

Sampling and Meshing a Surface with Guaranteed Topology and Geometry*

Siu-Wing Cheng[†] Tamal K. Dey[‡] Edgar A. Ramos[§] Tathagata Ray[‡]

March 2, 2007

Abstract

This paper presents an algorithm for sampling and triangulating a generic C^2 -smooth surface $\Sigma \subset \mathbb{R}^3$ that is input with an implicit equation. The output triangulation is guaranteed to be homeomorphic to Σ . We also prove that the triangulation has well-shaped triangles, large dihedral angles, and a small size. The only assumption we make is that the input surface representation is amenable to certain types of computations, namely computations of the intersection points of a line and Σ , computations of the *critical points* in a given direction, and computations of certain *silhouette points*.

1 Introduction

The need for triangulating a surface is ubiquitous in science and engineering. A set of points needs to be sampled from the input surface and then connected to generate such a triangulation. The underlying space of the resulting triangulation should have the same topology as that of the input surface. Nice geometric properties such as bounded aspect ratio and large dihedral angles are also desirable. The input surface can be specified in various ways and each leads to a different problem of surface triangulation.

When the surface is given by a set of point samples, the problem is known as *surface reconstruction* for which algorithms with topological and geometrical guarantees have been proposed [1, 2, 5, 18]. When the surface is polyhedral, i.e., made out of planar patches, the Delaunay refinement techniques solve the problem elegantly [11, 12, 13, 15, 31].

The case in which the input surface is smooth and specified by an implicit equation occurs in a variety of applications in geometric modeling, computer graphics, and finite element methods [34, 37, 32]. Obtaining a surface triangulation that has the correct topology and nice geometric properties (e.g. bounded aspect ratio and large dihedral angles) is an important issue in these applications. In this paper we present an algorithm to triangulate a generic C^2 -smooth

*Research partly supported by NSF grants DMS-0310642, CCR-0430735, and by Research Grant Council, Hong Kong, China (HKUST 6181/04E). A preliminary version of the paper appeared in the Proc. of 20th Ann. Sympos. Comput. Geom.

[†]Department of Computer Science and Engineering, HKUST, Hong Kong. Email: scheng@cse.ust.hk

[‡]Department of Computer Science and Engineering, Ohio State University, Ohio, USA. Email: tamaldey@cse.ohio-state.edu and rayt@cse.ohio-state.edu

[§]Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana-Champaign, USA. Email: eramosn@cs.uiuc.edu

implicit surface without boundary. All vertices are sampled from the input surface and the following guarantees are offered for any fixed $\lambda \leq 0.06$:

- (i) The underlying space of the output triangulation is homeomorphic to the input surface.
- (ii) All angles of the triangles are at least $\arcsin\left(\frac{1}{2+2\lambda}\right)$.
- (iii) The dihedral angle between adjacent triangles is at least $\pi - O(\lambda)$.
- (iv) The number of vertices is no more than $O\left(\frac{\varepsilon^2}{\lambda^2}\right)$ times the size of an ε -sample¹ for any $\varepsilon < \frac{1}{5}$.

Although the output triangulation does not have bounded Hausdorff distance from the input surface as enjoyed by an ε -sample [1, 2], property (iv) shows that our algorithm does not sample an excessive number of vertices.

Because of the importance of the problem, a number of algorithms have been proposed for meshing implicit surfaces across various application areas. These algorithms may give satisfactory experimental results, but they do not have guarantees on the validity of the topology and/or the quality of the output triangulation. We briefly survey a selected subset. The problem of triangulating implicit surfaces has been investigated by Bajaj et al. [3], Bloomenthal [4], Tristano, Owen, and Canann [35], Lau and Lo [27], and Cuillière [16]. The marching cube algorithm of Lorensen and Cline [28] can be used to triangulate an implicit surface. The algorithm determines the edges of a cubic grid intersecting the surface and then generate a tessellation by connecting these intersection points. Although the algorithm is very simple, there is no guarantee that the output has the same topology as that of the surface. Stander and Hart [34] proposed to vary the value of the implicit function from $-\infty$ to ∞ and dynamically maintain a triangulation of the changing isosurface. It is necessary to track all critical points of the implicit function. Maintaining triangulations of isosurfaces is a huge overhead given that only one isosurface, namely the input, needs to be triangulated. Witkin and Heckbert [37] proposed to spread particles governed by differential equations on the implicit surface. At equilibrium, the particles can be connected to form the surface triangulation, but it is unclear how to ensure that the surface topology is captured. The above algorithms do not offer any guarantee on the triangle shape though some of them include heuristics and illustrate their effectiveness experimentally.

