_P(x))) = \frac{\xi f_0}{h_P(x_0)} \left(\frac{h_P(x)}{h_P(x_0)} \right)^{\xi-1}.$$

We denote this maximum angle by ψ . We show that in this case, the flow vector $v(z)$ points to the interior of the ball $B_f(x)$. This implies that the flow $\phi(z)$ enters $B_f(x)$ and therefore by what we showed above, $B_f(x)$ must contain a medial axis, contradicting our choice of z .

The ball $B(x, h_P(x))$ contains no sample points but includes z , a medial axis point, and therefore by Lemma 12, the ball $B_0 = B(x, (1 - 4\varepsilon^2)h_P(x))$ does not intersect Σ . Consider the plane Π tangent to $\partial B_f(x)$ at z . This plane, intersects the ball B_0 in a disk of radius $R_0 \geq \sqrt{((1 - 4\varepsilon^2)h_P(x))^2 - f(h_P(x))^2}$. We will show that there are surface points at distance less than R_0 from z , i.e. $h_S(z) < R_0$. It can be easily observed that any such surface point must lie on the side of Π opposite to the one containing $B_f(x)$ and therefore $\text{conv}(A(x))$ resides on the side of Π opposite to $B_f(x)$. Since $d(x) \in \text{conv}(A(x))$, this implies that $v(x)$ points to the interior of $B_f(x)$, as desired. So, all left to show is that $h_S(z) < R_0$. To prove this, we show that for at least one of the points $y \in A_P(x)$, $\|z - y\| < R_0$. Since $|A_P(x)| \geq 2$, every plane containing the points x and $d_P(x)$ that does not intersect $A_P(x)$ must contain at least one point of $A_P(x)$ on each side. The maximum distance between z and $A_P(x)$ can therefore occur when $|A_P(x)| = 2$ and z is on the bisector plane of the segment connecting the two points in $A_P(x)$. To find this maximum distance we do a change of coordinates, putting the origin at x and the driver $d(x)$ on the x -axis, the two points p_1, p_2 in $A_P(x)$ on the x, y -plane, and z on the x, z -plane. We can calculate the coordinates of p_1, p_2 , and z as follows:

$$p_1 = (h \cos \theta, h \sin \theta, 0), \quad p_2 = (h \cos \theta, -h \sin \theta, 0), \quad z = (f(h) \cos \psi, 0, f(h) \sin \psi),$$

where by θ we denote the the driving angle $\theta_P(x)$ and by h we denote $h_P(x)$. Thus we get for the distance between z and p_1 (same distance between z and p_2):

$$\begin{aligned}\|z - p_1\|^2 &= (h_P(x) \cos \theta - f(h_P(x)) \cos \psi)^2 + (h_P(x) \sin \theta)^2 + (f(h_P(x)) \sin \psi)^2 \\ &= h_P(x)^2 + f(h_P(x))^2 - 2h_P(x)f(h_P(x)) \cos \theta \cos \psi\end{aligned}$$

Denoting $h_P(x_0)$ by h_0 and $h_P(x)$ by h and using the lower bounds for $\cos \psi$ and $\cos \theta$ we get:

$$\|z - p_1\|^2 < h^2 + f_0^2 \left(\frac{h}{h_0}\right)^{2\xi} - 2chf_0 \left(\frac{h}{h_0}\right)^\xi \frac{\xi f_0}{h_0} \left(\frac{h}{h_0}\right)^{\xi-1} = h^2 + (1 - 2c\xi)f_0^2 \left(\frac{h}{h_0}\right)^{2\xi}$$

Thus in order for $\|z - p_1\| < R_0$, it suffices to have

$$h^2 + (1 - 2c\xi)f_0^2 \left(\frac{h}{h_0}\right)^{2\xi} \leq (1 - 4\varepsilon^2)h^2 - f_0^2 \left(\frac{h}{h_0}\right)^{2\xi}.$$

Since $(1 - 4\varepsilon^2)^2 > 1 - 8\varepsilon^2$, the above inequality is satisfied when the following one is:

$$4\varepsilon^2 h^2 \leq (c\xi - 1)f_0^2 \left(\frac{h}{h_0}\right)^{2\xi} \quad \text{or identically} \quad \frac{4\varepsilon^2}{c\xi - 1} \leq \frac{f_0^2}{h_0^2} \left(\frac{h}{h_0}\right)^{2\xi-2}.$$

Since $h/h_0 \geq 1$, the above inequality holds if $f_0 \geq \left(\frac{2\varepsilon}{\sqrt{c\xi-1}}\right) h_0$. This is guaranteed by the bound on ξ prescribed in the statement of the Theorem. ■

Corollary 5 *Let x be a point in the 2-skeleton of $\text{Vor}(P)$ such that $\|x - \check{x}\| \leq \frac{2\sqrt{\varepsilon}}{1-2\sqrt{\varepsilon}} h_P(x)$, then for every point y on the flow bath $\phi_P(x)$, there is a medial axis point within distance*

$$\frac{2\sqrt{\varepsilon}}{1-2\sqrt{\varepsilon}} h_P(x) \left(\frac{h_P(y)}{h_P(x)}\right)^\xi$$

from y , where $\xi = 1 + O(\sqrt{\varepsilon})$.

