Geometric and Topological Guarantees for the Wrap Reconstruction Algorithm*

Edgar A. Ramos† Bardia Sadri‡

Abstract

We describe a variant of Edelsbrunner's Wrap algorithm for surface reconstruction, for which we can prove geometric and topological guarantees within the ε-sampling model. The Wrap algorithm is based on ideas from Morse theory applied to the flow map induced by certain distance function. The variant is made possible by a previous result on the "separation" of critical points for a related distance function that directly applies in this case. Though the variant is easily proposed, in order to prove the quality guarantees for the output, we need to closely investigate the geometric properties of the flow map.

1 Introduction

The problem of *surface reconstruction* calls for computing a surface that passes through a given set of points P in $\mathbb{R}^3$ sampled from an unknown underlying smooth surface Σ (in generalizations of the problem, the assumptions on the smoothness of the surface may be lifted, the sample data may be allowed to be noisy, and the constructed surface may only be be required to pass only sufficiently close to the data). This is an important practical problem that has received a great deal of attention from researchers in a number of computational such as vision, graphics, and geometry.

Particularly in computational geometry several approaches have been proposed; many of which with guarantees on the produced output. In order for an algorithm to be able to provide effective guarantees for the output, assumptions on the quality of the input sample P are inevitable. A very common sampling model is the ε-sampling framework proposed by Amenta, Bern and Eppstein [2] which defines a real-valued function f called the *local feature size* on Σ (as the distance to the *medial axis* $M(\Sigma)$ of Σ) and, essentially, requires that the sample set P leaves no ball centered at a surface point x and of radius $\varepsilon \cdot f(x)$ unsampled. Most of these algorithms are based on the Delaunay complex (and its dual Voronoi complex) of the sample points. For example, given an ε-sample of a smooth surface Σ, for a small enough value of ε, the algorithms of Amenta and Bern [1], Amenta, Choi, Dey and Leekha [3], Amenta, Choi and Kolluri [4], and Boissonnat and Cazals [5] reconstruct surfaces with the same topology as Σ that geometrically approximate it in terms of ε. Surveys of these algorithms can be found in [6, 7].

In [10], Edelsbrunner described an algorithm for solving problem by "wrapping finite sets in space". This algorithm was successfully implemented and comercialized (as `geomagic wrap`), and will be referred here generically as the Wrap algorithm. The Wrap algorithm does not make assumptions on the input data, and hence can only make some very general statement about the type of output. Specifically, for the main variant of his algorithm, the output is a *pinched sphere*. The significance of Wrap is in that it is based on discrete methods inspired by concepts in continuous mathematics, such as Morese functions and gradient fields, and develops

*A full version of this paper is available through the authors' web-pages.

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

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

an array of tools and concepts that aims at solving the surface reconstruction problem in a systematic and thorough way.

Edelsbrunner's approach employes a certain distance function together with ideas from Morse theory and creates a decomposition of the space coarser than the Delaunay complex that is exploited for producing a reconstruction. This approach makes essential use of the Delaunay complex $\mathrm{Del}\, P$ and the Voronoi complex $\mathrm{Vor}\, P$ of P. More specifically, for q in the set Q of Voronoi vertices, let w_q be the distance to its closest point in P (the radius of the Delaunay tetrahedron dual to q). The distance function considered is[1]

$$h_Q(x) = \min_{q \in Q}(\|x - q\|^2 - w_q^2).$$

Though this distance function is not smooth everywhere, there is a well developed theory of *critical points* (local extrema and saddle points) of such functions [13] (Edelsbrunner independetly addresses this by a smoothing and pass to the limit process). The critical points of h_Q are precisely those points where a Voronoi cell of P intersetcs its dual Delauany cell. Using Morse theory, h_Q has at any point x a direction of steepest ascent $v_Q(x)$ which induces a *flow path* described by a *flow map* $\phi_Q(x, t)$ and decompositions of the space into so called *stable* and *unstable* sets of the critical points (collection of points that flow towards and "away" from them). The Wrap reconstruction algorithm essentially works by computing an "approximation" to the unstable sets of the critical point at infinity, through Delaunay simplices, and by removing it leaves an output which is a Delaunay subcomplex.

Independently, Giesen and John [11] developed an alternative reconstruction approach based on a more natural distance function, the distance to the points:

$$h_P(x) = \min_{p \in P} \|x - p\|.$$

The algorithm is based on the corresponding stable sets. The critical points of h_P turn out to exactly coincide with those of h_Q. These approaches are in a sense dual to each other. In the primal context of Giesen and John, Dey et al [8] proved a *separation* result for the critical points of h_P: they are concentrated near Σ and near its medial axis $M(\Sigma)$. The separation can be computed effectively and was the basis in the same paper for a provable version of the reconstruction algorithm of Giesen and John, in the context of ε-sampling.

Our goal in this paper is to advance the understanding of the Wrap algorithm. In particular, we are interested in describing a variant for which we can prove geometric and topological guarantees within the ε-sampling model. Specifically, the separation of critical points proved in [8] can be transferred to h_Q, and it allows us to describe a simple modified version of the Wrap, that approximates the volume S bounded by Σ, and for which guarantees can be proved: The output of the modified Wrap differs from S only within a thin tubular neighborhood of Σ and has the same homotopy type. In order to prove these guarantees, understanding of the geometric behavior of the flow map appears to be crucial.

Overview of Paper and Results

In Section 2, we review some preliminary concepts. Further properties of flow maps induced by weighted points are presented in Section 3. Based on these, we describe in Section 4 the modified verision of Wrap. On input P, the modified Wrap outputs a subcomplex I of $\mathrm{Del}\, P$ that is expected to approximate S. Briefly, it collects the set O of all the Delaunay simplices that can be reached through flow lines from the Delaunay simplices corresponding to "exterior" critical points (those that lie outside but not close to Σ), but not reachable from the Delaunay simplices corresponding to other critical point, and then it outputs the rest, i.e. $\mathrm{Del}\, P \setminus I$. The rest of the paper is then devoted to proving guarantees for I. Whenever possible, we develop the results in general dimension, but at the end, we achieve guarantees only for the case of a

[1] Edelsbrunner actually considers $-h_Q$. For technical reasons, we prefer to work with h_Q; in particular, this makes our approach consistent with our work in [8, 12].

surface embedded in $\mathbb{R}^3$. In particular, the goal is to prove that I is homotopy equivalent to S, the bounded volume enclosed by the surface Σ, and that the error between I and S lies within a thin tubular neighborhood Σ_δ of Σ, where $\delta \approx \varepsilon^2$ measures the width relative to the local feature size. As a first step, in Section 5, we show that the flow map $\phi_Q(x,t)$ is continuous on both variables, and as a corollary, we also obtain an integral equation relating h_Q with v_Q and $\phi_Q(x,t)$. This result is an essential part in the homotopy proof in Section 8. Two geometric components are also needed for the homotopy proof to go through. The first, shown in Section 6 states that Σ_δ is flow-tight for h_Q, that is, no flow lines of ϕ_Q leave Σ_δ. The second geometric component, presented in Section 7, which also represents the geometric guarantee of the output, is that the symmetric difference between I and S lies in Σ_δ. This result is only proved for $\mathbb{R}^3$ and is the sole limit to the extension of the results to higher dimensions. Finally, the homotopy equivalence proof is presented in Section 8. It proceeds in three steps. Two of them follow a line of argument used by Lieutier [16] to prove that the medial axis of a bounded open set captures its homotopy type. Specifically, the flow map is used to establish the homotopy. The majority of the proofs are omitted due to lack of space and are available in the full-version of the paper.

2 Preliminaries

Basic Notions For a vector $v \in \mathbb{R}^n$, $\|v\|$ denotes the ℓ^2 norm of v, i.e. $\|v\| = \sqrt{\langle v, v\rangle}$ where $\langle \cdot, \cdot \rangle$ is used to denote the usual inner product on $\mathbb{R}^n$. Given a set T of points in $\mathbb{R}^n$, the *distance* of a point x to T is denoted $\mathrm{dist}(x,T)$ is defined as

$$\mathrm{dist}(x,T) = \inf_{y \in T} \|x - y\|.$$

Given two sets R and T of points, the distance between R and T is the distance between the closest pair of points x and y chosen from R and T respectively.

$$\mathrm{dist}(R,T) = \inf_{x \in R, y \in T} \|x - y\| = \inf_{x \in R} \mathrm{dist}(x,T).$$

For a point $x \in \mathbb{R}^n$ and $r > 0$, $B(x,r)$ denotes the open ball of radius r centered at x, i.e. $B(x,r) = \{y \in \mathbb{R}^n : \|x - y\| < r\}$. The *closed ball* $\overline{B}(x,r)$ is the closure of $B(x,r)$, i.e., $\overline{B}(x,r) = \{y \in \mathbb{R}^n : \|x - y\| \le r\}$.

Shapes and Surface We consider closed connected smooth $(n-1)$-manifolds embedded in $\mathbb{R}^n$. We call such manifolds *surfaces.* A surface Σ partitions its complement $\mathbb{R}^n \setminus \Sigma$ into two open sets: a bounded or *inner* component S and an unbounded or *outer* component S^*. The inner component is sometimes called the *shape* enclosed by Σ. Throughout this paper, Σ represents an unknown but fixed surface. Likewise, throughout, S and S^* refer respectively to the inner and outer shapes determined by Σ. Note that $\Sigma = \partial S = \partial S^*$ where ∂S represents the boundary of S. For a point $x \in \Sigma$, n_x denotes the direction of the normal to Σ at x. Of the two possible orientations of this direction, n_x^+ is the normal vector pointing toward the exterior of Σ, (into S^*) and n_x^- is the one pointing toward its interior (into S).

Medial Axis In general, the *medial axis* $M(T)$ of a set $T \subset \mathbb{R}^n$ is the set of points in T that have at least two closest points in ∂T. Formally, if we define for a point $x \in T$ the set $A_T(x)$ as

$$A_T(x) = \{y \in \partial T : \|x - y\| = \mathrm{dist}(x,T)\},$$

then the medial axis of T is given by

$$M(T) = \{x \in T : |A_T(x)| > 1\}.$$

Thus the medial axis $M(S)$ of the inner shape S determined by a surface Σ, also called the *inner medial axis* of Σ, is the set of all points in S that have at least two closest points in $\partial S = \Sigma$.

Note that since Σ is compact, $A_S(x)$ is well-defined and non-empty for every $x \in S$. The *outer medial axis* of Σ, $M(S^*)$ can be defined similarly. By medial axis $M(\Sigma)$ of Σ (or just M when sigma is understood), we refer to the union of the inner and outer medial axes of Σ, i.e.

$$M(\Sigma) = M(S) \cup M(S^*).$$

In other words, M is the set of all points in $\mathbb{R}^n$ that have at least two closest points in Σ.

Feature Size and Surface Samples

By definition, every point of $\mathbb{R}^n \setminus M$ has a unique closest point in Σ. For any point $x \in \mathbb{R}^n \setminus M$, $\hat{x}$ denotes the unique closest surface point to x, i.e.,

$$\hat{x} = \operatorname*{argmin}_{y \in \Sigma} \|x - y\|,$$

and $\check{x} \in M$ denotes 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 $\mu : \mathbb{R}^n \setminus (\Sigma \cup M) \to \mathbb{R} \cup \{\infty\}$ defined as $\mu(x) = \|\hat{x} - \check{x}\|$. The function

$$f : \Sigma \to \mathbb{R}, \quad x \mapsto \inf_{y \in M} \|x - y\|,$$

which assigns to each point in Σ its distance to M, is called the *local feature size*. Notice that for $x \in \mathbb{R}^n \setminus (\Sigma \cup M)$ 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,

$$f_{\min} = \inf_{x \in \Sigma} f(x),$$

called the *reach* of Σ, is strictly positive. This requirement is automatically fulfilled when Σ is a $C^{1,1}$ or smoother surface.

Sampling Criteria We consider two highly popular sampling criteria. The most natural and by far the most common in literature is known as the *uniform ε-sampling.* For a constant $\varepsilon > 0$, a point set P is a uniform ε-sample of a surface Σ if every point of Σ has a sample point in its closed ε-neighborhood, i.e.

$$\forall x \in \Sigma, \qquad \overline{B}(x, \varepsilon) \cap P \neq \emptyset.$$

Uniform sampling is insensitive to the size of surface features and can in this sense be deemed wasteful. It is therefore desirable to allow the sample to become sparser at larger features of the surface. The local feature size function $f(\cdot)$ as defined above can be regarded a well-behaving (Lipschitz) conservative local measure of feature size — it can be shown that for every surface point x, $1/f(x)$ upper bounds the largest principle curvature of the surface at x. Following Amenta and Bern [2], for a constant $\varepsilon > 0$, a finite sample $P \subset \Sigma$ is called a relative (or *adaptive*) ε-sample of Σ if every point $x \in \Sigma$ has a sample point within distance $\varepsilon f(x)$ from x, i.e.

$$\forall x \in \Sigma, \qquad \overline{B}(x, \varepsilon f(x)) \cap P \neq \emptyset.$$

Reduced Shapes and Tubular Neighborhoods Offset surfaces and shapes play a crucial role in our approach. Given the surface Σ, one can define tubular neighborhoods of Σ an the corresponding offset surfaces in two different ways.

Given a parameter $0 < \delta$, the *uniform δ-tubular neighborhood* $\hat{\Sigma}_\delta$ as the set

$$\hat{\Sigma}_\delta = \{x \in \mathbb{R}^n : \operatorname{dist}(x, \Sigma) \leq \delta\}.$$

In other words $\hat{\Sigma}_\delta$ is simply the Mikowski sum of Σ and $\overline{B}(0, \delta)$. Consider now what remains when we remove $\hat{\Sigma}_\delta$ from $\mathbb{R}^n$. Each of the two shapes S and S^* determined by Σ will be further

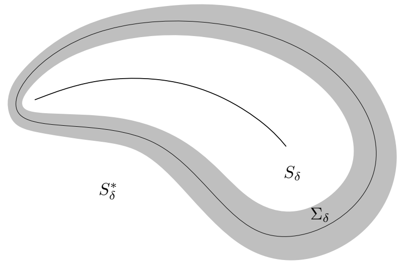

Figure 1: Σ_δ, S_δ, and S^*_δ.

trimmed. We define the uniformly reduced shape $\hat{S}_\delta$ as $S \setminus \hat{\Sigma}_\delta$. The outer uniformly reduced δ-shape $\hat{S}^*_\delta$ is defined similarly.

For a constant $0 < \delta < 1$, the *(relative) δ-tubular neighborhood* Σ_δ of Σ (See Figure 1) is the set

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

It can be shown that like Σ, Σ_δ partitions its complement $\mathbb{R}^n \setminus \Sigma_\delta$ into two components which we call the *δ-reduced shapes.* The bounded or the inner δ-reduced shape S_δ is defined as $S_\delta = S \setminus \Sigma_\delta$. Likewise, the outer δ-reduced shape is $S^*_\delta = S^* \setminus \Sigma_\delta$. Notice that the definition of Σ_δ, puts the medial axis $M(S)$ in S_δ and $M(S^*)$ in S^*_δ. However, it can be shown that points of $M(S)$ are in fact interior point of S_δ and the same relation exists between $M(S^*)$ and S^*_δ. Furthermore, it can be shown that the boundary of Σ_δ consists of exactly those points for which the equality $\|x - \hat{x}\| = \delta f(\hat{x})$ holds (See [12] for proofs and more details). The boundary of Σ_δ has two components: an inner component which coincide with the boundary of S_δ and an outer component which equal the boundary of S^*_δ.