In computational geometry, Chew [14] described an algorithm based on the “furthest-point” strategy: among the intersections between a Voronoi edge and the input surface, select and insert the furthest one from the sites defining the Voronoi edge. In effect, this algorithm attempts to compute the *restricted Delaunay triangulation* of the surface. Edelsbrunner and Shah [22] showed that a *topological ball property* is sufficient for the restricted Delaunay triangulation to be homeomorphic to the input surface. The algorithm of Chew does not guarantee this property or any other that ensures topological correctness.

Following the “furthest-point” strategy, Cheng, Dey, Edelsbrunner, and Sullivan [10] proposed an algorithm for triangulating the *skin surface* [20] that provides both topological and geometric guarantees. This algorithm exploits the fact that the *local feature size* is easily computable for skin surfaces. The local feature size of a point x on the surface is the distance from x to the medial axis.

¹defined in Section 2.1

Boissonnat and Oudot [8] carried forward the “furthest-point” strategy for general curved surfaces. They showed how to grow from an initial seed triangle on each surface component to a full triangulation with topological and geometric guarantees. The algorithm assumes that one can compute the local feature size of any point on the surface. Computing the medial axis is hard and hence computing the local feature size exactly is difficult, if not impossible, for surfaces in general. Of course, one can approximate the medial axis with existing algorithms [1, 5, 17]. However, these algorithms require a dense sampling with respect to the local feature size in the first place. An alternative suggested by Boissonnat and Oudot is to compute the minimum local feature size and then to run their algorithm to obtain a dense sample from which the medial axis can be approximated. In a second pass a new mesh can be computed with appropriate density using the approximated local feature size.

A related work by Boissonnat, Cohen-Steiner, and Vegter [7] considered triangulating the isosurfaces of a function $E : \mathbb{R}^3 \rightarrow \mathbb{R}$. Their method evaluates E at grid points and triangulates a box in \mathbb{R}^3 recursively to provide a piecewise linear interpolant \widehat{E} of E . The isosurface $E = 0$ is approximated with the isosurface $\widehat{E} = 0$. The authors provided conditions on sampling to guarantee that the computed surface $\widehat{E} = 0$ is isotopic to the surface $E = 0$. Their method samples the function E rather than the surface $E = 0$. Moreover, it computes the critical points of E as well as their indices. There are two other algorithms for producing isotopic triangulations, one by Mourrain and T  court [29] and another by Plantinga and Vegter [30]. More details about these isotopic triangulation algorithms can be found in the survey [6].

In this paper, we eliminate the need for local feature size computation. We show that it suffices to identify the critical points and a silhouette of the surface or a cross-section of the surface in a given direction. These computations are less demanding. (The critical points of the surface in a given direction should not be confused with the critical points of the implicit function as computed by the algorithms of Stander and Hart [34], and Boissonnat, Cohen-Steiner, and Vegter [7].) To this end, we depart from the strategy of Boissonnat and Oudot [8, 9] in a fundamental way. Topological ball property violations drive the refinement in our algorithm whereas they are used only for analysis in [8, 9].

Our algorithm incrementally grows a set of point samples and maintains the restricted Delaunay triangulation of the samples. In the topology recovery phase, we use a simple “topological-disk” test and certain critical and silhouette point computations to guide the sampling of points from the surface. Our approach extends the “furthest-point” strategy that selects samples only from the intersections between Voronoi edges and the surface. This extension allows us to sample points adaptively and prove that the output triangulation has the same topology as that of the surface. In the geometry recovery phase, we enforce the bounded aspect ratio and smoothness of the surface triangulation.

2 Preliminaries

2.1 Input surface and assumptions

The input is a compact surface $\Sigma \subset \mathbb{R}^3$ without boundary. We assume that Σ is specified as the zero-level set of a function $E : \mathbb{R}^3 \rightarrow \mathbb{R}$ such that the second partial derivatives of E are continuous (i.e., Σ is C^2 -smooth), and $E(x)$ and $\nabla E(x)$ do not vanish simultaneously at any $x \in \mathbb{R}^3$.