The proof of the above Corollary is also given in Appendix A. An immediate consequence of this Corollary is that the core and the flow closures converge to being contained in the medial axis as $\varepsilon \rightarrow 0$. As a result, when the core is used and extended using the filtering conditions of [DZ03], the computed approximate medial axis converges to the true medial axis in the limit.

5 Experiments and Concluding Remarks

We introduced the notion of the “core” as the union of unstable manifolds of medial axis critical points of the distance to and ε -sampling of a surface and established its homotopy equivalence to the sampled shape. We further showed how this core can be safely augmented to provide topological guarantee for any algorithm that geometrically approximates (a subset of) the medial axis. Finally we showed that the computed core and related flow closures do converge to subsets of the medial axis when the input sample grows infinitely dense. A result of the implementation of the unstable flow complex data structure, the core, and the extension of the core using conditions similar to those in [DZ03] is shown in Figure 1. As can be observed from the picture, the extension of the core has filled the holes in the junctions of the geometric approximation of the medial axis of the 3-holes model computed by filtering Voronoi facets by a condition similar to the angle condition in [DZ03].

Acknowledgments. The authors would like to thank Bálint Miklós and Jingyi Jin for the implementation and the presented output pictures.




Figure 1: Left: The core computed for the 3-holes model. The red lines are either unstable manifolds of index-2 saddle points, or the one dimensional parts (hairs) of index-1 saddle points. Middle: Filtered Voronoi facets based on a condition similar to the angle condition in [DZ03]. Right: Extended core: the core, plus the flow closures of the facets in the middle picture.

References

- [AB99] Nina Amenta and Marshall W. Bern. Surface reconstruction by voronoi filtering. *Discrete & Computational Geometry*, 22:481–504, 1999.
- [ABE98] Nina Amenta, Marshall W. Bern, and David Eppstein. The crust and the β -skeleton: Combinatorial curve reconstruction. *Graphical Models and Image Processing*, 60(2):125–135, 1998.
- [ABE04] Dominique Attali, Jean-Daniel Boissonnat, and Herbert Edelsbrunner. Stability and computation of the medial axis — a state-of-the-art report. *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*, 2004.
- [ACK01] Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. The power crust, unions of balls, and the medial axis transform. *Computational Geometry*, 19(2-3):127–153, 2001.
- [BC02] Jean-Daniel Boissonnat and Frédéric Cazals. Smooth surface reconstruction via natural neighbour interpolation of distance functions. *Computational Geometry*, 22(1-3):185–203, 2002.
- [CKM04] Tim Culver, John Keyser, and Dinesh Manocha. Exact computation of the medial axis of a polyhedron. *Computer Aided Geometric Design*, 21(1):65–98, 2004.
- [CL05] Frédéric Chazal and André Lieutier. The λ -medial axis. *Graphical Models and Image Processing*, 67(4):304–331, 2005.
- [DGRS05] Tamal K. Dey, Joachim Giesen, Edgar A. Ramos, and Bardia Sadri. Critical points of the distance to an epsilon-sampling of a surface and flow-complex-based surface reconstruction. In *Symposium on Computational Geometry (SCG)*, pages 218–227, 2005.
- [DZ03] Tamal K. Dey and Wulue Zhao. Approximating the medial axis from the voronoi diagram with a convergence guarantee. *Algorithmica*, 38:179–200, 2003.
- [GJ03] Joachim Giesen and Matthias John. The flow complex: A data structure for geometric modeling. In *Symposium on Discrete Algorithms (SODA)*, pages 285–294, 2003.
- [Gro93] Karl Grove. Critical point theory for distance functions. *Symposia in Pure Mathematics*, 54(3):357–385, 1993.
- [Hat01] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2001.

[Lie04] André Lieutier. Any bounded open subset of \mathbb{R}^n has the same homotopy type as its medial axis. *Computer-Aided Design*, 36(11):1029–1046, 2004.

A Missing Proofs

Proof of Lemma 1. Take $x \in M(S)$ and let $B = B(x, \omega)$ where $\omega = \frac{1-\delta}{1+2\delta} h_S(x) > 0$. We will show that $B \subset S_\delta$ implying that x is an interior point of S_δ . Take any point $y \in B$. If $y \in M(S)$, then $y \in S_\delta$ by definition. Thus assume that $y \notin M(S)$ and therefore has a unique closest point \hat{y} in Σ . Notice that since $B' = B(x, h_S(x)) \subset S$ but $\partial B'$ intersects Σ , we get

$$h_S(x) - \omega \leq h_S(y) \leq h_S(x) + \omega. \quad (3)$$

Moreover since $x \in M(S)$, we get

$$f(\hat{y}) \leq h_S(y) + \|x - y\| \leq h_S(y) + \omega \leq 2\omega + h_S(x). \quad (4)$$

Combining Equations (3) and (4) and using the definition of ω we obtain

$$h_S(y) \geq h_S(x) - \omega = \delta h_S(x) + 2\delta\omega \geq \delta f(\hat{y}).$$

■

Proof of Lemma 2. Lemma 1 shows that no point of $M(S)$ can be a boundary point for S_δ . For every other point $x \in S$, there is a unique $\hat{x} \in \Sigma$. Consider the segment $\hat{x}\check{x}$. For every point in this segment the closest point in Σ is the same, i.e., \hat{x} . Thus, all points in the relative interior of the segment $\hat{x}\check{x}$ lie in Σ_δ and all points of the relative interior of $x\check{x}$ lie in S_δ . This implies that x is a boundary point for S_δ .