Squared Distance and its Critical Points

Given a sample P of Σ, the *square distance* function to P defined as

$$h_P : \mathbb{R}^n \to \mathbb{R}, \qquad x \mapsto \min_{p \in P} \|x - p\|^2,$$

assigns to each point of space the square of its distance to the sample P. Distance functions have been widely studied and are known to carry a great deal of information about their inducing sets and their embeddings in the space. In particular, the *critical points* of such distance functions are of great interest. The the set of points at distance $h_P(x)$ from a point $x \in \mathbb{R}^n$ is denoted by $A_P(x)$. Critical points of h_P are those points x that satisfy $x \in \operatorname{conv} A_P(x)$, i.e. the points that are contained in the convex hull of their closest points in P. Equivalently, a point x is a critical point of h_P if it is the intersection of a Voronoi face of $\operatorname{Vor} P$ and its dual Delaunay face in $\operatorname{Del} P$ (See [13, 16, 15, 8] for more information about the critical point theory of distance functions and their applications).

In this paper we study square distance functions that are induced by *weighted points* in $\mathbb{R}^n$. The graph of an ordinary square distance function can be thought of a collection of similar paraboloids all tangent to the horizontal plane in $\mathbb{R}^{n+1}$. When considering weighted points, we allow these paraboloids to move up or down vertically and thus no longer be based at the horizontal plane. This concept is explored in detail in Section 3.

Separation of Critical Points Dey, et al. [8] observed that if P is a dense enough sample of the smooth surface Σ, then the critical points of the discrete distance function h_P have to be either very close to Σ or very close to M.

Theorem 2.1 [8] *Let P be an adaptive ε-sample of a smooth surface Σ for $\varepsilon < 1/\sqrt{3}$. Then for every critical point c of h_P, either (i) $\|c - \hat{c}\| \leq \varepsilon^2 f(\hat{c})$, or (ii) $\|c - \check{c}\| \leq 2\varepsilon^2 f(\hat{c})$.*

If the sample P is uniform with a sampling density ε proportional to smallest local feature size $f_{\min}$, the above theorem can be restated as follows. Notice that only the bound on critical points near the surface is improved.

Theorem 2.2 [8] *Let P be a uniform ε-sample of a smooth surface Σ for $\varepsilon < 1/\sqrt{3} \cdot f_{\min}$ where $f_{\min} = \inf_{x\in\Sigma} f(x)$. Then for every critical point c of h_P, either (i) $\|c - \hat{c}\| \leq \varepsilon^2$, or (ii) $\|c - \check{c}\| \leq 2\varepsilon^2 f(\hat{c})$.*

Thus the 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 to the second class as *medial axis critical points*. We further subdivide the medial axis critical points of 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 $M(S^*)$.

3 Flow Induced by Weighted Points

By a *weighted* point in the n dimensional Euclidean space is a pair $(q, w_q) \in \mathbb{R}^n \times \mathbb{R}$ where $q \in \mathbb{R}^n$ is an ordinary point to which the real *weight* w_q is assigned. An *unweighted* point is assumed to have weight zero. The *squared distance* or *power distance* to a weighted point q is given by the function

$$\pi_q : \mathbb{R}^n \to \mathbb{R}, \quad x \mapsto \|x - q\|^2 - w_q.$$

Given a discrete set Q of weighted points, the weighted squared distance to the set Q is defined by the map

$$h_Q : \mathbb{R}^n \to \mathbb{R}, \quad x \mapsto \min_{q\in Q} \pi_q(x) = \min_{q\in Q} \left(\|x - q\|^2 - w_q\right).$$

Sometimes we need to refer to the unweighted squared distance to a weighted set of points which is simply the map obtained by setting all point weights to zero and is defined as

$$h_Q^\circ : \mathbb{R}^n \to \mathbb{R}, \quad x \mapsto \min_{q\in Q} \|x - q\|^2.$$

Given a weighted set of points $Q \subset \mathbb{R}^n$, and a point $x \in \mathbb{R}^n$, $A_Q(x)$ denotes the preimage of h_Q, i.e. $A_Q(x) = \{q \in Q : h_Q(x) = \pi_q(x)\}$. We also define $A_Q^\circ(x)$ in a similar manner but with respect to h_Q°: $A_Q^\circ(x) = \left\{q \in Q : h_Q^\circ(x) = \|x - q\|^2\right\}$.

Voronoi and Delaunay Complexes of Weighted Points Consider the binary relation "$\sim_Q$" between points in $\mathbb{R}^n$ given by

$$x \sim_Q y \quad \Leftrightarrow \quad A_Q(x) = A_Q(y).$$

Trivially, "$\sim_Q$" is an equivalence relation. The equivalence classes of "$\sim_Q$" partition the space into convex regions. Each such region is an open set with respect to the relative topology of its affine hull. Closures of these regions form faces of a cell complex commonly known as the *weighted* Voronoi complex, or the *power complex* of Q and is denoted $\operatorname{Vor} Q$. For each $q \in Q$, the closure of the set of points x for which $A_Q(x) = \{q\}$ is called the Voronoi *cell* of q, denoted by V_q. Unless it is empty, V_q is a full dimensional convex set, meaning that its affine hull is the entire $\mathbb{R}^n$. The Delaunay complex $\operatorname{Del} Q$ of Q is defined as a cell complex dual to $\operatorname{Vor} Q$. The faces of $\operatorname{Del} Q$ are convex hulls of those subsets $U \subseteq Q$ for which $U = A_Q(x)$ for some $x \in \mathbb{R}^n$. For any point $x \in \mathbb{R}^n$, $V_Q(x)$ denotes the equivalence class of x under the relation "$\sim_Q$". It follows that the closure of $V_Q(x)$ is the lowest dimensional face of $\operatorname{Vor} Q$ that contains x in its relative interior. We also denote by $D_Q(x)$ the dual to $V_Q(x)$ in $\operatorname{Del} Q$. Thus, $D_Q(x) = \operatorname{conv} A_Q(x)$. For a cell complex K, the underlying space of K, i.e. the union of all faces in K, is denoted by $|K|$.

Generalized Gradients

The map $h_Q : \mathbb{R}^n \to \mathbb{R}$, is continuous(it can be shown to be the signed distance between the lower hull of a set of points and a paraboloid interpolating them; See proof of Proposition 6.1). However, h_Q is not globally differentiable. In particular, the gradient of h_Q is undefined at any point x for which $|A_Q(x)| > 1$. Nevertheless, Giesen and John [11] showed that there is a unique direction of steepest ascent at every regular point of h_Q, extending the gradient of h_Q to a vector field $v_Q : \mathbb{R}^n \to \mathbb{R}^n$ defined over the entire space $\mathbb{R}^n$. This generalized gradient vector field v_Q can be characterized as follows: for a point $x \in \mathbb{R}^n$, let $d_Q(x)$, called the *driver* of x, be the closest point to x in $D_Q(x) = \operatorname{conv} A_Q(x)$. As was shown by Giesen and John, moving opposite to $d_Q(x)$ locally maximizes over all directions the growth rate of h_Q. We thus define the vector field v_Q as

$$v_Q(x) = 2(x - d_Q(x)).$$

The set of points for which $v_Q(x) = 0$ or equivalently $x = d_Q(x)$, are called the *critical* points of h_Q. All other points are called *regular* points. It can be verified that the critical points of h_Q are the intersection points of weighted Voronoi faces and their dual Delaunay faces (when they do intersect). Alternatively, a point x is a critical point of h_Q if and only if $x \in \operatorname{conv} A_Q(x)$. The fact that $A_Q(x)$ is the same for every point x in the relative interior of every Voronoi face σ entails that that all such points have the same driver. We thus use the notations $A_Q(\sigma)$ and $d_Q(\sigma)$ to respectively denote the set of closest points and the driver common to all points in the relative interior of a Voronoi face σ.

It is important to observe that when all the points have weight zero, the above definitions and characterizations (module adjusting the length of the vectors $v_Q(x)$) exactly match those in [15, 16, 8] and others which consider only unweighted points.

Integration of the Generalized Gradient into a Flow Map

At every point x of a fixed Voronoi face τ, $v_Q(x) = 2(x - d_Q(\tau))$ and thus the vector field v_Q can be integrated to result integral lines which follow a pencil of half lines diverging from the driver $d_Q(\tau)$ of τ. As a result the vector field v_Q can be integrated throughout $\mathbb{R}^n$ resulting *integral lines* or *flow paths* that are in general piece-wise linear curves that may turn only at points where they reach the relative interior of a new Voronoi face. The induced flow map $\phi_Q : (\mathbb{R}^+ \cup \infty) \times \mathbb{R}^n \to \mathbb{R}^n$ assign to each pair (t, x) with $t \in \mathbb{R}^+ \cup \infty$ interpreted as *time* and $x \in \mathbb{R}^n$, the point $\phi_Q(t, x)$ along the flow path out of x that is reached at time t. The map ϕ_Q has the standard properties of a flow map, i.e. $\phi_Q(0, x) = x$, and $\phi_Q(s + t, x) = \phi_Q(s, \phi_Q(t, x))$. Moreover, at every point $\phi_Q(t, x)$ of a flow line $\phi_Q(x)$, except turning points, $\dot{\phi}_Q(t, x) = v_Q(\phi_Q(t, x))$, where $\dot{\phi}_Q(t, x)$ is used to denote $d\phi_Q(t, x)/dt$. In general $v_Q(\phi_Q(t, x))$ is the *right derivative* of $\phi_Q(t, x)$.

Flow Orbit, Flow-Tight, and Flow-Repellant Sets

A non-critical point is also called a *regular* point. For a given flow map ϕ_Q, the *flow orbit* of a regular point x, denoted $\phi_Q(x)$ is the set of all points $\phi_Q(t, x)$ for all $t > 0$. The flow orbit of a critical point is the critical point itself. The flow orbit $\phi_Q(T)$ of a set T is the union of flow orbits of all its points, i.e. $\phi_Q(T) = \bigcup_{x \in T} \phi_Q(x)$. Notice that by this definition $T \subseteq \phi_Q(T)$. A set T is said to be *flow-tight* if $\phi_Q(T) = T$. We call a set T *flow-repellant* if T^c is flow-tight.

Stable and Unstable Manifolds

Then the flow path $\phi_Q(t, x)$ of every point x converges to a critical point c of h_Q or goes to infinity as $t \to +\infty$. Notice that we consider c to be also in the flow orbit of x. For a critical point c of h_Q, the set of all points x whose flow orbits converges to c is called the *stable manifold* of c and is denoted by $\operatorname{Sm}(c)$. In other words, $\operatorname{Sm}(c) = \{x \mid \phi_Q(+\infty, x) = c\}$. 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*

Algorithm (original)Wrap(sample point-set P)

Let $\Delta \subseteq D$ be the set of *critical* simplices.
Let $O = \{\tau \in D : \omega \trianglelefteq \tau \text{ and } \forall \sigma \in \Delta : \sigma \not\trianglelefteq \tau\}$.
Return $I = D \setminus O$.

Figure 2: The original Wrap algorithm.

manifold $\mathrm{Um}(c)$ of a critical point c, as the set of all points into which points arbitrarily close to c flow. Formally, $\mathrm{Um}(c) = \bigcap_{\varepsilon>0} \bigcup_{y \in B(c,\varepsilon)} \phi_Q(y)$. With an abuse of terminology, we say that c "flows" into the points of $\mathrm{Um}(c)$.

4 The Wrap algorithm

We are now ready to describe the Wrap algorithm. Our presentation of the algorithm differs slightly from that of Edelsbrunner and this is due to our description of flow lines which correspond to *limit curves* studied by Edelsbrunner only with opposite orientation.

Given an ε-sample $P \subset \Sigma$ of a smooth[2] manifold Σ embedded in $\mathbb{R}^n$ and of codimension one, we define a set Q of weighted points consisting of the Voronoi vertices in $\mathrm{Vor}\, P$ with every vertex $q \in Q$ given the weight $w_q = h_P(q)$, i.e. the square of the distance between q and its closest points in P. It is a classical result that when restricted to $\mathrm{conv}\, Q$, $\mathrm{Vor}\, P = \mathrm{Del}\, Q$ and when restricted to $\mathrm{conv}\, P$, $\mathrm{Del}\, P = \mathrm{Vor}\, Q$ and in this sense the point sets P (unweighted) and Q (weighted), as well as their Voronoi and Delaunay complexes are called *dual* to each other.

An immediate consequence of this duality is that the critical points of h_Q are exactly the same as those of h_P. We call those simplices of $\mathrm{Del}\, P$ that contain a critical point, i.e. those that intersect their dual faces in $\mathrm{Vor}\, P$, *critical* simplices (or *centered* in Edelsbrunner's terminology). The intersection of a flow line of ϕ_Q and a full-dimensional simplex of $\mathrm{Del}\, P = \mathrm{Vor}\, Q$ is a line segment (if not empty). With lower dimensional simplices there is a second possibility, namely, a flow line can *cross* the simplex and thus intersect it in a single point. In such a case, the simplex is called *transversal* (or *equivocal* according to Edelsbrunner). Of course, the flow line can just as well intersect a non-full-dimensional simplex in a line segment in which case we say that the flow is *tangential* on the simplex in question or simply call the simplex *tangential* (Edelsbrunner calls such simplices *confident*).

We say a simplex τ precedes a simplex σ and denote it by $\tau \prec \sigma$ if τ and σ are incident simplices, i.e. τ is either a face or a coface of σ, and some flow line of ϕ_Q enters the relative interiors of σ immediately after leaving the relative interior of τ. More formally, when τ is a coface of σ, $\tau \prec \sigma$ if there exists a point x and and a time $t_0 > 0$ and a real number $0 < \alpha < t_0$ such that $\phi(t_0) \in \mathrm{rel\ int}\, \tau$ and $\phi(t, x) \in \mathrm{rel\ int}\, \sigma$ for $\alpha < t < t_0$. Similarly, when τ is a proper face of σ, $\tau \prec \sigma$ if for some point x there exist time $t_0 > 0$ and real $\alpha > t_0$ such that $\phi(t_0, x) \in \mathrm{rel\ int}\, \tau$ and $\phi(t, x) \in \mathrm{rel\ int}\, \sigma$ for $t_0 < t < \alpha$. We define the relation "$\preccurlyeq$" as the reflexive transitive closure of "$\prec$", namely, $\tau \preccurlyeq \sigma$ if there is a sequence $\tau = \tau_0 \prec \cdots \prec \tau_k = \sigma$ with $k \geq 0$.

Remark Edelsbrunner's definition of the precedence relation [10], which is denoted by "$\lhd$" is slightly different from ours, in that $\tau \lhd \sigma$ if $\tau \prec \sigma$, and in addition, the flow on one of τ or σ is transversal. This definition thus invalidates $\tau \lhd \sigma$ in the case where $\tau \prec \sigma$ but the flow is tangential in τ and reaches the face σ of τ and continues tangentially on σ. Note that no other case is possible; flow cannot cross two incident simplices transversally and cannot move tangentially from a face to a coface. The reflexive transitive closure of "$\lhd$" is denoted by "$\trianglelefteq$".