A maximal ball whose interior is disjoint from Σ is called a *medial ball*. The *medial axis* of Σ

is the set of the centers of medial balls. We borrow some definitions from Amenta and Bern [1] who proposed them in the context of surface reconstruction. The *local feature size* $f(x)$ at a point $x \in \Sigma$ is the distance from x to the medial axis. Since Σ is assumed to be C^2 -smooth and compact, $\min_{x \in \Sigma} f(x)$ is non-zero. The function f is 1-Lipschitz, that is, $f(x) \leq f(y) + \|x - y\|$ for any two points x and y on Σ . A point set $S \subset \Sigma$ is an ε -sample if for any $x \in \Sigma$, there is a point $p \in S$ such that $\|p - x\| \leq \varepsilon f(x)$.

Throughout this paper, we use x_1 , x_2 , and x_3 to denote the three orthogonal directions forming the coordinate frame. Given any point $y \in \mathbb{R}^3$, we use (y_1, y_2, y_3) to denote the coordinates of y . For any vector d in \mathbb{R}^3 , we use (d_1, d_2, d_3) to denote its components in the x_1 -, x_2 -, and x_3 -directions. Given two vectors d and d' , let $\langle d, d' \rangle$ and $d \times d'$ denote their inner and cross products, respectively. The gradient $\nabla E(x)$ is the vector $\left(\frac{\partial E(x)}{\partial x_1}, \frac{\partial E(x)}{\partial x_2}, \frac{\partial E(x)}{\partial x_3} \right)$. If the point x lies on Σ , $\nabla E(x)$ is parallel to the unit outward surface normal at x . The *critical points of Σ in a direction d* are the points $x \in \Sigma$ such that $\nabla E(x)$ is parallel to d .

We assume that Σ has finitely many critical points in any direction. We also assume that the Hessian at any point $x \in \Sigma$ is non-singular, i.e., for any two orthogonal tangent directions u_1 and u_2 at $x \in \Sigma$, the matrix $\left(\frac{\partial^2 E(x)}{\partial u_i \partial u_j} \right)$ is non-singular at x .

A unit vector d induces a height function h_d on Σ : $h_d(x) = \langle \nabla E(x), d \rangle$ for any $x \in \Sigma$. The set $h_d^{-1}(0)$ consists of the points $x \in \Sigma$ such that $\nabla E(x)$ is orthogonal to d , i.e., d is a tangent direction at x . Let v be a tangent direction at x orthogonal to d . Orient space so that d aligns with the x_1 -axis and v aligns with the x_2 -axis. We have $h_d(x) = \frac{\partial E(x)}{\partial d}$ and so $\left(\frac{\partial h_d(x)}{\partial d}, \frac{\partial h_d(x)}{\partial v} \right) = \left(\frac{\partial^2 E(x)}{\partial d^2}, \frac{\partial^2 E(x)}{\partial d \partial v} \right) \neq 0$ because the Hessian is assumed to be non-singular at $x \in \Sigma$. In other words, the points in $h_d^{-1}(0)$ are not critical for h_d and so it follows from the inverse function theorem in differential topology [25] that $h_d^{-1}(0)$ is a collection of smooth closed curves. We call these curves the *silhouette of Σ with respect to d* and denote it by J_d .

Take any direction d' orthogonal to d . A point $x \in J_d$ is critical in direction d' if the tangent to J_d at x is orthogonal to d' . We assume that J_d has finitely many critical points in direction d' . Given a plane Π , the critical points of the intersection curve(s) in $\Sigma \cap \Pi$ in any direction parallel to Π can be similarly defined. We also assume that the intersection curve(s) between Σ and any plane Π have finitely many critical points in any direction parallel to Π .

2.2 Generic intersection and topological ball property

Let S be a finite point set in \mathbb{R}^3 . The *Voronoi cell* of a point $p \in S$ is defined as $V_p = \{x \in \mathbb{R}^3 : \forall q \in P, \|p - x\| \leq \|q - x\|\}$. A Voronoi cell is a convex polyhedron. For $2 \leq j \leq 4$, the closed faces shared by j Voronoi cells are called $(4 - j)$ -dimensional *Voronoi faces*. The 0-, 1-, 2-dimensional Voronoi faces are called *Voronoi vertices, edges, and facets*, respectively. The *Voronoi diagram* $\text{Vor } S$ is the collection of all Voronoi faces.