We now show that S_δ has no other boundary points. First we show that no point $x \in S_\delta \setminus M(S)$ with $h_S(x) = \lambda f(\hat{x})$ where $\lambda > \delta$ can be a boundary point of S_δ . To this end, we show that the open ball centered at x and with radius $\omega = \frac{1-\delta}{1+\delta}(\lambda - \delta)f(\hat{x})$ is entirely contained in S_δ . Consider the ball $B = B(c, f(\hat{x}))$ tangent to Σ at \hat{x} at the same side of Σ as x and with radius $f(\hat{x})$. By definition $B \subset S$ (both medial balls tangent to Σ at any point have radius larger than the local feature size of that point). Let y be a point at distance less than ω from x . The angle $\alpha = \angle xcy$ becomes larger when y moves further away from x in the direction of xy and it is maximized when $\angle cyx = \pi/2$ in which case $\sin \alpha \leq \omega / ((1 - \lambda)f(\hat{x}))$. The point \hat{y} can be at most as far away from y as \hat{x} is. Therefore, \hat{y} lies in the ball $B' = B(y, \|\hat{x} - y\|)$. Since B is an empty ball, this implies that the distance between \hat{y} and \hat{x} is at most $2f(\hat{x}) \sin \alpha = 2\omega / (1 - \delta)$ and from this we get

$$f(\hat{y}) \leq f(\hat{x}) + \|\hat{x} - \hat{y}\| \leq f(\hat{x}) + \frac{2\omega}{1 - \delta}.$$

On the other hand, $h_S(y) \geq h_S(x) - \omega$. Therefore,

$$h_S(y) \geq h_S(x) - \omega = \delta \left(f(\hat{x}) + \frac{2\omega}{1 - \delta} \right) \geq \delta f(\hat{y}).$$

Since y was taken arbitrarily from $B(x, \omega)$, it follows that $B(x, \omega) \subset S_\delta$, meaning x is an interior point of S_δ . Finally we show that no point of Σ_δ can be a boundary point of S_δ . Let $x \in \Sigma_\delta$ be a point for which $h_S(x) = \lambda f(\hat{x})$ where $\lambda < \delta$. Consider the ball of radius $\omega = \frac{1-\delta}{1+\delta}(\delta - \lambda)f(\hat{x})$ and let y be a point in this ball. Similar to the previous part, we can show that $\|\hat{y} - \hat{x}\| \leq 2\omega / (1 - \delta)$ and therefore $f(\hat{y}) \geq f(\hat{x}) - 2\omega / (1 - \delta)$. On the other hand $h_S(y) \leq h_S(x) + \omega$. Combining these we get

$$h_S(y) \leq h_S(x) + \omega = \delta \left(f(\hat{x}) - \frac{2\omega}{1 - \delta} \right) < \delta f(\hat{y}).$$

■




Figure 2: Left: Proof of Lemma 5; Right: Proof of Lemma 6.

Proof of Lemma 3. By Proposition 2, if $\phi_S(\Delta, x) \notin M(S)$, it must be that $\phi_S(t, x) \notin M(S)$ for all $t \in [0, \Delta]$. Since for every point x outside the medial axis $r_S(x) = 0$ and therefore by Eq. (1) $\|v_S(x)\| = 1$, if $\phi_S(\Delta, x) \notin M(S)$ by Eq. (2) we get

$$\begin{aligned} h_S(\phi_S(\Delta, x)) &= h_S(x) + \int_0^\Delta \|v_S(\phi_S(\tau, x))\|^2 d\tau \\ &\geq h_S(x) + \int_0^\Delta d\tau = h_S(x) + \Delta. \end{aligned}$$

However, $h_S(x) < \Delta$ for every point $x \in S$, a contradiction. \blacksquare

Proof of Lemma 4. By Lemma 2, the boundary points of S_δ are exactly those x satisfying $h_S(x) = \delta f(\hat{x})$. The flow vector $v_S(x)$ for any such point is exactly $(x - \hat{x})/h(x)$ which is a unit vector in the direction of normal to Σ at \hat{x} . In fact, the flow direction remains constant on the relative interior of the segment $\hat{x}\check{x}$. For every such point y , $\hat{y} = \hat{x}$ and therefore $f(\hat{y}) = f(\hat{x})$ remains the same while the distance to \hat{x} grows larger than $\delta f(\hat{x})$ as soon as the flow passes through x . Thus, the flow direction at every boundary point of S_δ points toward the interior of S_δ . \blacksquare