[2]Our analysis requires only $C^{1,1}$ smoothness for the target manifold Σ.

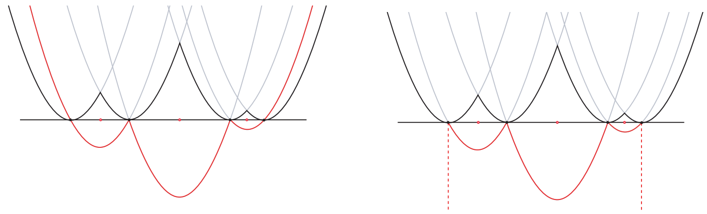

Figure 3: The function h_Q and its modification for the Wrap algorithm.

A subtle technical issue in the definition of the Wrap algorithm involves the assumption of a critical point at infinity; a *minimum*. Since the point is at infinity, its weight, i.e. its distance to the point set P is $+\infty$. This is in contrast with our definition of the function h_Q as can be seen in the one dimensional example of Figure 3. In the figure on the left the points in P are represented by solid bullets on the horizontal line. The graph of the distance function h_P (the black curve) is the the lower envelope of the the parabolas placed at every point of P. The points in Q, i.e. the Voronoi vertices of $\mathrm{Vor}\, P$ are midpoints of consecutive pairs of points in P and are shown by hollowed red bullets. The weight w_q of a point q in Q is $h_P(q)$, i.e. the vertical distance between q and the black curve. Thus the weight of every point in Q is positive. If we now base a parabola with offset w_q below every point $q \in Q$, the graph of the function h_Q (the red curve) is exactly the lower envelope of this second set of parabolas. Observe that h_Q is non-positive everywhere inside $\mathrm{conv}\, P$ but goes to $+\infty$ when the distance to $\mathrm{conv}\, P$ grows infinitely large and as can be seen in the figure, h_Q has no critical points outside $\mathrm{conv}\, P$. However, If we assume in addition that there is a critical point of weight $+\infty$ at infinity, the graph of the function h_Q can be modified to look like the red curve in the figure on the right. The dotted segments in the figure are vertical lines that correspond to the two arms of a parabola placed infinitely far away at an infinite offset below the horizontal axis.

The effect of this modification to h_Q on Wrap is that the algorithm treats $\mathbb{R}^n \setminus \mathrm{conv}\, P$ as a special *abstract* critical simplex ω which corresponds, and contains, this critical point. Since this critical point is infinitely far away, every simplex τ of $\mathrm{Del}\, P$ that is contained in the boundary of $\mathrm{conv}\, P$ is considered preceded by ω.

With these preliminaries covered, Edelsbrunner's Wrap algorithm can now be stated as shown in Figure 2.

As stated above, the output of Wrap is not guaranteed to agree topologically with the sampled surface Σ. In fact, Edelsbrunner proves in [10] that the produced output I is the boundary of a contractible volume. However, Edelsbrunner also suggests methods for extending the algorithm in order to make possible the production of non-simply-connected output. For example, he suggests to consider, in addition to the simplices that are preceded by ω, those that are preceded by other "significant" critical simplices. In essence, the primary result of this paper is that this intuition can be made into an algorithm with provable guarantees. We present a way of determining the significance of critical simplices in a way that such a modification can certify the geometric closeness and topological accuracy of the output.

In Figure 5, we present a modified version of Wrap which can capture the topology of Σ or rather the bounded volume S enclosed by it. The modification rests primarily on the notion of *separation of critical points* as discovered in [8] and stated in Theorem 2.1. Separation of critical points allows us to filter out the so-called *surface critical points* which are in essence *artifacts of discretization*. Furthermore, [8] provides an algorithmic way of distinguishing between medial axis critical points that are contained in S as opposed to those contained in S^* by building the incidence graph of stable manifolds of critical points in the stable flow complex induced by P and looking at the connected components of the subgraph induced by the vertices corresponding to medial axis critical points. The critical point at infinity is naturally assumed to be in the second group, i.e. contained in S^*. Our algorithm shown in Figure 5 amends Wrap by adding to ω all other outer medial axis critical simplices.

The rest of this paper is dedicated to proving that this modified version of Wrap produces

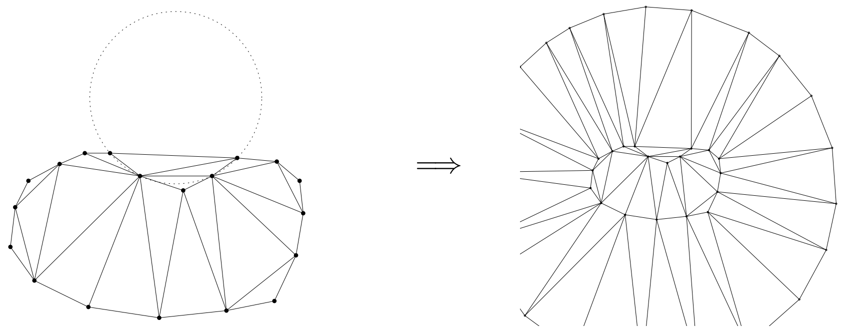

Figure 4: Extension of the sample points for pushing the flow toward conv P near it.

an output that is geometrically close to S and is has the same homotopy type as S provided that the input P to the algorithm is an ε-sample of Σ for a sufficiently small value of ε. In the sequel WRAP refers to this modified version.

4.1 The Extended Sample

Before we proceed we address another technical difficulty related to the assumption of the existence of a critical point at infinity. Our proofs heavily use the flow ϕ_Q induced by the set Q of Voronoi vertices in Vor P as a homotopy. It is therefore crucial to ascertain the continuity of ϕ_Q over the bounded region it is used. Per se, the argument given for continuity of ϕ_Q in Section 5 has no way of handling a critical point at infinity. It turns out that we cannot ignore this critical point either. Since ϕ_Q is a steepest ascent flow, and because h_Q without the extra critical point at infinity, goes to infinity when the distance to conv P grows infinitely large, the flow ϕ_Q escapes conv P outside of it. This is in contrast to the behavior the modified h_Q suggests and it turns out that the modified h_Q in which the steepest ascent outside the convex hull of P is toward conv P is the proper choice for our topological proofs. However, assuming an infinitely heavy critical point at infinity gives us an infinitely fast moving flow map near and outside conv P.

Several of our proofs in what follows study the behavior of ϕ_Q on the outer boundary of Σ_δ. These proofs are based on the assumption that every point on this boundary is contained in some Delaunay n-simplex. This assumption trivially fails when the point in question is outside conv P. We handle these problem by adding a component Σ_0 to the surface Σ in such a way that the local feature size of every point in Σ with respect to $M(\Sigma \cup \Sigma_0)$ is the same as its local feature size with respect to $M(\Sigma)$ alone. This will then imply that if P_0 is an ε-sample of Σ_0, $P_1 = P \cup P_0$ is an ε-sample of $\Sigma_1 = \Sigma \cup \Sigma_0$. Moreover, the component of the tubular neighborhood Σ_δ around Σ in the two-component surface $\Sigma \cup \Sigma_0$ will be identical to the one for the surface Σ alone.

Thus consider a ball $B = B(c, R)$ enclosing Σ and therefore the sample $P \subset \Sigma$. Let R_D be the circumradius of the largest Delaunay ball in Del P. Consider the ball $B_0 = B(c, R_0)$ where $R_0 \geq 4R + 2R_D + 2F$ in which $F = \sup_{x \in \Sigma} f(x)$. let $\Sigma_0 = \partial B$ be the a new component of the surface and let P_0 be an ε-sample of it. It is easy to see that every Delaunay ball of Del P is entirely contained in the interior of B and therefore remains empty of the points in P_0. Consequently, Del P is a subcomplex of Del P_1 where $P_1 = P \cup P_0$ is the extended point set which samples the two-component surface $\Sigma_1 = \Sigma \cup \Sigma_0$ (See Figure 4).

Let z be a point of $M_1 = M(\Sigma_1)$, i.e. the medial axis of the extended surface. If $z \in B(c, R_0/2)$, then every closest point of Σ_1 to z has to be in Σ and therefore $z \in M = M(\Sigma)$. Thus, for a point $x \in \Sigma$, the distance to M_1 is at least as large as the distance between x and M since the points outside $B(c, R_0/2)$ are at least $R + R_D + F$ distance away from x while x has a point within distance F in M. Thus $P_1 = P \cup P_0$ is a valid ε-sample of Σ_1. This in particular ensures that the separation of critical point given Theorem 2.1 remains valid for separation of critical points of h_{P_1}. We intend to use the set Q_1 of all Voronoi vertices in Vor P_1, weighted as seen above, instead of the Q solely for our analysis of the WRAP algorithm. Since the critical points of h_{P_1} and h_{Q_1} are the same, Theorem 2.1 implies that the critical points of h_{Q_1} are contained either in the ε^2 tubular neighborhood of Σ_1, or near its medial axis.

Proposition 4.1 *The surface critical points of h_{Q_1} contained in Σ_{ε^2} are exactly the same as the critical points of h_Q contained in Σ_{ε^2}.*

Proof. Let $c \in \Sigma_{\varepsilon^2}$ be a critical point of h_{Q_1}. Since $c \in \Sigma_{\varepsilon^2}$, $\|c - \hat{c}\| \leq \varepsilon^2 f(\hat{c})$. By the ε-sampling condition, $\hat{c}$ has a point $p \in P$ within distance $\varepsilon f(\hat{c})$. Thus

$$\sqrt{h_{P_1}(c)} \leq \sqrt{h_P(c)} \leq \varepsilon^2 f(\hat{c}) + \varepsilon f(\hat{c}) \leq 2\varepsilon F < 2\varepsilon R_1.$$

Thus $A_{P_1}(c) \subset P$. Since c is a critical point, $c \in \operatorname{conv} A_{P_1}(c) \subset \operatorname{conv} P$. Thus c is a critical point of h_P. ∎

It is useful to study the change of drivers when a flow line moves from a Voronoi face into another. Let τ and σ be two consecutive Voronoi faces in $\operatorname{Vor} Q$ visited by a flow path $\phi_Q(x)$ in the same order, thus in our notation $\tau \prec \sigma$. Since τ and σ are incident faces of $\operatorname{Vor} Q$, τ is either a face ($\tau < \sigma$) or a coface of σ ($\tau > \sigma$). If $\tau < \sigma$, then drivers for τ and σ are the same point. To see this first observe that the flow in σ has to be tangential and therefore, the drivers $d_Q(\tau)$ and $d_Q(\sigma)$ must both be in the affine hull of σ. Also, σ cannot be critical since otherwise no flow line can enter it. Thus $d_Q(\sigma)$ is outside σ. Let y be the point in τ where $\phi_Q(x)$ leaves τ and let z be a point on $\phi_Q(x)$ in the relative interior of σ. Notice that by definition, $\pi_q(z) > \pi_p(z)$ for all $p \in A_Q(\sigma)$ and $q \in A_Q(\tau) \setminus A_Q(\sigma)$. Thus the hyperplane orthogonal to $v_Q(z)$ and containing the affine hull of $A_Q(\sigma)$ separates σ and $A_Q(\tau) \setminus A_Q(\sigma)$. The point $d_Q(\sigma)$ is the closest point to z on $\operatorname{conv} A_Q(\sigma)$. With the the above observation, $d_Q(\sigma)$ is just as well the closest point to z on $\operatorname{conv} A_Q(\tau)$. This is true for every $z \in \phi_Q(x)$ and in the relative interior of σ. In particular, z can be chosen arbitrarily close to y. Consequently, $d_Q(\sigma)$ is the closest point to y in $\operatorname{conv} A_Q(\tau)$ which is in turn $d_Q(\tau)$.

On the other hand, if $\tau > \sigma$, $A_Q(\tau) \subset A_Q(\sigma)$, and in particular $\operatorname{conv} A_Q(\tau) \subseteq \operatorname{conv} A_Q(\sigma)$. This implies that the new driver ($d_Q(\sigma)$) is no farther than the old one, i.e. $d_Q(\tau)$.

These observations lead to the following proposition.

Proposition 4.2 *If a flow line $\phi_Q(x)$ consecutively intersects the relative interiors of faces τ and σ in $\operatorname{Vor} Q$, then for every point $z \in \phi_Q(x) \cap \sigma$, $\|z - d_Q(\sigma)\| \leq \|z - d_Q(\tau)\|$ with equality holding only if σ is a proper coface of τ.*

It is not hard to see that proof given in [10] for the acyclicity of the precedence relation does indeed generalize to "$\prec$" as defined here and therefore "$\prec$" defines a partial order on the simplices of $\operatorname{Del} P$. If τ and σ are incident simplices in $\operatorname{Del} P$, they have at least one vertex $p \in P$ in common. Using p as z in the statement of the Proposition 4.2 givse us the following observation.

Observation If $\tau \prec \sigma$ then $h_P(d_Q(\tau)) \leq h_P(d_Q(\sigma))$ with equality holding only if τ is a proper face of σ, i.e. $\dim \tau < \dim \sigma$.

The above lemma establishes the acyclicity of the relation "$\prec$", similar to the way Edelsbrunner shows for the slightly different version of of this relation in [10], by showing that the for if $\tau \prec \sigma$, the pair $(h_P(d_Q(\tau)), \dim \tau)$ is strictly larger lexicographically than $(h_P(d_Q(\sigma)), \dim \sigma)$.

Our topological proofs heavily make use of the fact that the map ϕ_Q is indeed continuous on both variables. This and other properties of ϕ_Q are proved in the subsequent chapter.

5 Continuity of the Flow Map

Clearly, the map ϕ_Q changes continuously with changing of its first parameter; the image of $\phi_Q(\cdot, x)$ is a continuous curve.

Algorithm (modified)Wrap(sample point-set P)

```
C ← set of the critical points of h_P (or h_Q).
(C_M^+, C_M^-, C_Σ) ← Separate(C).
Let Δ ⊂ D be the set of critical simplices.
Let Δ* be the set of critical simplices corresponding to C_M^+ plus ω.
Let O = {τ ∈ D : ∃σ ∈ Δ*, σ ≼ τ and ∀σ ∈ Δ \ Δ* : σ ⋠ τ}.
Return I = D \ O.
```

Figure 5: The modified Wrap algorithm.

We shall first characterize the rate of flow along each linear piece of a flow path and then prove that ϕ_Q is indeed continuous on its second parameter as well.