Assuming general position, the convex hull of $j \leq 4$ points in S defines a $(j - 1)$ -dimensional *Delaunay simplex* σ if the vertices of σ define a $(4 - j)$ -dimensional Voronoi face in $\text{Vor } S$. We use V_σ to denote this Voronoi face. We call σ and V_σ the *dual* of each other. The 1-, 2-, and 3-dimensional Delaunay simplices are called *Delaunay edges, triangles, and tetrahedra*, respectively. The Delaunay simplices define a decomposition of the convex hull of S called the *Delaunay triangulation* of S . We denote it by $\text{Del } S$.

The set of Delaunay simplices whose dual Voronoi faces intersect Σ form the *restricted*

Delaunay triangulation $\text{Del} S|_\Sigma$ of S with respect to Σ . Formally, $\text{Del} S|_\Sigma = \{\sigma \in \text{Del} S : V_\sigma \cap \Sigma \neq \emptyset\}$. For any simplex $\sigma \in \text{Del} S|_\Sigma$, the intersection $V_\sigma \cap \Sigma$ is called a *restricted Voronoi face*. The *restricted Voronoi diagram* $\text{Vor} S|_\Sigma$ is the collection of all restricted Voronoi faces.

We say that a Voronoi face V_σ satisfies *generic intersection property* (GIP) if either $V_\sigma \cap \Sigma = \emptyset$ or V_σ intersects Σ *transversally*, that is, the *affine space* of V_σ is not tangent to Σ at the points in $V_\sigma \cap \Sigma$. In particular, this implies that a Voronoi vertex should not lie on Σ . The Voronoi diagram $\text{Vor} S$ satisfies GIP for Σ if all of its faces satisfy GIP.

We say that a Voronoi face V_σ satisfies the *topological ball property* (TBP) if either $V_\sigma \cap \Sigma = \emptyset$ or the restricted Voronoi face $V_\sigma \cap \Sigma$ is a closed topological ball of dimension $\dim(V_\sigma) - 1$, where $\dim(V_\sigma)$ is the dimension of V_σ . The Voronoi diagram $\text{Vor} S$ satisfies TBP for Σ if all its Voronoi faces satisfy TBP. Our meshing algorithm is based on the following result of Edelsbrunner and Shah that relates the topology of $\text{Del} S|_\Sigma$ to Σ .

Theorem 2.1 [22] *The underlying space of $\text{Del} S|_\Sigma$ is homeomorphic to Σ if $\text{Vor} S$ satisfies TBP and GIP for Σ .*

Notice that the above theorem is originally proved for non-degenerate point set, that is, no five points are co-spherical. However, one may drop this requirement by appealing to the SOS technique [21] that simulates generic conditions by perturbing the points symbolically.

2.3 Background results

We state a few geometric results in the literature that we use frequently. Let ℓ and ℓ' be two line segments, vectors, or lines. We use $\angle \ell, \ell'$ to denote the acute angle between the support lines of ℓ and ℓ' . For any point $x \in \Sigma$, we use n_x to denote the unit outward surface normal at x . For any triangle pqr , we use n_{pqr} to denote a unit normal to pqr . Define two functions $\alpha(\lambda)$ and $\beta(\lambda)$ where

$$\alpha(\lambda) = \frac{\lambda}{1 - 3\lambda} \quad \text{and} \quad \beta(\lambda) = \alpha(2\lambda) + \arcsin \lambda + \arcsin \left(\frac{2 \sin(2 \arcsin \lambda)}{\sqrt{3}} \right).$$

The key property of $\alpha(\lambda)$ and $\beta(\lambda)$ is that both are $O(\lambda)$ and they approach zero as λ does so.

Lemma 2.1 ([1]) *Let x and y be two points on Σ . If $\|x - y\| \leq \lambda f(x)$ for some $\lambda < 1/3$, $\angle n_x, n_y \leq \alpha(\lambda)$.*²

Lemma 2.2 ([2]) *Let pqr be a triangle with vertices on Σ . If the circumradius of pqr is less than $\lambda f(p)$ for $\lambda \leq \frac{1}{\sqrt{2}}$, $\angle n_{pqr}, n_p \leq \beta(\lambda)$.*

Lemma 2.3 ([10]) *Let x and y be two points in the intersection of a line ℓ and Σ . Then $\|x - y\| \geq 2f(x) \cos(\angle \ell, n_x)$.*

²The slightly stronger condition of $\|x - y\| \leq \lambda \min\{f(x), f(y)\}$ is stated in [1], but the proof uses $\|x - y\| \leq \lambda f(x)$ only.