Proof of Lemma 5. Let c be the point on the line segment $\hat{x}\check{x}$ at distance $f(\hat{x})$ from \hat{x} . Without loss of generality we assume that $v = y - x$, where $y \in S$ is that close to x such that the inner angle θ of the triangle xcy at c is less than $\pi/2 - \alpha$. Note that by definition α is the inner angle of the triangle xcy at x , see also Figure 2 (Left). Let $B = B(c, f(\hat{x}))$ and let $B' = \overline{B}(y, \|\hat{x} - y\|)$. B does not contain any point from Σ and $\hat{y} \in B'$. Therefore $\hat{y} \in B' \setminus B$ and we have

$$\|\hat{x} - \hat{y}\| \leq 2f(\hat{x}) \sin \theta,$$

which in turn together with the fact that f is 1-Lipschitz implies that

$$f(\hat{y}) \leq f(\hat{x})(1 + 2 \sin \theta).$$

Let w be the intersection point of the boundary of B with the ray through $y - c$ and let x' be the projection of x on this ray. By construction $x' \in B$. Our assumption on θ implies that that on the ray through $c - y$

the point y comes before the point x' as seen from c . Putting these facts together we get

$$\begin{aligned}
\|y - \hat{y}\| &\geq \|y - w\| \\
&= \|y - x'\| + \|x' - w\| \\
&> \|y - x'\| + \|x - \hat{x}\| \\
&= \|y - x'\| + \delta f(\hat{x}) \\
&= \|x - x'\| \cot(\alpha + \theta) + \delta f(\hat{x}) \\
&= (1 - \delta)f(\hat{x}) \sin \theta \cot(\alpha + \theta) + \delta f(\hat{x}).
\end{aligned}$$

Thus,

$$\begin{aligned}
\frac{\|y - \hat{y}\|}{f(\hat{y})} &> \frac{(1 - \delta)f(\hat{x}) \sin \theta \cot(\alpha + \theta) + \delta f(\hat{x})}{f(\hat{x})(1 + 2 \sin \theta)} \\
&= \frac{(1 - \delta) \sin \theta \cot(\alpha + \theta) + \delta}{1 + 2 \sin \theta}.
\end{aligned}$$

In order for y to be in S_δ we want this fraction to be larger than δ . This amounts to $\cot(\alpha + \theta) > 2\delta/(1 - \delta)$ or equivalently $\tan(\alpha + \theta) < (1 - \delta)/2\delta$. Since we assumed $\tan \alpha < (1 - \delta)/2\delta$ we can find by the continuity of the tangent function $\theta_0 > 0$ such that $\tan(\alpha + \theta) < (1 - \delta)/(2\delta)$ for all $0 < \theta < \theta_0$. This implies the existence of $\lambda_0 > 0$ such that

$$(1 - \lambda)x + \lambda y = x + \lambda(y - x) = x + \lambda v \in S_\delta$$

for all $0 \leq \lambda < \lambda_0$, i.e., the vector v points into S_δ at x . ■

Proof of Lemma 6. Let c be the point on the line segment $\hat{x}\tilde{x}$ at distance $f(\hat{x})$ from \hat{x} . Let $B = B(c, f(\hat{x}))$ and let $B' = \overline{B}(x, (\delta + \varepsilon)f(\hat{x}))$. The driver $d_P(x)$ of x has to be contained in the convex hull of $B' \setminus B$. Let y be a point in the intersection of ∂B and $\partial B'$. Consider the triangle $cx y$, see also Figure 2 (Right), the inner angle of this triangle at x is at least $\pi - \alpha$. From the cosine theorem we get

$$\begin{aligned}
\cos(\pi - \alpha) &\leq \frac{(1 - \delta)^2 f(\hat{x})^2 + (\delta + \varepsilon)^2 f(\hat{x})^2 - f(\hat{x})^2}{2(1 - \delta)(\delta + \varepsilon)f(\hat{x})^2} \\
&= \frac{2\delta(\delta + \varepsilon - 1) + \varepsilon^2}{2(1 - \delta)(\delta + \varepsilon)}.
\end{aligned}$$

It follows

$$\cos \alpha \geq \frac{2\delta(1 - \delta - \varepsilon) - \varepsilon^2}{2(1 - \delta)(\delta + \varepsilon)},$$

which implies the statement of the lemma. ■

Proof of Lemma 9. Let $R = \mu(x)$. Let $B = B(\hat{x}, \varepsilon f(x))$. By the ε -sampling condition, the closest sample point to \hat{x} is at distance no more than $\varepsilon f(\hat{x}) \leq \varepsilon R$ from \hat{x} . In other words, B must contain at least one point from P .

Let $B_i = B(\tilde{x}, \mu^-(x))$, where μ^- is the medial feature size of the \hat{x} on the side of Σ opposite to x , and let B_o be the outer ball of radius $f(\hat{x})$, tangent to Σ at \hat{x} . Furthermore, let x_o be center of B_o . Since B_i and B_o cannot contain any sample points, the closest sample point to x must lie in the region $S = B \setminus (B_i \cup B_o)$. It is easy to see that the farthest points in S from \tilde{x} are those in $\partial B \cap \partial B_o$. Let y be one such point and consider the plane Π containing x , \hat{x} and y , and consequently \tilde{x} and x_o (See Figure 3). The intersection of B , B_i and B_o with Π are circles to which we shall refer in Figure 3 with the names of the corresponding balls. Let z and z' be the points where the bisector of $\angle \hat{x}x_o y$ meets the segment $\hat{x}y$ and the bisector of $\angle x_i x_o$, respectively.