Lemma 5.1 *Let $x \in \mathbb{R}^n$ and $T \in \mathbb{R}^+$ be such that the flow path $\phi_Q([0,T], x)$ consists of a single line segment. Then for all $t \in [0,T]$,*

$$\phi_Q(t,x) = x + \frac{1}{2} v_Q(x) \left(e^{2t} - 1\right).$$

Proof. Notice that the $\phi_Q(t,x)$ is a point on the line segment connecting x and $\phi_Q(T,x)$. This direction of this segment is indicated by $v_Q(x)$. Thus, for simplicity we align the real line with $v_Q(x)$ in such a way that $d_Q(x)$ becomes the origin and x lies on the positive side of the real line. We indicate the distance to origin on this real line of the point $\phi_Q(t,x)$ by $\psi(t)$. Thus $\phi_Q(t,x)$ is related to $\psi(t)$ by the equation

$$\phi_Q(t,x) = d_Q(x) + \frac{v_Q(x)}{\|v_Q(x)\|} \psi(t). \tag{1}$$

We have

$$\begin{aligned} \psi(0) &= \|x - d_Q(x)\| \\ \dot{\psi}(t) &= 2\psi(t), \end{aligned}$$

where $\dot{\psi}$ denotes $d\psi(t)/dt$. This gives the differential equation

$$\frac{d\psi}{\psi} = 2dt,$$

which has the solution

$$\psi(t) = \psi(0) e^{2t}.$$

Replacing in (1) gives us

$$\begin{aligned} \phi_Q(t,x) &= d_Q(x) + \frac{v_Q(x)}{\|v_Q(x)\|} \cdot \|x - d_Q(x)\| e^{2t} \\ &= x - \frac{1}{2} v_Q(x) + \frac{v_Q(x)}{\|v_Q(x)\|} \cdot \frac{1}{2} v_Q(x) e^{2t} \\ &= x + \frac{1}{2} v_Q(x) \left(e^{2t} - 1\right). \end{aligned}$$

■

A vector field v defined on $\mathbb{R}^n$ is *k-semi-Lipschitz* if for every pair of vectors x and y

$$\langle x - y, v(x) - v(y) \rangle \leq k \cdot \|x - y\|^2.$$

Lemma 5.2 *Let Q be a set of weighted points in $\mathbb{R}^n$. Then the vector field v_Q is 2-semi-Lipschitz, i.e. for any $x, y \in \mathbb{R}^n$,*

$$\langle x - y, v_Q(x) - v_Q(y)\rangle \leq 2\|x - y\|^2.$$

Proof. Let p be a point in $A_Q(x)$ and $q \in A_Q(y)$. We show that

$$\langle x - y, p - q\rangle \geq 0. \tag{2}$$

Since $p \in A_Q(x)$, $\pi_p(x) \leq \pi_q(x)$ or equivalently

$$\|x - p\|^2 - \|x - q\|^2 \leq w_p - w_q.$$

Similarly since $q \in A_Q(y)$, $\pi_p(y) \geq \pi_q(y)$ and therefore

$$w_p - w_q \leq \|y - p\|^2 - \|y - q\|^2.$$

Eliminating $w_p - w_q$ from the above two inequality we obtain

$$\langle x - p, x - p\rangle - \langle x - q, x - q\rangle \leq \langle y - p, y - p\rangle - \langle y - q, y - q\rangle.$$

Expanding we get

$$\langle x, q\rangle - \langle x, p\rangle \geq \langle y, q\rangle - \langle y, p\rangle,$$

which gives (2) by rearranging.

Inequality (2) can be written also as

$$\langle x - y, p - y\rangle \geq \langle x - y, q - y\rangle. \tag{3}$$

The above inequality in particular holds for the point $p_0 \in A_Q(x)$ minimizing $\langle x - y, p_0 - y\rangle$ and the point $q_0 \in A_Q(y)$ maximizing $\langle x - y, q_0 - y\rangle$. Therefore the hyperplane $\mathit{\Pi} = \{z \in \mathbb{R}^n : \langle x - y, z - y\rangle = \gamma\}$, where $\gamma = \frac{1}{2}(\langle x - y, p_0 - y\rangle + \langle x - y, q_0 - y\rangle)$, separates the two sets $A_Q(x)$ and $A_Q(y)$ as well as their convex hulls and in particular $d_Q(x) \in \mathrm{conv}(A_Q(x))$ and $d_Q(y) \in \mathrm{conv}(A_Q(y))$. Thus we get

$$\langle x - y, d_Q(x) - d_Q(y)\rangle \geq 0. \tag{4}$$

Now to prove the statement of the lemma we write

$$\begin{aligned}
\langle x - y, v_Q(x) - v_Q(y)\rangle &= \langle x - y, 2(x - d_Q(x)) - 2(y - d_Q(y))\rangle \\
&= 2\langle x - y, (x - y) - (d_Q(x) - d_Q(y))\rangle \\
&= 2\|x - y\|^2 - 2\langle x - y, d_Q(x) - d_Q(y)\rangle \\
&\leq 2\|x - y\|^2,
\end{aligned}$$

where the final inequality follows (4). ■

Lemma 5.3 *Let Δ be an upper bound for the diameter of $\mathrm{conv}\, Q$. Then for any point x,*

$$\|v_Q(x)\| \leq 2(\Delta + \mathrm{dist}(x, \tilde{Q})).$$

Proof. By definition $v_Q(x) = 2(x - d_Q(x))$. Let p be a point in $\mathrm{conv}\, \tilde{Q}$. By triangle inequality $\|x - d_Q(x)\| \leq \|x - p\| + \|p - d_Q(x)\|$. Since $d_Q(x) \in \mathrm{conv}\, A_Q(x) \subset \mathrm{conv}\, \tilde{Q}$, $\|p - d_Q(x)\| \leq \Delta$. Thus $\|x - d_Q(x)\| \leq \Delta + \mathrm{dist}(x, \tilde{Q})$ and this implies the statement of the lemma. ■

Our main goal in this section is to prove that the flow ϕ_Q is continuous on its second variable. This, along with its continuity on its first variable, guarantees that the map $\phi_Q : \mathbb{R}^+ \times \mathbb{R}^n \to \mathbb{R}^n$ is a continuous map and can therefore be used for the role of map H in Proposition 8.1 for establishing homotopy equivalences.

Since the flow path $\phi_Q([0, T], x)$ is a piece-wise linear curve, we can study the flow in individual linear pieces of a flow path.

Lemma 5.4 *The map $t \mapsto \|v_Q(\phi(t,x))\|^2$ is the right-derivative of $t \mapsto h_Q(\phi_Q(t,x))$. In other words, for any x and $\varepsilon > 0$, there exists a $t_0 > 0$, such that for any $t \in [0, t_0]$,*

$$\left|h_Q(\phi_Q(t,x)) - \left(h_Q(x) - t\|v_Q(\phi_Q(t,x))\|^2\right)\right| < t\varepsilon. \tag{5}$$

Proof. Again it is enough to prove the lemma for the first linear piece of $\phi_Q(x)$. Let $y = \phi_Q(t,x)$ and let $\lambda = \|x - y\|$. Let p be a point in $A_Q(x)$ and in the hyper-plane Π containing $d_Q(x)$ and orthogonal to $v_Q(x)$. Notice that such a point must exist since by definition $d_Q(x)$ is the closest point to x on $\operatorname{conv} A_Q(x)$ and Π is a supporting hyper-plane orthogonal to $x - d_Q(x)$. Thus $\operatorname{conv} A_Q(x)$ must have a vertex on Π. Let $\alpha = \measuredangle(x - d_Q(x), x - p)$. Since y is in the same linear pieces of the flow line as x, $A_Q(x) = A_Q(y)$ and therefore $p \in A_Q(y)$. Using the cosine law on triangle xyp we get

$$\|y - p\|^2 = \|x - p\|^2 + \lambda^2 + 2\lambda\|x - p\| \cdot \cos\alpha.$$

Subtracting w_p from both sides gives

$$h_Q(y) = h_Q(x) + \lambda^2 + 2\lambda\|x - p\| \cdot \cos\alpha.$$

From Lemma 5.1, $\lambda = \frac{1}{2}\|v_Q(x)\|\gamma$ where $\gamma = e^{2t-1}$. Thus we have for the left hand side of the inquality (5)

$$\begin{aligned}
&\left|h_Q(y) - h_Q(x) + t\|v_Q(y)\|^2\right| \\
&\quad = \left|\lambda^2 + 2\lambda\|x - p\| \cdot \cos\alpha - t\|v_Q(y)\|^2\right| \\
&\quad = \left|\lambda^2 + 2\lambda\|x - p\| \cdot \cos\alpha - t(\|v_Q(x)\| + 2\lambda)^2\right| \\
&\quad = \left|\lambda^2 + 2\lambda\|x - p\| \cdot \cos\alpha - 4t\lambda^2 - t\|v_Q(x)\|^2 - 4t\lambda\|v_Q(x)\|\right| \\
&\quad = \left|\|v_Q(x)\|^2\left(\frac{1}{4}\gamma^2 - t\gamma^2 - 2t\gamma - t\right) + \|v_Q(x)\|\gamma\|x - p\|\cos\alpha\right|.
\end{aligned}$$

From our choice of p and the discussion above,

$$\|x - p\| \cdot \cos\alpha = \|x - d_Q(x)\| = \frac{1}{2}\|v_Q(x)\|.$$

Thus we have

$$\left|h_Q(y) - h_Q(x) + t\|v_Q(y)\|^2\right| \quad = \quad \|v_Q(x)\|^2\left|\left(\frac{1}{4} - t\right)\gamma^2 - 2t\gamma - t + \frac{1}{2}\gamma\right|$$

Therefore, for (5) to hold we must have for any $\varepsilon > 0$,

$$\|v_Q(x)\|^2\left|\left(\frac{1}{4} - t\right)\gamma^2 - 2t\gamma - t + \frac{1}{2}\gamma\right| < t\varepsilon,$$

when t is sufficiently small. Substituting $e^{2t} - 1$ for γ the above inequality becomes

$$\|v_Q(x)\|^2\left|\left(\frac{1}{4} - t\right)\left(e^{2t} - 1\right)^2 - 2t\left(e^{2t} - 1\right) - t + \frac{1}{2}\left(e^{2t} - 1\right)\right| < t\varepsilon,$$

which leads to the following inequality after simplification and rearrangement

$$\frac{1}{t}\left|\left(\frac{1}{4} - t\right)e^{4t} - \frac{1}{4}\right| < \frac{\varepsilon}{\|v_Q(x)\|^2}.$$

The left hand side of the above inequality is an increasing continuous function on $t \in (0, +\infty)$ which converges to 0 when $t \to^+ 0$ and this completes the proof of the Lemma.

■

The above Lemma immediately implies the following Corollary.

Corollary 5.5 *For any $x \in \mathbb{R}^n$ and any $t \in [0, +\infty)$,*

$$h_Q(\phi_Q(t,x)) = h_Q(x) + \int_0^t \|v_Q(\phi_Q(\tau,x))\|^2 d\tau.$$

In particular, the map $t \mapsto h_Q(\phi_Q(t,x))$ is strictly increasing.

Theorem 5.6 *The flow map ϕ_Q is continuous on its second variable. In other words for all $x \in \mathbb{R}^n$ and for all $t \geq 0$ and for all $\varepsilon > 0$, there exists a $\delta > 0$, such that for every $y \in \mathbb{R}^n$ satisfying $\|x-y\| < \delta$, $\|\phi_Q(t,x) - \phi_Q(t,y)\| < \varepsilon$.*

Proof. Consider two points x and y and real number t and let for each $0 \leq \tau \leq t$, $x(\tau) = \phi_Q(t,x)$ and $y(\tau) = \phi_Q(t,y)$. The two flow lines $\tau \mapsto x(\tau)$ and $\tau \mapsto y(\tau)$ for $\tau \in [0,t]$ are piecewise linear curves each with a finite number of turns. Let $0 = t_1 < t_2 < \cdots < t_r = t$ be the set of times at which at least one of these two curves makes a turn (enters a new Voronoi cell and switches drivers). We show that for each of the intervals $[t_i, t_{i+1}]$, $1 \leq i \leq r-1$,

$$\|x(t_{i+1}) - y(t_{i+1})\| \leq \|x(t_i) - y(t_i)\| \cdot e^{t_{i+1}-t_i}.$$

This will implies that

$$\|x(t_r) - y(t_r)\| \leq \|x(0) - y(0)\| \cdot \prod_{i=1}^{r-1} e^{t_{i+1}-t_i} = e^t \cdot \|x(0) - y(0)\|.$$

In particular, for any $\varepsilon > 0$, if $y \in B(x,\delta)$ where $\delta \leq \varepsilon / e^t$, then

$$\phi_Q(t,y) \in B(\phi_Q(t,x), \varepsilon).$$

It suffices to prove the above claim only for the first interval $[t_1, t_2]$; the claim will then hold for the subsequent intervals by a reparameterization of time. Equivalently, it will be enough to show that the claim is valid for the case where each of the two flow paths $x([0,t])$ and $y([0,t])$ is a single line segment.

For $0 \leq \tau \leq t$, we define the function $\lambda(\tau)$ as the square of the distance between $x(\tau)$ and $y(\tau)$, i.e.

$$\lambda(\tau) = \|x(\tau) - y(\tau)\|^2.$$

Then we have

$$\begin{aligned}
d\lambda(\tau) &= \lambda(\tau + d\tau) - \lambda(\tau) \\
&= \|x(\tau + d\tau) - y(\tau + d\tau)\|^2 - \|x(\tau) - y(\tau)\|^2 \\
&= \|x(\tau) + v_Q(x(\tau))d\tau - y(\tau) - v_Q(y(\tau))d\tau\|^2 - \|x(\tau) - y(\tau)\|^2 \\
&= 2\langle x(\tau) - y(\tau), v_Q(x(\tau)) - v_Q(y(\tau))\rangle d\tau + \|v_Q(x(\tau)) - v_Q(y(\tau))\|^2 (d\tau)^2 \\
&\leq 2\|x(\tau) - y(\tau)\|^2 d\tau + \|v_Q(x(\tau)) - v_Q(y(\tau))\|^2 (d\tau)^2,
\end{aligned}$$

in which the last inequality follows from Lemma 5.2. This results

$$\frac{d\lambda(\tau))}{d\tau} \leq 2\lambda(\tau) + \|v_Q(x(\tau)) - v_Q(y(\tau))\|^2 d\tau$$

Lemmas 5.3 and 5.4 bound $\|v_Q(x(\tau)) - v_Q(y(\tau))\|^2$ as a function of τ. Since $d\tau$ is infinitesimally small, we get the differential inequality

$$\frac{d\lambda(\tau)}{d\tau} \leq 2\lambda(\tau),$$

which for $0 \leq \tau \leq t$, yields to the solution

$$\|x(\tau) - y(\tau)\|^2 = \lambda(\tau) \leq \lambda(0) \cdot e^{2\tau}.$$

In particular for $\tau = t$,

$$\|x(t) - y(t)\| \leq e^t \cdot \|x(0) - y(0)\|.$$

∎

Corollary 5.7 *For any finite set Q of weighted points in $\mathbb{R}^n$, the induced flow map ϕ_Q is continuous on $[0, +\infty) \times \mathbb{R}^n$.*

6 Flow on Tubular Neighborhoods

As mentioned before, our topological proofs hinge upon using the continuous map ϕ_{Q_1} as a homotopy. The main result of this section is that the tubular neighborhood Σ_δ of Σ is tight for the flow ϕ_{Q_1} when δ is at a suitable range. This appears to be crucial in proving the homotopy equivalence results in Section 8.

The power of a point x with respect to a ball $B = B(c, R)$ is denoted by $\pi_B(x)$ and is defined as

$$\pi_B(x) := \|x - c\|^2 - R^2.$$

In other words, the power of a point x with respect to a ball of radius R centered at c is equivalent to the square of the distance between x and the point c with weight R^2. Thus $\pi_B(x)$ is positive outside B, zero on ∂B and negative inside B.