3 Overview

Our algorithm’s goal is to obtain a sufficiently dense sample S on Σ for which the GIP and TBP hold. The approach to obtain S is incremental. First, we initialize S to contain the critical points of Σ in the x_3 -direction. Then, while either the GIP or TBP does not hold, we add one more point to S : a *witness* to the *violation*. The correctness of the algorithm is then trivial as long as it halts, which follows from a lower bound on how close a new inserted point can be to the existing ones.

As it turns out, it is not so easy to identify a TBP violation for a facet or cell. Therefore we take a conservative approach: a witness is always returned if there is a violation, but a witness may also be returned even in some cases where there is no violation. However, we guarantee that no harm is done in inserting the false witnesses since these false witnesses are also far away from existing points in S .

A GIP is violated when the affine space of a Voronoi edge or facet is tangent to Σ or when a Voronoi vertex lies on Σ . Notice that any tangential contact between Σ and a Voronoi edge or facet is isolated by our assumption that Σ is generic. The case of the affine space of a Voronoi edge or facet being tangent to Σ can be handled simply by returning the tangency point as witness, which can be determined by solving an appropriate system of equations. We will show that this witness is sufficiently far from the existing sample points. In contrast, a Voronoi vertex lying on Σ can happen at any sampling density. It is a degenerate intersection between Voronoi facets and cells with Σ . The algorithm pretends that Σ is perturbed locally to get around the degeneracy. Nevertheless, we have to deal with the Voronoi vertices on Σ directly in the proofs. The reason is that a local perturbation may change the local feature size a lot. Since we need to obtain distance lower bounds in terms of the local feature size with respect to Σ , we cannot assume the local perturbation in the analysis.

It is natural to test for TBP violations in increasing order of Voronoi face dimensions. Testing a TBP violation at an edge is comparatively easier than at a facet or a cell. We assume GIP for all edges and facets:

Edge e : Testing for a TBP violation at an edge e is simply a matter of counting the number of intersections between e and Σ , which is easily determined by computing all intersections between e and Σ . If there is a violation, the witness returned is the intersection point furthest from any sample point in S that generates e .

Facet F : It is assumed that TBP holds for edges. $F \cap \Sigma$ is a collection of closed or open curves (with endpoints in the boundary of F). TBP is violated if either $F \cap \Sigma$ includes more than one open curve or a closed curve. In the first case, there are more than two intersections between Σ and the boundary of F , and the witness returned is the furthest intersection point from any sample point in S that generates F . In the second case, checking whether $F \cap \Sigma$ contains a closed curve is not easy, so we settle for a necessary witness instead. We find the critical points of $F \cap \Sigma$ in some direction parallel to F . Then we compute the lines in the plane of F that are normal to $F \cap \Sigma$ at these critical points. If any such line intersects $F \cap \Sigma$ in two or more points, the furthest one is returned as the witness. Clearly, such a witness exists if $F \cap \Sigma$ contains a closed curve. While this witness may exist even if there is no closed curve, it will be shown that in either case the witness is sufficiently far from the existing sample points.

Cell V_p : It is assumed that TBP holds for edges and facets. $V_p \cap \Sigma$ is a collection of surface

patches with or without boundaries. There is a violation to TBP if $V_p \cap \Sigma$ is not a topological disk. In particular, a violation occurs if the boundary of $V_p \cap \Sigma$ consists of more than one closed curve. This is easily determined by checking whether the dual triangles incident to p form one or more topological disks. $V_p \cap \Sigma$ cannot contain a surface component without any boundary; otherwise, we would have placed the critical points of this component in the x_3 -direction as seeds inside V_p . The hardest case is that the boundary of $V_p \cap \Sigma$ is a single closed curve, but $V_p \cap \Sigma$ has positive genus. We handle this case using the silhouette J_{n_p} , i.e., the set of points $x \in \Sigma$ such that n_x is orthogonal to n_p . It is known that J_{n_p} is a smooth closed curve and the points in J_{n_p} are far from p . So the test can be done as follows: either J_{n_p} intersects some facet of V_p at a point w , or J_{n_p} contains an extreme point w in a direction orthogonal to n_p . The point w is returned as the witness in either case.