The two triangles $x_o \hat{x} z'$ and $z \hat{x} z'$ are similar and therefore $\alpha = \angle z x_o \hat{x} = \angle z \hat{x} z'$. Furthermore, since $\|\hat{x} - y\| = \varepsilon f(\hat{x})$, $\sin \alpha = \varepsilon/2$. We have for the distance of x and y :




Figure 3: proof of Lemma 9

$$\begin{aligned}
\|x - y\|^2 &= \|x - \hat{x}\|^2 + \|\hat{x} - y\|^2 - 2\|\check{x} - \hat{x}\| \cdot \|\hat{x} - y\| \cos(\pi/2 + \alpha) \\
&= h_S(x)^2 + \varepsilon^2 f(\hat{x})^2 + 2\varepsilon h_S(x) f(\hat{x}) \sin \alpha \\
&= h_S(x)^2 + \varepsilon^2 f(\hat{x})^2 + \varepsilon^2 h_S(x) f(\hat{x}).
\end{aligned}$$

■

Proof of Lemma 10. Let $\lambda = \|x - \check{x}\|/\mu(x)$. The ball $B = B(\check{x}, \mu(x))$ contains no sample points. On the other hand, there is a sample point within distance $\varepsilon\mu(x)$ from \hat{x} . By triangle inequality, this implies that there is a sample point within distance $(1 - \lambda + \varepsilon)\mu(x)$ from x . Consider the ball $B' = B(x, (1 - \lambda + \varepsilon)\mu(x))$. All sample points in $A_P(x)$ are contained in $B' \setminus B$. This implies that the angle $\angle(-v(x), y - x)$ is at least θ for any $y \in \partial B' \cap \partial B$. Thus using the cosine rule on the triangle $\check{x}xy$ for any such y , we get the following inequality.

$$\mu(x)^2 \leq \lambda^2 \mu(x)^2 + (1 - \lambda + \varepsilon)^2 \mu(x)^2 + 2\lambda\mu(x)(1 - \lambda + \varepsilon)\mu(x) \cos \theta,$$

or equivalently

$$2(1 - \cos \theta)\lambda^2 - 2(1 + \varepsilon)(1 - \cos \theta)\lambda + \varepsilon(2 + \varepsilon) \geq 0.$$

Solving this inequality for λ we conclude that either,

$$\begin{aligned}
\lambda &\leq \frac{1 + \varepsilon}{2} \left(1 - \sqrt{1 - \frac{2\varepsilon(2 + \varepsilon)}{(1 - \cos \theta)(1 + \varepsilon)^2}} \right) \\
&\leq \frac{1 + \varepsilon}{2} \cdot \frac{2\varepsilon(2 + \varepsilon)}{(1 - \cos \theta)(1 + \varepsilon)^2} \\
&\leq \frac{2\varepsilon}{1 - \cos \theta},
\end{aligned}$$

or

$$\begin{aligned}
\lambda &\geq \frac{1+\varepsilon}{2} \left(1 + \sqrt{1 - \frac{2\varepsilon(2+\varepsilon)}{(1-\cos\theta)(1+\varepsilon)^2}} \right) \\
&\geq \frac{1+\varepsilon}{2} \left(2 - \frac{2\varepsilon(2+\varepsilon)}{(1-\cos\theta)(1+\varepsilon)^2} \right) \\
&\geq 1 + \varepsilon - \frac{\varepsilon(2+\varepsilon)}{(1-\cos\theta)(1+\varepsilon)} \\
&\geq 1 - \frac{2\varepsilon}{1-\cos\theta},
\end{aligned}$$

A slicker analysis as the one used in the proof of the separation of critical points in [DGRS05] can give a stronger result in the latter case. ■

Proof of Lemma 11. Let $x' = \phi_P(t, x)$ be a point on L , satisfying $\|x - x'\| = \tau$ where τ is infinitesimally small. By definition, $A_P(x') = A_P(x)$. Let y be any point at distance $f(h_P(x))$ from x making an angle of ψ with $d_P(x) - x$. We have for the distance of y to x' :

$$\begin{aligned}
\|y - x'\|^2 &= \|y - x\|^2 + \tau^2 + 2\tau\|x - y\| \cos \psi \\
&= f(h_P(x))^2 + \tau^2 + 2\tau f(h_P(x)) \cos \psi.
\end{aligned}$$

For y not to be contained in $B_f(x')$ it must hold that $\|y - x'\| > f(h_P(x'))$, or equivalently:

$$f(h_P(x))^2 + \tau^2 + 2\tau f(h_P(x)) \cos \psi \geq f(h_P(x'))^2.$$

By rearranging we get

$$\tau + 2f(h_P(x)) \cos \psi < \frac{f(h_P(x) + \tau)^2 - f(h_P(x))^2}{\tau}.$$

Taking the limit when $\tau \rightarrow 0$ we get

$$\lim_{\tau \rightarrow 0} \tau + 2f(h_P(x)) \cos \psi < \lim_{\tau \rightarrow 0} \frac{f(h_P(x) + \tau)^2 - f(h_P(x))^2}{\tau},$$

which gives

$$\begin{aligned}
2f(h_P(x)) \cos \psi &< \frac{d}{dh_P} (f(h_P(x)))^2 \\
&= 2f(h_P(x)) \frac{d}{dh_P} f(h_P(x)).
\end{aligned}$$
■