The following proposition is a well-known result on the structure of the Delaunay complex. The proof is included for completeness.

Proposition 6.1 *Let P be a set of points in $\mathbb{R}^n$ and let x be point in a Delaunay n-simplex $\tau \in \mathrm{Del}\, P$. Let B be the ball circumscribing τ and let B' be an arbitrary empty ball. Then $\pi_B(x) \leq \pi_{B'}(x)$. In other words, of all the empty balls, B is the one with respect to which the power of x is the smallest.*

Proof of Proposition 6.1. Consider a mapping from $\mathbb{R}^n \to \mathbb{R}^{n+1}$ given by $x \mapsto x^*$ where

$$x^* = \left(x_1, \ldots, x_n, \sum_{i=1}^{n} x_i^2\right),$$

when $x = (x_1, \ldots, x_n)$. In other words, every point is lifted to the standard paraboloid with equation $x_n = x_1^2 + \cdots + x_n^2$. It can be verified that lifted image of a sphere ρ in $\mathbb{R}^n$ lies in a hyperplane H_ρ in $\mathbb{R}^{n+1}$ and a sphere ρ in $\mathbb{R}^n$ is empty if and only if H_ρ is below the image p^* of every $p \in P$. Furthermore, given a sphere ρ and a point x, $\pi_\rho(x)$ is the signed vertical distance between x^* and H_ρ. So, for a given point x, the empty ball B containing x that minimizes $\pi_B(x)$ must have $H_{\partial B}$ below every p^* with $p \in P$ but must vertically be as far away as possible above x^*. This makes $H_{\partial B}$ a supporting hyperplane of the lower hull of $P^* = \{p^* : p \in P\}$ and thus B must be a Delaunay ball. ■

Lemma 6.2 *Consider an ε-sample of the surface Σ_1 with $\varepsilon \leq 1/10$. For a $0 < \delta < 1$, let x be a point on the boundary of Σ_δ, i.e. $\|x - \hat{x}\|/f(\hat{x}) = \delta$ with $\hat{x} \in \Sigma$. Let τ be a Delaunay n-simplex in $\mathrm{Del}\, P_1$ that contains x and let c be the circumcenter of τ. If c is at the same side of Σ as x, then for the angle $\alpha = \measuredangle(c - x, x - \hat{x})$, $\cos\alpha \geq 1 - \delta - 6\varepsilon$.*

Proof of Lemma 6.2. Let u and v be the two points on the line normal to Σ at $\hat{x}$ and at distance $f(\hat{x})$ from it, with u being on the same side of Σ as x. The balls $B_u = B(u, f(\hat{x}))$ and $B_v = B(v, f(\hat{x}))$ are disjoint from Σ and their boundaries are tangent to it.

Consider the line segment $\overline{cv}$. Since we assumed c is in the same side of Σ as x (and u), $\overline{cv}$ intersects Σ in at least one point. Let z be an arbitrary point in $\overline{cv} \cap \Sigma$. Let us denote $\|c - v\|$ by d and $\|z - v\|$ by b. Notice that $d \geq b \geq f(\hat{x})$.

Since the local feature size function is 1-Lipschitz, using the triangle inequality we get

$$f(z) \leq \|z - \hat{x}\| + f(\hat{x}) \leq \|z - v\| + \|v - \hat{x}\| + f(\hat{x}) = b + 2f(\hat{x}).$$

By the ε-sampling condition, there must be a sample point in P within $\varepsilon f(z) \leq \varepsilon(b + 2f(\hat{x}))$ from z. The ball $B_c = B(c, R)$ circumscribing τ does not intersect P and therefore, $\|z - c\| + \varepsilon f(z) \geq R$,

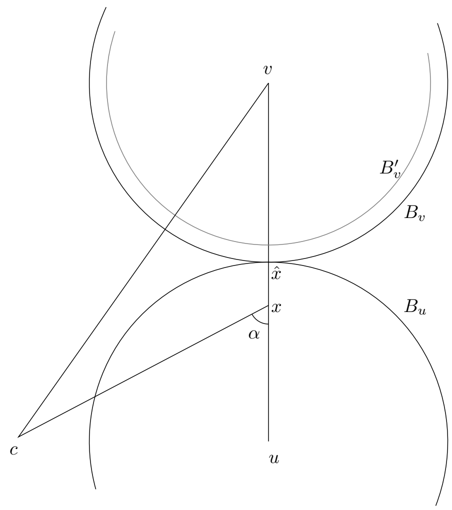

Figure 6: Proof of Lemma 6.2.

or else z does not meet the sampling condition. Using $\|z-c\| = d-b$ and the above upper bound for $f(z)$, we get

$$\begin{aligned} R &\leq d-b+\varepsilon(b+2f(\hat{x})) \\ &= d-(1-\varepsilon)b+2\varepsilon f(\hat{x}) \\ &\leq d-(1-\varepsilon)f(\hat{x})+2\varepsilon f(\hat{x}) \\ &= d-f(\hat{x})+3\varepsilon f(\hat{x}). \end{aligned} \tag{6}$$

Consequently if B'_v denote the ball centered at v with radius $(1-3\varepsilon)f(\hat{x})$, then $B_c \cap B'_v = \emptyset$.

On the other, using Proposition 6.1 for Delaunay ball B and empty ball B_u we can write

$$\|x-c\|^2 - R^2 \leq (1-\delta)^2 f(\hat{x})^2 - f(\hat{x})^2. \tag{7}$$

Using the cosine rule on the triangle cxv for angle $\measuredangle(x-c, x-v) = \pi - \alpha$ we obtain

$$\cos(\pi-\alpha) = \frac{\|x-v\|^2 + \|x-c\|^2 - \|c-v\|^2}{2\|x-v\|\|x-c\|} = \frac{(1+\delta)^2 f(\hat{x})^2 + \|x-c\|^2 - \|c-v\|^2}{2(1+\delta)f(\hat{x})\|x-c\|}.$$

Combining this with inequalities (6) and (7), and defining $r = R/f(\hat{x})$ results

$$\cos\alpha \geq \frac{2r(1-3\varepsilon) - 6\varepsilon - 2\delta^2 + 9\varepsilon^2}{2(1+\delta)\sqrt{r^2+(1-\delta)^2-1}}.$$

The right hand side of the of the above inequality is a function of r (taking ε and δ as constants) defined for $r^2 > 1-(1-\delta)^2$ (notice that $r = R/f(\hat{x})$ is always positive). It can be verified that its derivative has a unique root corresponding to a global minimum. Calculating the value of the function at this minimum gives us

$$\cos\alpha \geq \frac{-2\delta^2 - 6\varepsilon + 9\varepsilon^2 + \frac{4(1-3\varepsilon)(2-9\varepsilon+9\varepsilon^2)}{2+3\varepsilon}}{2(1+\delta)\sqrt{-1+(1-\delta)^2+\frac{4(2-9\varepsilon+9\varepsilon^2)^2}{(2+3\varepsilon)^2}}}$$

Elementary algebraic simplifications assuming that $\varepsilon \leq 1/10$ entails the statement of the Lemma.
■

Let us observe that if the sample considered in Lemma 6.2 is a uniform ε-sample for $\varepsilon \leq \frac{1}{10} f_{\min}$, one can redo the proof at every point x at distance $\delta \cdot f_{\min}$ from Σ by assuming $f(\hat{x}) = f_{\min}$. Thus the above result for uniform samples can be states as follows.

Lemma 6.3 *Let P be a uniform ε-sample of the surface Σ_1 with $\varepsilon \leq 1/10 \cdot f_{\min}$. For a $0 < \delta < 1$, let x be a point on the boundary of Σ_δ, i.e. $\|x - \hat{x}\| = \delta \cdot f_{\min}$ with $\hat{x} \in \Sigma$. Let τ be a Delaunay n-simplex in $\mathrm{Del}\, P_1$ that contains x and let c be the circumcenter of τ. If c is at the same side of Σ as x, then for the angle $\alpha = \measuredangle(c - x, x - \hat{x})$, $\cos\alpha \geq 1 - \delta - 6\varepsilon$.*

Lemma 6.4 *Let P_1 be an ε-sample of the surface Σ_1 with $\varepsilon < 1/3$. Let x be a point on the boundary of Σ_δ for $\delta > 9\varepsilon^2$. Finally, let τ be a Delaunay n-simplex in $\mathrm{Del}\, P_1$ containing x and let $B = B(c, R)$ be the circumsphere of τ. Then c and x are on the same side of Σ.*

Proof of Lemma 6.4. Suppose to the contrary that x and c are in opposite sides of Σ. Let u be the point on normal to Σ at $\hat{x}$ at the same side of Σ as x, satisfying $\|\hat{x} - u\| = f(\hat{x})$. Since c is on the opposite side of Σ from x (and u), the segment $\overline{cu}$ must intersect Σ at a some point z. Let $d = \|c - u\|$ and $b = \|z - u\| \geq f(\hat{x})$. Using the 1-Lipschitzness of local feature size, we can bound $f(z)$ as

$$f(z) \leq \|z - \hat{x}\| + f(\hat{x}) \leq \|z - u\| + \|u - \hat{x}\| + f(\hat{x}) = b + 2f(\hat{x}).$$

Thus z must have a sample point within $\varepsilon f(z) \leq \varepsilon(b + 2f(\hat{x}))$ in P. This puts an upper bound on the radius R of the empty ball B:

$$R \leq \|c - z\| + \varepsilon f(z) \leq d - b + \varepsilon(b + 2f(\hat{x})).$$

Since $b \geq f(\hat{x})$ we get:

$$R \leq (d - 1 + 3\varepsilon) f(\hat{x}). \tag{8}$$

On the other hand, since x is in τ and $B_u = B(u, f(\hat{x}))$ is empty, $\pi_{B_c}(x) \leq \pi_{B_u}(x)$, i.e.

$$\|x - c\|^2 - R^2 \leq \|x - u\|^2 - f(\hat{x})^2, \tag{9}$$

Using $\|x - u\| = (1 - \delta) f(\hat{x})$ and the cosine rule on the triangle cux we can write $\|x - c\|^2$ as

$$\|c - x\|^2 = (1 - \delta)^2 f(\hat{x})^2 + d^2 - 2d(1 - \delta) f(\hat{x}) \cos\alpha,$$

where $\alpha = \measuredangle(u - c, u - x)$. Combining with Equation (9) and using $\cos\alpha \leq 1$ gives us

$$R^2 \geq d^2 - 2d(1 - \delta) f(\hat{x}) + f(\hat{x})^2. \tag{10}$$

Combining Equations (8) and (10) using using the fact that $d \geq f(\hat{x})$, implies $\delta \leq 9\varepsilon^2$, which is a contradiction. ■

Lemma 6.5 *Let x be a point on the boundary of Σ_δ for $0 < \delta < 1$ and let v be a vector satisfying $\tan\alpha \leq \frac{1-\delta}{2\delta}$, where $\alpha = \measuredangle(x - \hat{x}, v)$. Then there is a real number $t_0 > 0$, such that $x + vt \in \Sigma_\delta$ for every $0 \leq t \leq t_0$.*

Proof of Lemma 6.5. For simplicity we take $f(\hat{x})$ as unit length. Let c be a point between $\hat{x}$ and $\check{x}$ satisfying $\|c - \hat{x}\| = f(\hat{x}) = 1$. Consider a point $y = x + tv$, close enough to x so that $y \in B_c = B(c, 1)$. Let θ represent the angle $\measuredangle(c - x, c - y)$. It is easy to see that the parameter t in the statement of the lemma can be replaced with the angle θ corresponding to t. Indeed, we prove that there exists a θ_0 such that all points on the segment xy_0 where y_0 represents the point corresponding to angle θ_0, are all in Σ_δ.

Since $\hat{y}$ can be no farther away from y that $\hat{x}$, $\hat{y}$ must be inside the ball $B_y := B(y, \|y - \hat{x}\|)$. Considering the fact that B_c is disjoint from Σ, $\|\hat{x} - \hat{y}\|$ cannot be larger than the diameter of the spherical cap $\partial B_y \setminus B_c$. Thus $\|\hat{x} - \hat{y}\| \leq 2\sin\theta$ and therefore,

$$f(\hat{y}) \geq f(\hat{x}) - \|\hat{x} - \hat{y}\| \geq 1 - 2\sin\theta. \tag{11}$$

Let us denote $\|c - y\|$ by ℓ. Also, assume without loss of generality that v is a unit vector and therefore $\|x - y\| = t$. From the sine law on the triangle cxy, we have

$$\ell = (1 - \delta)\sin\alpha / \sin(\alpha - \theta).$$

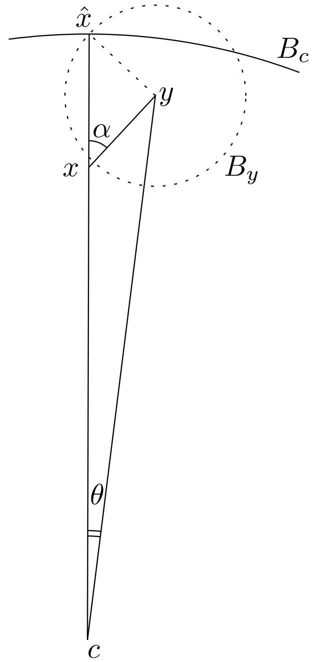

Figure 7: Proof of Lemma 6.5.

Since $\|y-\hat{y}\| \leq \|y-\hat{x}\|$, sing the cosine law on triangle $cy\hat{x}$ we have

$$\begin{aligned} \|y-\hat{y}\|^2 &\leq \|y-\hat{x}\|^2 \\ &= 1+\ell^2-2\ell\cos\theta \\ &= 1+\frac{(1-\delta)^2\sin^2\alpha}{\sin^2(\alpha-\theta)}-\frac{2(1-\delta)\sin\alpha\cos\theta}{\sin(\alpha-\theta)}. \end{aligned} \tag{12}$$

Let us denote the right hand side of (12) as $g(\theta)$ (taking α and δ as constants). Thus from (11) and (12) we have

$$\frac{\|y-\hat{y}\|^2}{f(\hat{y})^2} \leq \frac{g(\theta)}{(1-2\sin\theta)^2}.$$

The statement of the Lemma follows if the function $h(\theta) := g(\theta)/(1-2\sin\theta)^2 \leq \delta^2$ when θ belongs to some interval $[0,\theta_0)$. However, since $h(0)=\delta^2$, this amounts to verifying that $(dh/d\theta)(0) \leq 0$ and $(d^2h/d\theta^2)(0) < 0$ when $(dh/d\theta)(0) = 0$. These claims can be verified algebraically when the specified bound on α is applied. ■

Corollary 6.6 *For $\varepsilon \leq 1/10$ and $9\varepsilon^2 < \delta \leq 3/10 - 2\varepsilon$, no flow line of ϕ_{Q_1} leaves the δ-tubular neighborhood Σ_δ. In other words, Σ_δ is flow-tight for ϕ_Q.*

Proof of Corollary 6.6. By Lemmas 6.5 and 6.4, it suffices choose δ in such a way that $\beta = \measuredangle(\hat{x}-x, v_Q(x))$ is smaller than the angle α required in Lemma 6.2. In other words, it suffices to have

$$\tan\beta \leq \frac{1-\delta}{2\delta},$$

or equivalently

$$\frac{1}{\cos^2\beta} \leq 1+\frac{(1-\delta)^2}{4\delta^2}.$$

By Lemma 6.2, $\cos\beta \geq 1 - \delta - 6\varepsilon$. Thus it suffices to choose δ in such a way that

$$\frac{1}{(1-\delta-6\varepsilon)^2} \leq 1 + \frac{(1-\delta)^2}{4\delta^2}.$$

It can be verified that for $\varepsilon \leq 1/10$, the inequality is enforced when $9\varepsilon^2 < \delta \leq 3/10 - 2\varepsilon$. ■

7 Geometric Quality

The purpose of this section is to prove that the set O of simplices removed from $\text{Del}\, P$ by the WRAP algorithm advances close to the actual surface Σ. In particular, this entails that the symmetric difference between the output I of WRAP and the original shape S is contained in the tubular neighborhood $\Sigma_{9\varepsilon^2}$. One use of this result is proving a gap between the points in I and exterior medial axis critical points which is important in our homotopy proofs.

The following Lemma is close in spirit to Theorem 2.1. The proofs and the analyses are closely related but the analysis in the following lemma is slicker. In essence, it shows that if the Voronoi face dual to a Delaunay simplex τ is not entirely far from the surface, then τ is entirely very close to the surface.

Tangent Feature-size Balls

For a point x on Σ, B_x^+ and B_x^- respectively denote the outer and inner balls of radius $f(x)$ tangent to Σ at x.

Lemma 7.1 *Let x be a point on the boundary of Σ_λ for $0 < \lambda < 1$. Define the set $L(x) \subset \mathbb{R}^n$ as*

$$L(x) := B\left(x, \ell(\varepsilon,\lambda) f(\hat{x})\right) \setminus \left(B_{\hat{x}}^+ \cup B_{\hat{x}}^-\right),$$

where

$$\ell(\varepsilon,\lambda) = \sqrt{\lambda^2 + \varepsilon^2(1+\lambda)}.$$

Then,

1. *$A_{P_1}(x) = A_P(x) \subset L(x)$. In particular, the points in $A_P(x)$ are within distance $\ell(\varepsilon,\lambda) f(\hat{x})$ from x.*
2. *$L(x) \subset B\left(x, r(\varepsilon,\lambda) f(\hat{x})\right)$ where*

$$r(\varepsilon,\lambda) = \varepsilon\sqrt{\frac{1+\lambda}{1-\lambda}}.$$

3. *$L(x)$ is contained in a cone with apex x, axis $\hat{x} - x$, and half-angle γ where*

$$\psi = \psi(\varepsilon,\lambda) = \arcsin\left(\frac{r(\varepsilon,\lambda)}{\ell(\varepsilon,\lambda)}\right).$$

Proof. Let c^+ be a point on the line normal to Σ at $\hat{x}$ at the same side of Σ as x and at distance $f(\hat{x})$ from $\hat{x}$. Let c^- be the point symmetrically opposite to c^+ with respect to $\hat{x}$. For simplicity we normalize the distances so that $f(\hat{x}) = 1$. Let $B^+ = B(c^+, 1)$ and $B^- = B(c^-, 1)$ be balls of radius 1, i.e. $f(\hat{x})$, tangent to Σ at $\hat{x}$. By the definition of local feature size, these balls avoid Σ. By the ε-sampling assumption, there is a sample point in the ball $B_\varepsilon = B(\hat{x}, \varepsilon)$. Thus the closest point in P to x is within distance ℓ from x, where ℓ is the distance between x and y, a farthest point from p in $L_0 = B_\varepsilon \setminus (B^+ \cup B^-)$. Figure 8 shows a planar section of this setting. In the figure, the region L_0 is shaded with the darkest gray. Let $B_\ell = B(x, \ell)$. All closest points to x in P must lie in the region $L_1 = L(x) = B_\ell \setminus (B^+ \cup B^-)$. Let z be a point in this region farthest away from $\hat{x}$ and let $r = \|\hat{x} - z\|$. Let γ be the angle between $y - \hat{x}$ and the hyper-plane tangent to Σ at $\hat{x}$. It can be easily seen from the right hand side of Figure 8 that $\sin\gamma = \varepsilon/2$.

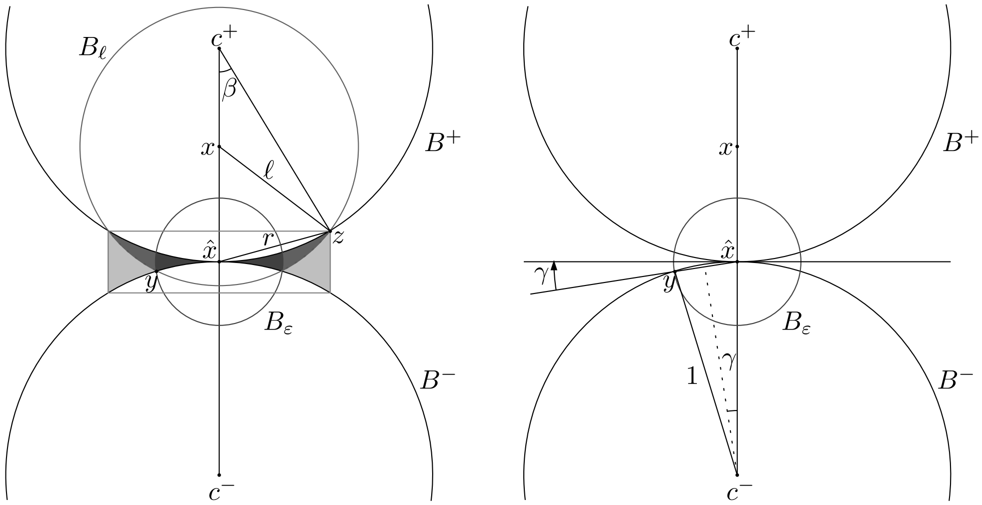

Figure 8: Proof of Lemma 7.2

Since the angle $\angle(y - \hat{x}, x - \hat{x})$ is $\pi/2 + \gamma$ we have using the cosine rule

$$\begin{aligned} \ell^2 = \|x - z\|^2 = \|x - y\|^2 &= \lambda^2 + \varepsilon^2 - 2\varepsilon\lambda \cos(\pi/2 + \gamma) \\ &= \lambda^2 + \varepsilon^2 - 2\varepsilon\lambda \sin \gamma \\ &= \lambda^2 + \varepsilon^2(1 + \lambda). \end{aligned}$$

Now by applying the cosine rule to the triangle xzc^+, we have for the angle $\beta = \angle(x - c^+, z - c^+)$

$$\begin{aligned} \cos \beta &= \frac{1 + (1 - \lambda)^2 - \ell^2}{2(1 - \lambda)} \\ &= \frac{1 + (1 - \lambda)^2 - \lambda^2 - \varepsilon^2(1 + \lambda)}{2(1 - \lambda)} \\ &= 1 - \frac{\varepsilon^2}{2} \cdot \frac{1 + \lambda}{1 - \lambda}. \end{aligned}$$

If we rewrite the above equality

$$\cos \beta = 1 - 2 \left(\frac{\varepsilon}{2} \sqrt{\frac{1 + \lambda}{1 - \lambda}} \right)^2,$$

and observe on the figure that $\sin(\beta/2) = r/2$, we can use the identity $\cos \beta = 1 - 2\sin^2(\beta/2)$ to obtain,

$$r = \varepsilon \cdot \sqrt{\frac{1 + \lambda}{1 - \lambda}}.$$

To complete the proof, we need to only show that the angle $\beta' = \measuredangle(\hat{x} - x, z - x)$ is smaller than $\psi(\varepsilon, \delta)$ as given in the statement of the Lemma. From the figure $\sin \beta' = h/\ell$ where h is the distance between z and the line supporting the segment $x\hat{x}$. Since $h \leq r$, $\sin \beta' \leq r/l = \sin \psi$. ■

Lemma 7.2 *Let P_1 be an ε-sample of the surface Σ_1 for $\varepsilon \leq 1/10$. For a simplex τ in $\operatorname{Del} P_1$, if the face of $\operatorname{Vor} P_1$ dual to τ intersects Σ_λ for $\lambda \leq 1/2$, then the simplex τ is contained entirely in Σ_δ for $\delta \geq \frac{5}{2}\varepsilon^2$.*

Proof of Lemma 7.2. We use the setting of the proof of Lemma 7.1 (refer to Figure 8). Let $r = r(\varepsilon, \lambda)$ and $\ell = \ell(\varepsilon, \lambda)$. Consider the hyper-plane H^+ passing through z and parallel to the tangent plane to Σ at $\hat{x}$. Let H^- be the hyper-plane parallel to H^+ and symmetric to it with respect to $\hat{x}$. The hyper-planes H^+ and H^- intersect the balls B^+ and B^- respectively in two similar disks C^+ and C^- both centered on the segment c^+c^-. The region $L(x)$ is contained in the convex hull of $C^+ \cup C^-$ which is a cylinder L_2 whose central axis lies on c^+c^-. Since L_2 and τ are both convex, $\tau \subset L_2$. The important observation is that since the disks C^+ and C^- are respectively contained in the balls B^+ and B^- in opposite sides of Σ, every segment parallel to c^+c^- connecting a point from C^+ to its corresponding point in C^- must intersect Σ. This is in particular true for those of such segments that intersect τ. Thus the length of these segments is an upper bound on the distance between any point in τ and Σ. This length, i.e. the distance between C^+ and C^- is simply $2(1 - \cos\beta)$. Thus for every point $u \in \tau$:

$$\|u - \hat{u}\| \leq 2(1 - \cos\beta) = \varepsilon^2 \cdot \frac{1+\lambda}{1-\lambda}.$$

On the other hand by the triangle inequality

$$\|\hat{u} - \hat{x}\| \leq \|\hat{u} - u\| + \|u - x\| \leq \|\hat{u} - u\| + r \leq \varepsilon^2 \cdot \frac{1+\lambda}{1-\lambda} + \varepsilon \cdot \sqrt{\frac{1+\lambda}{1-\lambda}}.$$

Since the local feature size is a 1-Lipschitz function

$$f(\hat{u}) \geq f(\hat{x}) - \|\hat{x} - \hat{u}\| \geq 1 - \varepsilon^2 \cdot \frac{1+\lambda}{1-\lambda} + \varepsilon \cdot \sqrt{\frac{1+\lambda}{1-\lambda}}.$$

Therefore

$$\frac{\|u - \hat{u}\|}{f(\hat{u})} \leq \frac{\varepsilon^2 \cdot \frac{1+\lambda}{1-\lambda}}{1 - \varepsilon \cdot \sqrt{\frac{1+\lambda}{1-\lambda}} - \varepsilon^2 \cdot \frac{1+\lambda}{1-\lambda}} \leq \frac{5}{2}\varepsilon^2,$$

for $\lambda < 1/2$ and $\varepsilon \leq 1/10$. ■

Lemma 7.3 *Let $\lambda < 1/2$ and let $\mathcal{D}$ be the subcomplex of $\operatorname{Del} P_1$ consisting of all n-simplices whose circumcenters are contaiend in S^*_λ, along with all proper faces of such simplices. Then $|\mathcal{D}|$ covers $\operatorname{conv} P_1 \cap S^*_\delta$ for $\delta > 9\varepsilon^2$.*

Proof of Lemma 7.3. Every point $x \in \operatorname{conv} P_1$ is in at least one Delaunay n-simplex. To prove the lemma, we show that for every $x \in \operatorname{conv} P_1 \cap S^*_\delta$, the circumcenter of every n-simplex $\tau \in \operatorname{Del} P_1$ containing x is in S^*_λ which implies that $\tau \in \mathcal{D}$.

Thus assume $x \in \operatorname{conv} P_1 \cap S^*_\delta$ and let τ be a Delaunay n-simplex that contains x. Since $x \in S^*_\delta$, $\|x - \hat{x}\| > 9\varepsilon^2 f(\hat{x})$. Thus by Lemma 6.4, the circumcenter c of τ is at the same side of Σ as x and therefore $c \in S^*$. Now, by Lemma 7.2, $\|c - \hat{c}\| \geq \frac{1}{2} f(\hat{c})$ since otherwise $\tau \subset \Sigma_{2.5\varepsilon^2}$ while $x \in \tau$ is in $S^*_{9\varepsilon^2}$, a contradiction. Thus, $c \in S^*_\lambda$ and $\tau \in \mathcal{D}$ and therefore the point x is covered by $|\mathcal{D}|$. Since this argument holds for every $x \in \operatorname{conv} P_1 \cap S^*_\delta$, the proof is complete. ■

We will need the following two lemmas from [1] and [9] respectively.

Lemma 7.4 *Let x and y be points on Σ with $\|x - y\| \leq \xi f(x)$ for $\xi \leq 1/3$. Then $\angle(n^+_x, n^+_y) \leq \frac{\xi}{1-3\xi}$. Likewise $\angle(n^-_x, n^-_y) \leq \frac{\xi}{1-3\xi}$.*

Lemma 7.5 *For points $p, q, r \in \Sigma$, let p be a vertex of the triangle τ with the largest angle. If the circumradius of τ is ρ, then the angle between n_τ, the normal to the plane of τ, and n_p, the normal to Σ at p, is at most $\beta(\rho/f(p))$ where*

$$\beta(\lambda) := \arcsin(\lambda) + \arcsin\left(\frac{2}{\sqrt{3}} \sin(2 \arcsin \lambda)\right).$$

In particular, when $\lambda \leq 1/4$, $\beta(\lambda) \leq 4\lambda$.

Orienting Triangle Normals

In general, for a triangle τ let n_τ denote the direction of the line normal to the plane of τ. It is possible to orient n_τ according to vector n_p^+ for a vertex p of τ. Thus of the two vectors parallel to the direction n_τ, we take the one with angle smaller than $\pi/2$ from n_p^+ as n_τ^+ and the other one as n_τ^-. It is of course generally possible to get a conflicting orientation for n_τ when we repeat this using a different vertex of τ. We call a Delaunay triangle τ *flat* (to the surface) if the orientations of n_τ with respect to surface normals at all three of its vertices are consistent. For flat triangles n_τ^+ and n_τ^- are well-defined without reference to any particular vertex. Next we show that all Delaunay simplices whose dual intersect S_λ^* for a large enough λ, are put in O by WRAP. This entails that WRAP progresses in removing simplices from Del P (and putting them in O) and reaches a close neighborhood of the surface. The proof of the main result of this section, i.e. Theorem 7.8 depends on the following lemma the proof of which makes use of of Lemma 7.5 from [9] which states that a Delaunay triangles with small circumradius is almost tangent to surface at its vertices. This lemma is only proven for $\mathbb{R}^3$ and is believed not to generalize to higher dimensions. Because of this, the guarantees provided in this paper only hold in three dimensions. However, this is the only weak link in the provided chain of arguments and a proof of the following lemma for arbitrary dimensions generalizes all the guarantees in this paper.