Proof of Lemma 12. Let $y \in B \cap \Sigma$. Since B intersects the medial axis, $f(y) \leq 2R$. Thus there is a sample point within distance $\varepsilon f(y) \leq 2\varepsilon R$ from y . Since B contains no sample points, y must be within distance $2\varepsilon R$ from ∂B . We grow a ball B' centered at x until its boundary touches Σ . Let R' be the radius of B' . By the above argument $R' \geq (1 - 2\varepsilon)R$. Let y be a point in which B' touches Σ . As indicated above, $f(y) \leq 2R$. Let B_o be the tangent ball of radius $f(y)$ at the opposite side of Σ with respect to x . With an argument similar to that of Lemma 9 we get for $h_S(x) = \|x - y\|$:

$$R^2 \leq h_S(x)^2 + \varepsilon^2 f(y)^2 + \varepsilon^2 h_S(x) f(y).$$

Using $f(y) \leq 2R$ we get

$$R^2 \leq h_S(x)^2 + 4\varepsilon^2 R^2 + 2\varepsilon^2 R h_S(x),$$

which by rearranging gives the following quadratic inequality for $h_S(x)$:

$$h_S(x)^2 + 2\varepsilon^2 R h_S(x) - (1 - 4\varepsilon^2) R^2 \geq 0.$$




Figure 4: Left: A Voronoi face intersected by its dual Delaunay edge. The intersection point is a critical point of h_P , which drives all points in the relative interior of the Voronoi face (the arrows indicate the induced vector field). Middle: The flow continues into another Voronoi face. Note that not the whole Voronoi face has to be part of the unstable manifold. Right: The unstable manifold of an index-2-saddle point can have one-dimensional parts, which are orbits of Voronoi vertices.

Since $h_S(x) \geq (1 - 2\varepsilon)R$, the only valid range for $h_S(x)$ in the above inequality is

$$h_S(x) \geq \varepsilon^2 R^2 + \sqrt{\varepsilon^4 R^2 + (1 - 4\varepsilon^2)R^2} \geq (1 - 4\varepsilon^2)R. \quad \blacksquare$$

Proof of Corollary 5. We first consider the case where y has a driving angle $\theta = \theta_P(y)$ with $\cos \theta \leq 1 - \sqrt{\varepsilon}$. By Lemma 10, $\|y - \tilde{y}\| \leq 2\sqrt{\varepsilon}\mu(y)$, or equivalently, $h_S(y) \geq (1 - 2\sqrt{\varepsilon})\mu(y)$. Since $h_P(y) \geq h_S(y)$, we get

$$\|y - \tilde{y}\| \leq 2\sqrt{\varepsilon}\mu(y) \leq \frac{2\sqrt{\varepsilon}}{1 - 2\sqrt{\varepsilon}}h_S(y) \leq \frac{2\sqrt{\varepsilon}}{1 - 2\sqrt{\varepsilon}}h_P(y).$$

When the cosine of the medial angle grows above $1 - \sqrt{\varepsilon}$ in a point y along the flow line $\phi_P(x)$, by Theorem 3 there always is a point within distance

$$\frac{2\sqrt{\varepsilon}}{1 - 2\sqrt{\varepsilon}}h_P(x) \left(\frac{h_P(y)}{h_P(x)} \right)^\xi,$$

from y , where

$$\xi = \frac{1}{1 - \sqrt{\varepsilon}} \left(1 + 4\varepsilon^2 \frac{(1 - 2\sqrt{\varepsilon})^2}{(2\sqrt{\varepsilon})^2} \right) \leq 1 + O(\sqrt{\varepsilon}). \quad \blacksquare$$

B Algorithms

In this section we design an algorithm to compute flow orbits of points under the flow ϕ_P for a finite point set $P \subset \mathbb{R}^3$. Here we are interested in the special case that P is a sampling of a smooth surface Σ . We use the flow orbits to compute unstable manifolds of critical points of h_P and to provide a data structure, a directed acyclic graph, that essentially describes a subdivision of Voronoi facets.

To devise an algorithm that computes the flow orbits we use the oracle $\text{DRIVER}(x)$ that gives us for any query point $x \in \mathbb{R}^3$ its driver. The oracle basically has to solve the point location problem in the Voronoi diagram of the sample points P to find the lowest dimensional Voronoi face that contains x and then compute the closest point on the Delaunay face dual to the Voronoi face. The oracle returns a tuple (d, V) , where d is the driver of x and V is the Voronoi face dual to the highest dimensional Delaunay face that contains d . Note that d is the driver of all points in the relative interior of V and it is also the driver of x , which can be in the boundary of V . The algorithm of Figure 5 computes the orbit of a point $x \in \mathbb{R}^3$. The algorithm

```

ORBIT( $x \in \mathbb{R}^3$ )
1   $O := \{x\}$ 
2   $(d, V) := \text{DRIVER}(x)$ 
3  while  $d \neq x$  and  $x \neq \infty$  do
4     $R :=$  ray from  $d$  in direction  $x$ 
5     $y :=$  last point behind  $x$  along  $R$  that is contained in  $V$ 
6     $O := O \cup xy$ 
7     $x := y$ 
8     $(d, V) := \text{DRIVER}(x)$ 
9  return  $O$ 

```

Figure 5: The algorithm to compute the flow orbit of a given point x in \mathbb{R}^3 .

ORBIT successively adds line segments xy to the orbit O , which gets initialized with the point x itself. This ensures that the algorithm works correctly also in the case that x is a critical point. If x is a critical point it is its own driver and the condition in line 3 will be violated. If x is not critical or a point at infinity the next line segment xy , which starts at x and ends in y is computed in the body of the **while** loop. Note that it is possible that y computed within this loop can be point at infinity. The algorithm will always encounter such a point if the orbit of x does not end in a critical point.