Lemma 7.6 *Let $\varepsilon \leq 0.03$ and $\delta = \varepsilon$. Let $\mathcal{D}$ be as defined in Lemma 7.3 for $\lambda = \delta = \varepsilon$. Let τ be a Delaunay triangle on the boundary of $|\mathcal{D}|$ and let v_i be the vertex of the Voronoi edge dual to τ that is contained in S_δ^*. Furthermore, let x be any intersection point of e and the outer boundary of Σ_δ. Then (1) τ is flat, (2) for every point $y \in \tau$, the angle between the vectors $x - y$ and n_τ^+ (n_τ^-), is at most*

$$\frac{r}{1-3r} + \beta\left(\frac{\ell}{1-r}\right) + \arcsin\left(\frac{r}{\ell}\right),$$

where $r = r(\varepsilon, \delta)$ and $\ell = \ell(\varepsilon, \delta)$ are defined in Lemma 7.1. Finally (3) v_i is on the same side of the plane of τ as x and is farther from this plane than x.

Proof. For simplicity, we assume that $f(\hat{x}) = 1$. By Lemma 7.1 $h_P(x) \leq \ell = \ell(\varepsilon, \delta)$. Since x is at equal distance from the vertices of τ, $h_P(x)$ is an upper bound for the circumradius of τ. Moreover, by Lemma 7.1, $\|\hat{x} - p\| \leq r = r(\varepsilon, \delta)$ which results, using Lipschitz property of the local feature size, that $f(p) \geq 1 - r$. Therefore, by Lemma 7.5

$$\measuredangle(n_\tau, n_p) \leq \beta\left(\frac{\ell}{1-r}\right). \tag{13}$$

Let n_τ^+ be the orientation of n_τ that makes an acute angle with n_p^+.

For any vertex q of τ, $\|q - \hat{x}\| \leq r$ and therefore using the Lemma 7.4 we obtain

$$\measuredangle(n_q^+, n_{\hat{x}}^+) \leq \frac{r}{1-3r}. \tag{14}$$

Since (14) holds for every vertex of τ, it also holds for p and therefore

$$\measuredangle(n_p^+, n_q^+) \leq \measuredangle(n_p^+, n_{\hat{x}}^+) + \measuredangle(n_q^+, n_{\hat{x}}^+) \leq \frac{2r}{1-3r}.$$

Thus we get

$$\measuredangle(n_\tau^+, n_q^+) \leq \beta\left(\frac{\ell}{1-r}\right) + \frac{2r}{1-3r} \leq 13^\circ,$$

for our choices of ε and δ. Thus n_τ^+ and n_τ^- are well-defined and τ is flat.

For a point $y \in \tau$, let $\alpha(y)$ be the angle between n_τ^+ and the vector $x - y$. Let x_0 be the intersection point of the affine hulls of e and τ. Since e and τ are orthogonal, $\cot \alpha(y)$ is exactly $\|x-x_0\|/\|y-x_0\|$. Thus $\alpha(y)$ depends only on the distance between y and x_0 and grows monotonically when $\|y-x_0\|$ grows. Therefore, over τ, $\alpha(y)$ achieves its maximum at every one of the three vertices of τ.

Therefore, it suffices to prove the statement of the lemma only for $y = p$ where p is the vertex of τ with the largest face angle in τ. Equations (13) and (14) put a bound on angle between n_τ^+ and $n_{\hat{x}}$:

$$\measuredangle(n_\tau^+, n_{\hat{x}}^+) \leq \frac{r}{1-3r} + \beta\left(\frac{\ell}{1-r}\right) < 11°,$$

for $\delta = \varepsilon \leq 0.03$. By Lemma 7.1 the angle between $x-p$ and $n_{\hat{x}}$ is at most

$$\measuredangle(n_{\hat{x}}^+, x-p) \leq \arcsin\left(\frac{r}{\ell}\right).$$

Combining, we get

$$\measuredangle(n_\tau^+, x-p) \leq \measuredangle(n_\tau^+, n_{\hat{x}}^+) + \measuredangle(n_{\hat{x}}^+, x-p) \leq \frac{r}{1-3r} + \beta\left(\frac{\ell}{1-r}\right) + \arcsin\left(\frac{r}{\ell}\right).$$

Next we show that x and v_i are on the same side of the plane Π of τ as x. First observe that the above argument remains valid for every intersection point x of e and the outer boundary of Σ_δ. Since for any such x, the angle between $x-p$ and n_τ^+ is less than $90°$, all these points are on the same side of Π. Among all such intersection points, we take x to be the one closest to v_i and prove that v_i is on the same side of Π as x. First notice that since $v_i \in S_\delta^*$ and x is the closest point to v_i on the outer boundary of Σ^δ, the segment xv_i is entirely contained in S_δ^*. On the other hand, as shown above, the angle between the segment xv_i (parallel to n_τ) and $n_{\hat{x}}$ is at most

$$\frac{r}{1-3r} + \beta\left(\frac{\ell}{1-r}\right).$$

For our choices of ε and δ, this is a smaller angle than $\theta = \arctan\left(\frac{1-\delta}{2\delta}\right)$ and therefore the segment xv_i falls inside the double-cone with apex x, axis $n_{\hat{x}}$ and half angle $\arctan\theta$. On the other hand, by Lemma 6.5, xv_i cannot be inside the cone opening toward $n_{\hat{x}}^-$ since every segment xz in this cone must have an initial segment xz' in Σ_δ while xv_i is entirely in S_δ^*. Thus v_i has to be in the cone opening toward $n_{\hat{x}}^+$. Since we showed above that $\measuredangle(n_{\hat{x}}^+, n_\tau^+)$ differ by at most $11°$, v_i has to be on the same side of Π as x and farther away from Π than x. ■

Lemma 7.7 *Let P_1 be an ε-sample of a surface Σ_1 in $\mathbb{R}^3$ for $\varepsilon \leq 0.03$. Let $\mathcal{D}$ be as defined in Lemma 7.3. Then no flow line of ϕ_{Q_1} starting from a simplex in $\mathrm{Del}\, P_1 \setminus \mathcal{D}$ enters the interior of $|\mathcal{D}|$.*

Proof. We prove the theorem by analyzing the flow direction on the boundary of $|\mathcal{D}|$. Specifically, for every point x on a Delaunay triangle or edge on $\partial|\mathcal{D}|$, we show that $v_Q(x)$ i's either tangent to $\partial|\mathcal{D}|$ or points toward its exterior. Notice that Delaunay vertices, i.e. points in P, are maxima (and thus sinks) of ϕ_Q and thus need not to be analyzed.

Let τ be a Delaunay triangle on $\partial|\mathcal{D}|$. By definition the Voronoi edge e dual to τ has a vertex v_o in Σ_δ and the other v_i in S_δ^* thus intersecting the outer boundary of Σ_δ at some point x. Pick x arbitrarily if there are multiple such intersections. By Lemma 7.6, τ is flat. Take the plane Π of τ as the horizontal plane and the direction of e as vertical. Of the two Delaunay tetrahedra incident to τ, let σ_i be the one included in $\mathcal{D}$ and let σ_o be the other one. Clearly, σ_i is dual to the Voronoi vertex v_i and σ_o is dual to v_o. If v_i and v_o are on opposite sides of Π, $d_Q(\tau)$, the closest point to σ on e will be the intersection of the affine hulls of τ and e and the flow on τ will be tangential in which case there is nothing to prove. Thus assume that v_i and v_o are on the same side of Π.

The three points v_i, x and v_o are on a line orthogonal to Π with x between v_i and v_o. The vertical order of σ_i and σ_o (in the direction of e) agrees with that of v_i and v_o. Thus the flow on τ is toward σ_o if and only if v_i is farther from Π than x. Lemma 7.6 guarantees that this is indeed the case.

Next, consider a Delaunay edge e on the boundary of $|\mathcal{D}|$ and let τ and τ' be two Delaunay triangles incident to e and on the boundary of $|\mathcal{D}|$. Let H and H' be half-planes, respectively supporting τ and τ', and sharing the line through e for boundary. Let W be the set of all tetrahedra incident to e but not in $\mathcal{D}$ and between τ and τ'. $|W|$ is a polytope that has e, τ, and τ' on its boundary. Let ϕ be the dihedral angle between τ and τ' measured from inside W. We will also refer to wedge made by τ and τ' and contained locally inside W as W.

Let p be an endpoint of e and thus a common vertex of τ and τ'. It is shown in the proof of Lemma 7.6 that each of the two angle $\measuredangle(n_\tau^+, n_p^+)$ and $\measuredangle(n_{\tau'}^+, n_p^+)$ are less than 13° when $\varepsilon \leq 0.03$. Thus the angle between n_τ^+ and $n_{\tau'}^+$ is at most 26°. Both n_τ^+ and $n_{\tau'}^+$ point toward the interior of W and therefore consistently orient τ and τ'. Thus ϕ, the dihedral angle between τ and τ', can only be between $180 - 26 = 154°$ and $180 + 26 = 206°$ (note that an inconsistent orientation of τ and τ' puts ϕ in the range 0–26° or 334–360°).

In the rest of this proof, we analyze this setting restricted only to the bisector plane Π of e. Note that the Voronoi facet dual to e is a planar polygon P contained in Π and so is the driver $d_Q(e)$, which is the closest point to e (or equivalently the mid-point m of e) on P. Let t and t' be intersections of τ and τ' with Π respectively. Thus t and t' are line segments in Π, each incident at one end to m. Voronoi edges s and s' respectively dual to τ and τ' are edges of the polygon P and are contained in Π. Also, let ℓ, ℓ', λ, and λ' be the supporting lines of t, t', s, and s' in Π respectively.

To simplify the argument, in drawing this arrangement on Π, we place m at the origin and draw the lines ℓ and ℓ' with a small (less than 26°) angle from the horizontal axis of the plane. The lines λ and λ' make similar angles with the vertical axis. It is thus meaningful to talk about the *top* or *bottom* planar wedge made by ℓ and ℓ' or by λ and λ'. In the sequel, the top wedge made by ℓ and ℓ' is denoted by $\ell \uparrow \ell'$ and the bottom wedge by $\ell \downarrow \ell'$. The wedge W corresponds to the planar wedge made by ℓ and ℓ' as determined by t and t'.

Observation 1 A first observation is that regardless of the position of t and t' along ℓ and ℓ' with respect to m, both s and s', the Voronoi edges dual to τ and τ' respectively, intersect $\ell \uparrow \ell'$. This is a consequence of Lemma 7.6 which gives an upper-bound the angle between $x - m$ and n_τ^+, where x is an intersection point of s and outer boundary of Σ_δ.

$$\measuredangle(x - m, n_\tau^+) \leq \frac{r}{1 - 3r} + \beta\left(\frac{\ell}{1 - r}\right) + \arcsin\left(\frac{r}{\ell}\right) < 57°,$$

for our choices of ε and δ. Thus then angle between $x - m$ and ℓ is at least $90 - 57 = 33°$. Recalling that n_τ^+ is the vector we get by orienting λ upward, shows that x is in $\ell \uparrow \ell'$. Similarly, any intersection point x' between s' and the outer boundary of Σ_δ lies in the same wedge.

The second key observation is based on the definitions of Delaunay and Voronoi complexes.

Observation 2 The order of λ and P along the direction ℓ agrees with that of t and m. The same holds with λ replaced with λ' and t with t'.

We first look at the case where $\phi \geq \pi$. Refer to Figure 9. In this figure the obtuse angle made by t and t' corresponds to W. We want to show that the driver $d_Q(m)$ of m, i.e. the closest point to m on P is *not* in $\ell \downarrow \ell'$. This would imply that $v_Q(m)$ points toward the interior or W, which is what we wish to prove. By Observation 2, Since t is to the left of m, P must be on the right of λ. Similarly, P must be on the left of λ'. Thus P must be in $\lambda \downarrow \lambda'$ (grayed in the figure). Since by Observation 1, P intersects $\ell \uparrow \ell'$, this wedge and $\lambda \uparrow \lambda'$ must intersect. Thus, the left case of Figure 9 cannot happen. Thus suppose $\lambda \uparrow \lambda'$ intersects $\ell \uparrow \ell'$. We consider two cases depending on whether m is in $\lambda \uparrow \lambda'$ or not. If not (Figure 9 middle), suppose to the contrary that the driver $d = d_Q(m)$ is in $\ell \downarrow \ell'$. Consider the line among λ and λ' that separates m and d, say λ' as in the figure, and consider the angle between $m - d$ and $x' - d$, where x' is a point of intersection of s' and the outer boundary of Σ_δ. It is easy to observe that this angle is acute. However, P being a convex polygon, the segment dx' is contained in P and this segment making an acute angle with mx is in contradiction with d being the closest point of P to m.

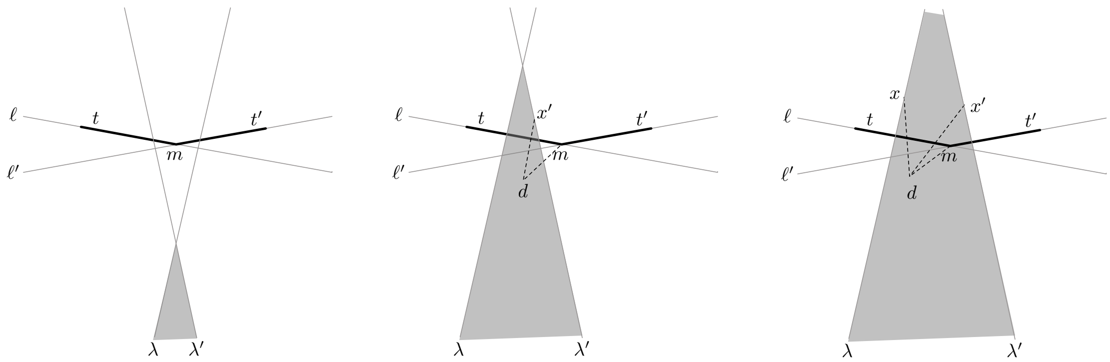

Figure 9: Proof of Lemma 7.7: $\phi \geq \pi$

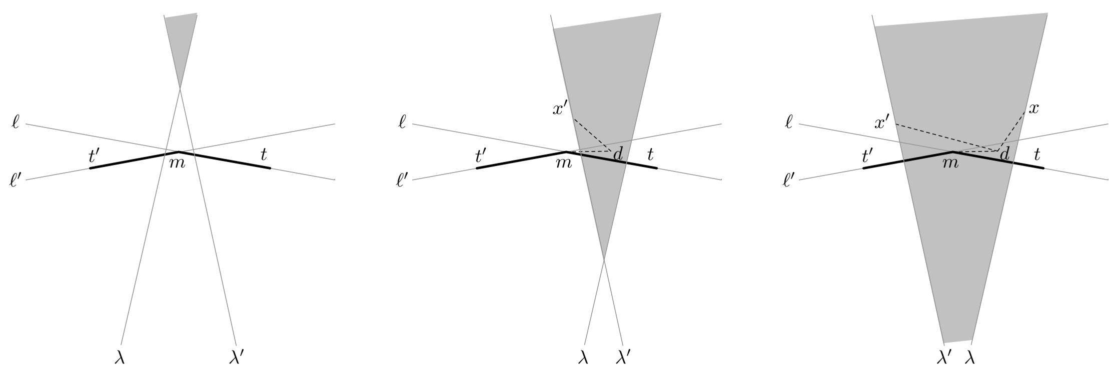

Figure 10: Proof of Lemma 7.7: $\phi \geq \pi$

In the case where m is contained in the wedge $\lambda \downarrow \lambda'$, suppose again to the contrary that the driver d of m is the wedge $\ell \downarrow \ell'$. It is easy to observe that in this case, at least one of the angle $\measuredangle(x-d, m-d)$ and $\measuredangle(x'-d, m-d)$ is acute. This results a contradiction in the same manner of the previous case.