Unstable Manifolds. Recall that the critical points of h_P are exactly the intersection points of Delaunay faces and their dual Voronoi faces, which we consider to be closed. It turns out that the dimension of the Delaunay face characterizes the type of the critical point: it is a local minimum if the Delaunay face is zero-dimensional, a saddle point if the Delaunay face is one- or two-dimensional and a local maximum if the the Delaunay face is three-dimensional (see [GJ03] for more details). For a critical point we call the dimension of its corresponding Delaunay face the index of the critical point.

In order to compute the set \mathcal{C} we have to compute the unstable manifolds of all medial axis critical points. Clearly, local minima, i.e., the points in P cannot be close to the medial axis. All other types of critical points can be and we have to consider all of them. The unstable manifold of a local maximum of h_P is just the local maximum itself, i.e., a single point. Let us now consider the unstable manifold of an index-2-saddle point.

Lemma 13 *Let c be an index-2-saddle point of h_P . The unstable manifold $\mathcal{U}(c)$ of c is a piecewise linear curve.*

Proof. Note that c is the intersection point of a two-dimensional Delaunay face with its dual Voronoi edge E . We show that in a sufficiently small neighborhood U of c we have

$$\mathcal{U}(c) \cap U = E \cap U.$$

Since c is the driver of all points in the relative interior of E we have $E \cap U \subset \mathcal{U}(c) \cap U$. It remains to show $\mathcal{U}(c) \cap U \subset E \cap U$. Every $x \in U \setminus E$ has to be contained in interior of one of the Voronoi faces that contain E in their boundary. Let V be such a Voronoi face and let d be the driver of all points in the relative interior of V . By the definition of the flow ϕ_P the only points $y \in V$ for which $\phi_P(t, y) = x$ for some $t \geq 0$ have to reside within the intersection of V with the line segment that connects d and x . No point of the latter line segment is contained in a sufficiently small neighborhood of c . That is, $x \notin \mathcal{U}(c) \cap U$ and thus $\mathcal{U}(c) \cap U \subset E \cap U$. This implies that $\mathcal{U}(c)$ is the union of E with the orbits of the endpoints of E , i.e., $\mathcal{U}(c)$ is a piecewise linear curve. ■

Using the algorithm ORBIT, it is straightforward to construct an algorithm to compute the closure of the unstable manifold of an index-2-saddle point c . This algorithm is given in Figure 6. We will also use the algorithm ORBIT to compute the unstable manifold of an index-1-saddle point. The latter unstable manifolds are “hairy” piecewise linear surfaces.

Lemma 14 *Let c be an index-1-saddle point of h_P . The unstable manifold $\mathcal{U}(c)$ of c is the union of a piecewise linear surface and the orbits of some Voronoi vertices in the boundary of the surface.*

```

INDEX-2-UNSTABLEMANIFOLD (c)
1  E := Voronoi edge that contains c
2  U := E
3  u, v := endpoints of E
4  if u ≠ ∞ do
5    U := U ∪ ORBIT(u)
6  if v ≠ ∞ do
7    U := U ∪ ORBIT(u)
8  return U

```

Figure 6: The algorithm for computing the unstable manifold of an index 2 saddle point.

Proof. Note that c is the intersection point of a Delaunay edge with its dual two-dimensional Voronoi face V . With a similar argument to the one given in proof of Lemma 13, one can show that in a sufficiently small neighborhood U of c ,

$$\mathcal{U}(c) \cap U = V \cap U.$$

Since c is the driver of all points in the relative interior of V we have $V \subseteq \mathcal{U}(c)$. The orbits of all points in the Voronoi edges incident to V also belong to $\mathcal{U}(c)$. Each such Voronoi edge E falls into one of two cases: either the affine hull of E intersects its dual Delaunay face or it does not.

In the former case we distinguish two sub-cases: either E itself (not only its affine hull) intersects its dual Delaunay face or it does not. In the first sub-case, the union of the orbits of the points in E is exactly the unstable manifold of the intersection point, which is by definition a critical point. If this unstable manifold is different from E , i.e., if the endpoints of E are not critical (local maxima), then the orbits of the endpoints might be hairs of $\mathcal{U}(c)$. In the second sub-case, it is not difficult to see that the union of the orbits of the points in E is the union of E and the orbit of one of its endpoints. The latter orbit also might be a hair of $\mathcal{U}(c)$.

In the latter case, where only the affine hull of E intersects its dual Voronoi face, the driver d of the points in the relative interior of E is on the boundary of the Delaunay face dual to E , i.e., on a Delaunay edge and therefore is also the driver of all the points in the relative interior of the Voronoi face V' dual to the Delaunay edge. Let u and v be the endpoints of E . Since d is by definition contained in the affine hull of V' we can determine the part of V' that belongs to $\mathcal{U}(c)$ by shooting rays from d in the direction of u and v , respectively. These rays enclose a wedge W in the affine hull of V' that contains E . We have $W \cap V' \subset \mathcal{U}(c)$. Note that $W \cap V'$ is linear surface patch. The boundary of this patch consists of Voronoi edges or line segments contained in Voronoi edges. The orbits of all points in these edges/segments also belong to $\mathcal{U}(c)$. We can compute the collection of these orbits recursively as a collection of surface patches and orbits of Voronoi vertices. Note that the “hairs”, i.e., orbits of Voronoi vertices that are not contained in one of the surface patches, can only originate from the boundary of the surface patch and therefore, the relative interior of the unstable manifold is a two-dimensional manifold. ■

We should point out here that the hairs are topologically important. Neglecting the hairs destroys the guarantee on homotopy equivalence of $M(S)$ and the core \mathcal{C} . The characterization of the unstable manifolds of index-1-saddle points immediately suggests the algorithm of Figure 7 to compute them.