The case $\phi < \pi$ is similarly handled. Refer to Firgure 10. Here the wedge W determined by t and t' is exactly $\ell \downarrow \ell'$ and we want to prove that the driver $d = d_Q(m)$ is in the wedge $\ell \uparrow \ell'$ which entails that the flow in at m enters W. Notice that in this case, Observation 2 implies that $P \subset \lambda \uparrow \lambda'$. Thus if $\lambda \uparrow \lambda'$ is contained in $\ell \uparrow \ell'$ (left figure), there is nothing to prove. If $\lambda \uparrow \lambda'$ is not contained in $\ell \uparrow \ell'$ but it doesn't include m (middle figure), then let λ' be the line that separates P and m and let x' be an intersection point of s' and the outer boundary of Σ_δ. If $d \in W$, then $\measuredangle(m-d, x'-d)$ is acute leading to a contradiction as in the previous case. If m is contained in $\lambda \uparrow \lambda'$ and $d \in W$, then as was done above, it is easy to observe that at least one of the angle $\measuredangle(m-d, x'-d)$ or $\measuredangle(m-d, x-d)$ resulting a contradiction with d being the driver of m.

■

Using the previous lemma, the geometric guarantee of the Wrap algorithm is given by the following theorem.

Theorem 7.8 *For $\varepsilon \leq 0.03$, the output I of the* Wrap *algorithm is contained in $S \cup \Sigma_{9\varepsilon^2}$ and includes $S_{9\varepsilon^2} = S \setminus \Sigma_{9\varepsilon^2}$.*

Proof. Lemma 7.7 shows that no simplices of the subcomplex $\mathcal{D}$ of $D = \mathrm{Del}\, P_1$ is preceded by a surface or an interior medial axis critical simplices. This means that all simplices in $\mathcal{D} \cap \mathrm{Del}\, P$ are picked by the WRAP algorithm and put in O. Since by Lemma 7.3, all simplices in $D \setminus \mathcal{D}$ are contained in $S \cup \Sigma_{9\varepsilon^2}$, the same is true for $I = \mathrm{Del}\, P \setminus O$, i.e. $|I| \subset S \cup \Sigma_{9\varepsilon^2}$.

To show that $I \supset S \setminus \Sigma_{9\varepsilon^2}$, assume to the contrary that a simplex $\tau \in O$ reaches $S \setminus \Sigma_{9\varepsilon^2}$. Since by Corollary 6.6 $\Sigma_{9\varepsilon^2}$ is flow-tight, any point $x \in \tau \cap S_{9\varepsilon^2}$ can only be on the unstable manifold of inner medial axis critical points and therefore τ is preceded by some inner medial axis critical simplex. This contradicts the choice of τ. ■

8 Topological Correctness

Following Lieutier [16] 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 [14].

Proposition 8.1 *Let X and $Y \subseteq X$ be arbitrary sets and let $H : [0,1] \times X \to X$ 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) $\forall y \in Y, \forall t \in [0,1]$, $H(t,y) \in Y$. *Then X and Y have the same homotopy type.*

Intuitively, we may interpret the first argument of the map H as time. Using a simple reparameterization in the first argument, we can replace the interval $[0,1]$ with any interval $[0,T]$ where $T > 0$ is a real number. The finiteness of the considered time interval is crucial to the definition. 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 by time T, they all arrive in Y. Notice that it is important that the points in Y stay in Y at all times.

Let Q be a set of (possibly weighted) points in $\mathbb{R}^n$. If X and Y are subsets of $\mathbb{R}^n$ with $Y \subset X$, in order to establish a homotopy equivalence between X and Y by applying Proposition 8.1 using ϕ_Q as H, one must show that

1. the flow orbit of every point in X stays in X, i.e. $\phi_Q(X) = X$ (and thus the map ϕ_Q can be restricted to X alone),
2. the flow orbit of every point in Y stays in Y, i.e. $\phi_Q(Y) = Y$, and
3. within a finite amount of time, the flow orbit of every point in X ends in Y.

Notice that the first condition of Proposition 8.1 is automatically satisfied for any flow map ϕ_Q because for every $x \in \mathbb{R}^n$, $\phi_Q(0,x) = x$. If X is bounded, the finiteness of flow time into Y can be guaranteed using Corollary 5.5 provided there is a lower bound c for $\|v_Q(x)\|$ for every $x \in X \setminus Y$: let Δ be an upper bound on the diameter of X and assume that for some constant $c > 0$, $\|v_Q(x)\| \geq c$ for every $x \in X \setminus Y$. Let $y = \phi_Q(t,x)$ be in $X \setminus Y$. Since Y is flow-tight for ϕ_Q, $y \notin Y$ implies that $\phi_Q(\tau,x) \notin Y$ for all $0 \leq \tau \leq t$. Consequently $\|v_Q(\phi_Q(\tau,x))\| \geq c$. Then by Corollary 5.5 we have for $y = \phi_Q(t,x)$

$$\begin{aligned} h_Q(y) &= h_Q(\phi_Q(t,x)) \\ &= h_Q(x) + \int_0^t \|v_Q(\phi_Q(\tau,x))\|^2 d\tau \\ &\geq h_Q(x) + \int_0^t c^2 d\tau \\ &= h_Q(x) + tc^2. \end{aligned} \tag{15}$$

But since X is flow-tight for ϕ_Q, $y \in X$ and therefore

$$\begin{aligned} h_Q(y) &= \min_{q\in Q} \left(\|y-q\|^2 - w_q\right) \\ &\le \max_{x\in X, q\in\tilde{Q}} \|x-q\|^2 - \min_{q\in Q} w_q \\ &\le \left(\Delta + \operatorname{dist}(X,\tilde{Q}) + \operatorname{diam}\tilde{Q}\right)^2 - \min_{q\in Q} w_q. \end{aligned}$$

The latter quantity is bounded and this put a finite upper bound on the value of t in the inequality (15). Thus we have proved the following theorem.

Theorem 8.2 *Let Q be a finite set of (possibly weighted) points in $\mathbb{R}^n$. If for sets $Y \subset X \subset \mathbb{R}^n$, X and Y are both flow-tight for ϕ_Q, i.e. $\phi_Q(X) = X$ and $\phi_Q(Y) = Y$, and there is a constant $c > 0$ for which $\|v_Q(x)\| \geq c$ for all $x \in X \setminus Y$, then X and Y are homotopy equivalent.*

The above theorem is the key to all homotopy equivalence proofs in the rest of this monograph. To invoke the theorem we need two flow-tight sets X and Y which we sometimes call the *source set* and the *sink set*, respectively. Stable and unstable manifolds, their unions, and their intersections are flow-tight by definition. However, these won't be the only examples of flow-tight sets we will consider.

Lemma 8.3 *The two sets $S' = S \cup \Sigma_\delta$ and $\operatorname{cl} S = S \cup \Sigma$ are homotopy equivalent, for any $0 < \delta < 1$. In fact, the latter is a strong deformation retract of the former.*

Proof of Lemma 8.3. Consider the retraction map $r : S' \to \operatorname{cl} S$ given by

$$r(x) = \begin{cases} \hat{x} & x \in S' \setminus \operatorname{cl} S \\ x & x \in \operatorname{cl} S \end{cases}$$

The map r is continuous on Σ_δ because $\Sigma_\delta \cap M(\Sigma) =$ and $M(\Sigma)$ consists of the only points in space where the map $x \mapsto \hat{x}$ is not continuous (in fact undefined). Since the map r is identity on $\operatorname{cl} S$, r is continuous on all of its domain.

If we now define the map $R : [0,1] \times S' \to S'$ as

$$R(t,x) = \begin{cases} (1-t)x + t\hat{x} & x \in S' \setminus \operatorname{cl} S \\ x & x \in \operatorname{cl} S, \end{cases}$$

the map R is a straight-line homotopy from the identity on S' to the retraction r. ∎

Lemma 8.4 *Let E be the union of unstable manifolds of all surface and inner medial axis critical points under ϕ_{Q_1}. Then $S \cup \Sigma_\delta$ is homotopy equivalent to E for $9\varepsilon^2 < \delta < 3/10 - 2\varepsilon$.*

Proof. By Corollary 6.6, $\Sigma_\delta \cup S$ is flow-tight for ϕ_{Q_1}. On the other hand E is also flow-tight for ϕ_{Q_1} since by definition $\phi_{Q_1}(E) = E$. Therefore Theorem 8.2 implies the desired homotopy equivalence if we only show that $\|v_{Q_1}(x)\|$ is bounded from below for every $x \in (S \cup \Sigma_\delta) \setminus E$. But $\|v_{Q_1}(x)\| = 2\|x - d_{Q_1}(x)\|$. When $D_{Q_1}(x) \cap V_{Q_1}(x) = \emptyset$, the distance between x and $d_{Q_1}(x)$ is at least the distance between $V_{Q_1}(x)$ and $D_{Q_1}(x)$. Since both $\operatorname{Vor} Q_1$ and $\operatorname{Del} Q_1$ are finite complexes, there is a lower bound on the distance between pairs of dual faces from the two complexes that do not intersect.

On the other hand, if $V_{Q_1}(x) \cap D_{Q_1}(x)$ do intersect, then their intersection is a critical point c coinciding with $d_{Q_1}(x)$. Thus x is by definition on the unstable manifold of c. If c is a surface or inner medial axis critical point, then $\operatorname{Um}(c) \subset E$ and thus $x \in E$ contradicting the choice of c. Thus c can only be an outer medial axis critical point. But then, by Theorem 2.1, c is not contained in $\Sigma_{1-2\varepsilon}$. The Hausdorff distance between $\Sigma_{1-2\varepsilon}$ and Σ_δ being strictly positive puts a lower bound on the distance between x and c. ∎

Lemma 8.5 *Let E be as in Lemma 8.4. If $\varepsilon \le 0.03$, then the output I of the algorithm* Wrap *is homotopy equivalent to E.*

Proof. We first observe that I is flow-tight for ϕ_{Q_1}. By definition, the set O in the Wrap algorithm consists of those simplices in $\mathrm{Del}\, P$ that are only preceded by exterior medial axis critical simplices (including ω). If a flow line leaves I it enters O. But by definition, simplices in I are reachable from critical simplices other than the outer medial axis critical ones. Therefore if a flow line enters O from I, it provides a path from a critical simplex of surface or inner medial axis to some simplex in O contradicting the definition of O.

On the other hand, E is also flow-tight for ϕ_Q as was shown in the proof of Lemma 8.4. We observe here that $E \subseteq I$. This is because E is by definition the locus of all points reachable by some flow path starting at a critical simplex of inner medial axis or surface. Trivially, these critical simplices are themselves in I. For every other point $x \in E$, there is a flow path $\phi_Q(y)$ starting at a point y in the relative interior of some critical simplex τ in the mentioned group reaching x. Let $\tau = \tau_0, \ldots, \tau_k$ be the set of simplices intersecting $\phi_Q(y)$ until it reaches x with τ_k being the simplex containing x in its relative interior. The existence of this path implies that $\tau_0 \prec \cdots \prec \tau_k$. Therefore τ_k cannot be placed in O by the Wrap algorithm and consequently $\tau_k \in I$. So, we have shown that $E \subseteq I$.

Thus in order to complete the proof we only need to show that at every $x \in I \backslash E$, $\|v_Q(x)\| > c$ for some $c > 0$ and then the homotopy equivalence between E and I follows from Theorem 8.2. Similar to the proof of Lemma 8.4, we only need to show that the points in $I \setminus E$ are at a positive distance from every outer medial axis critical point. This follows by a similar argument to the one in the proof of Lemma 8.4 using the fact that $I \subset S \cup \Sigma_\delta$ for $\delta = 9\varepsilon^2$ as is shown by Theorem 7.8. ■

The topological guarantee of the wrap algorithm follows from combining Lemmas 8.3, 8.4, 8.5.

Theorem 8.6 *Let P be an ε-sample of a surface Σ embedded in $\mathbb{R}^3$ with $\varepsilon < 0.03$. Then the output I of the* Wrap *algorithm is homotopy equivalent to the bounded shape S enclosed by Σ.*

9 Conclusions and Further Work

We have described a variant of the Wrap reconstruction algorithm that in 3D produces an output with geometric and topolgical guarantees. It remains unclear, even in 3D, whether the boundary of the shape output by the algorithm is a manifold. It appears that resolving this requires a better understanding of how the Delaunay complex intersects the unstable manifolds of the critical points on the surface. In this direction, a better general understanding of the unstable manifolds for weighted point sets may be needed. As for higher dimensions, we currently have no guarantees under the ε-sampling assumption. Our argument depends on showing that Wrap manages to get very close to Σ by removing enough simplices from $\mathrm{conv}\, P$. Our proof for this fact in 3D relies on a result that says a triangle with “small” circumradius lies “flat” to the surface [3]. In particular, this result is not known to generalize to higher dimensions.

References

[1] N. Amenta and M. W. Bern. Surface reconstruction by Voronoi filtering. *Discrete & Computational Geometry*, 22:481–504, 1999.

[2] N. Amenta, M. W. Bern, and D. Eppstein. The crust and the β-skeleton: Combinatorial curve reconstruction. *Graphical Models and Image Processing*, 60(2):125–135, 1998.

[3] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *International Journal of Computational Geometry and Applications*, 12(1-2):125–141, 2002.

[4] N. Amenta, S. Choi, and R. K. Kolluri. The power crust, unions of balls, and the medial axis transform. *Computational Geometry*, 19(2-3):127–153, 2001.

[5] J.-D. Boissonnat and F. Cazals. Smooth surface reconstruction via natural neighbour interpolation of distance functions. *Computational Geometry*, 22(1-3):185–203, 2002.

[6] T. K. Dey. Curve and surface reconstruction. In J. E. Goodman and J. O. Rourke, editors, *Handbook of Discrete and Computational Geometry*. CRC Press, second edition, 2004.

[7] T. K. Dey. *Curve and Surface Reconstruction : Algorithms with Mathematical Analysis*. Cambridge University Press, 2006.

[8] T. K. Dey, J. Giesen, E. A. Ramos, and B. 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.

[9] T. K. Dey and P. Kumar. A simple provable algorithm for curve reconstruction. In *Symposium on Discrete Algorithms (SODA)*, pages 893–894, 1999.

[10] H. Edelsbrunner. Surface reconstruction by wrapping finite point sets in space. *Discrete & Computational Geometry*, 32:231–244, 2004.

[11] J. Giesen and M. John. The flow complex: A data structure for geometric modeling. In *Symposium on Discrete Algorithms (SODA)*, pages 285–294, 2003.

[12] J. Giesen, E. A. Ramos, and B. Sadri. Medial axis approximation and unstable flow complex. In *Symposium on Computational Geometry*, pages 327 – 336, 2006.

[13] K. Grove. Critical point theory for distance functions. *Symposia in Pure Mathematics*, 54(3):357–385, 1993.

[14] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2001.

[15] M. John. *Flow Complexs: Structure, Algorithms, and Applications*. PhD thesis, Swiss Federal Institute of Technology, ETH Zürich, 2003.

[16] A. 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.