Unstable Flow Complex. Since unstable manifolds of two (or more) critical points can have non-empty intersections, a union of several unstable manifolds, represented as the union of cell-complexes forming the individual unstable manifolds does not enjoy the structure of a geometric cell complex embedded in the Euclidean space in a non-degenerate fashion. This is because many cells in this union may (partially) overlap. The unstable flow complex as defined in Section 2 offers a way to cleanly resolve this issue.

Every regular point x on the 2-skeleton of $\text{Vor}(P)$ is in the unstable manifolds of at least two critical points since it is clearly contained in the unstable manifolds of all the minima in $A_P(x)$. On the other hand, the points in the interior of any Voronoi cell V_p of a point $p \in P$ are only flown into from p . This implies that the interior points of V_p , i.e. points that are strictly closer to p than to any other point in P form the cell of the \mathbf{U} associated to the singleton set $\{p\}$ of critical points and that these are the only cells of \mathbf{U} that

```

INDEX-1-UNSTABLEMANIFOLD ( $c$ )
1   $F :=$  Voronoi face that contains  $c$ 
2   $U := F$ 
3   $Q := \emptyset$ 
4  for each Voronoi edge  $E = vw$  incident to  $F$  do
5     $Q.\text{push}(E, v, w)$ 
6  while  $Q \neq \emptyset$ 
7     $(E, v, w) := Q.\text{pop}()$ 
8     $(d, V) := \text{DRIVER}((v + w)/2)$ 
9    if  $E \subset V$  do
10     compute wedge  $W$  from  $v, w$  and  $d$ 
11      $U := U \cup (W \cap V)$ 
12     push all triples  $(E', v, w') \neq (E, v, w)$  for
        maximal line segments  $v'w'$  in the
        boundary of  $W \cap V$  onto  $Q$ 
13  else
14      $U := U \cup E \cup \text{ORBIT}(v) \cup \text{ORBIT}(w)$ 
15  return  $U$ 

```

Figure 7: The algorithm for computing the unstable manifold of an index 1 saddle point.

are associated to singletons. In other words, the 3-cells of \mathbf{U} coincide with those of $\text{Vor}(P)$ and therefore the rest of the cells of \mathbf{U} , are contained in the 2-skeleton of $\text{Vor}(P)$. Indeed these cells subdivide the 2-skeleton of $\text{Vor}(P)$ and the problem of computing the unstable flow complex of P boils down to computing this subdivision. This subdivision can be trivially achieved through computing individual unstable manifolds, as follows: Start with the 2-skeleton of $\text{Vor}(P)$ as \mathbf{U} and iterate over the critical points of index 1 or higher, in an arbitrary order. For each such critical point c , compute the unstable manifold $\mathcal{U}(c)$ and subdivide along the boundary of $\mathcal{U}(c)$, the cells of the current complex \mathbf{U} that are partially contained in $\mathcal{U}(c)$.

Flow Closures and Flow DAG. An important task performed using unstable flow complex is the computation of the flow closure of a collection of its cells. Given a face f of \mathbf{U} , it is easy to compute the set of all points in the space the points in f flow into. Given a Voronoi face g containing a critical point c , since the flow lines on g all start at c , g is not cut by the unstable manifold of any critical point and is therefore wholly contained in \mathbf{U} . The unstable manifold $\mathcal{U}(c)$ is then computed by taking the flow closure of all the points in g (or simply the closure of g).

An important property of the unstable flow complex \mathbf{U} , directly derived from its definition, is that if the orbit $\phi_P(x)$ of a point x of a face f of \mathbf{U} flows into a neighboring face f' , then every point of f' is flown into from some point of f . In other words, every face of \mathbf{U} flows into a set of whole faces. Thus, in computing the flow closure of a face, no other face needs to be subdivided — notice that this is not the case, for example, for the the faces of the Voronoi 2-skeleton. Consequently, one can construct a *flow graph* \mathcal{F} , in fact a directed acyclic graph (DAG), whose vertices are the faces of the unstable flow complex and two vertices f and f' are connected by an arc if the face of \mathbf{U} corresponding to f is incident to and flows into the one corresponding to f' . Given this graph, the unstable manifold of a critical point c can be easily computed by locating the corresponding face of \mathbf{U} , and finding all the faces of \mathbf{U} whose corresponding vertices in \mathcal{F} are reachable from that of c . The core \mathcal{C} of the medial axis can thus be easily computed as the set of all faces of \mathbf{U} whose corresponding vertices in \mathcal{F} are reachable from some vertex of \mathcal{F} corresponding to a face of \mathbf{U} containing a medial axis critical point.

The complex \mathbf{U} also allows us to compute the flow closure of any collection of its faces. In particular, it enables us to compute the the flow closure of any face of $\text{Vor}(P)$ (recall that \mathbf{U} is a refinement of $\text{Vor}(P)$) by simply taking the union of flow closures of the faces of \mathbf{U} that are contained in the given face of $\text{Vor}(P)$.