Topology recovery as described above generate a sample of the surface. This sample is further refined to ensure some geometric properties of the output triangulation such as aspect ratios of triangles and large dihedral angles at the edges. The rest of the paper is organized as follows. In Section 4, we present the analytic tools to cope with the violations of GIP and TBP. In Section 5, the details of the topology recovery part of our algorithm is given. Figure 1 below illustrates the the dependencies between the lemmas in Sections 2.3 and 4, the subroutines in Section 5, and TBP. In Section 6, we show how to enforce bounded aspect ratio and large dihedral angles by inserting new points that are far from existing sample points. After inserting some point(s) to repair the geometry, we have to rerun the topology recovery part because GIP or TBP may no longer hold. Thus, the entire algorithm alternates between repairing topology and geometry. The full analysis of the algorithm is presented in Section 6.

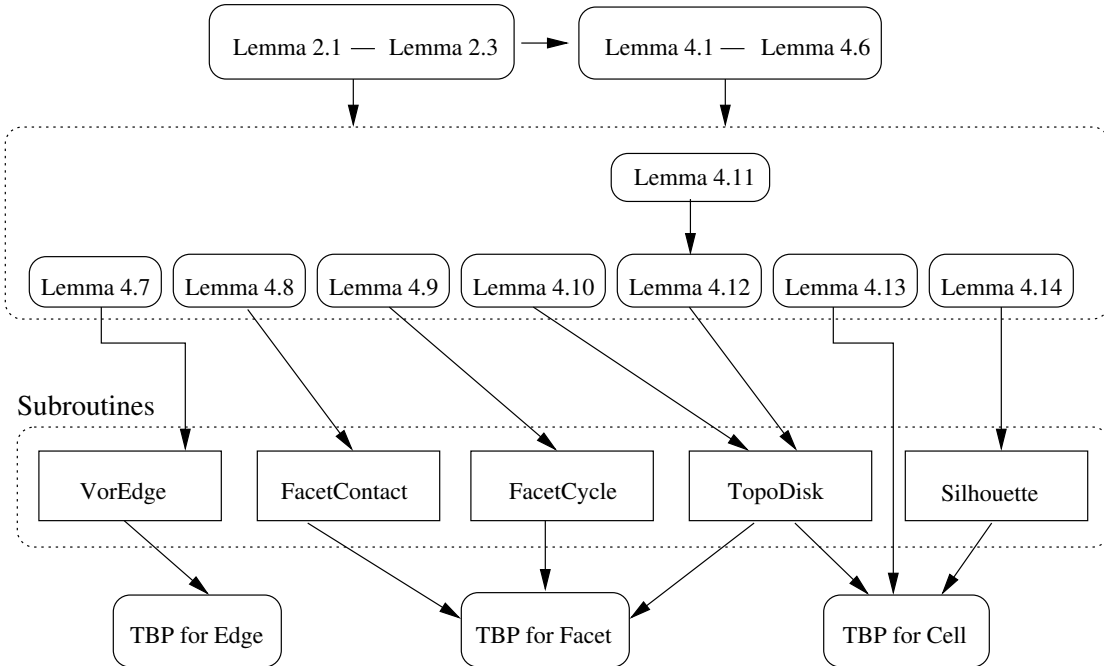


Figure 1: Dependencies between the subroutines, lemmas, and TBP.

4 Violation

Theorem 2.1 is our main tool to recover topology. The following subsections treat the violations of TBP for the cases of V_σ being a Voronoi edge, facet, and cell separately. In each case, we show how to identify a point $x \in V_\sigma \cap \Sigma$ such that $\|p - x\| \geq \lambda f(p)$ for any vertex p of σ where

$$\begin{aligned} \frac{\lambda}{1-\lambda} &< \cos(\alpha(\lambda) + 3\beta(\lambda)) \\ \alpha(\lambda) + \beta(\lambda) &< \pi/3 \\ \arccos \lambda &> \alpha(\lambda) + \beta(\lambda). \end{aligned}$$

The above inequalities hold for any $\lambda \leq \lambda_0 = 0.06$ which we assume throughout the rest of the paper.

The GIP may be violated if some Voronoi vertex lies on Σ . We defer this discussion to Section 5.1. The results in this section hold regardless of whether some Voronoi vertex lies on Σ .

4.1 Technical results

We first prove a few technical results that we use later.

Lemma 4.1 *Let p and x be two points on Σ . If $\|p - x\| \leq \lambda f(p)$, then $f(x) \geq (1 - \lambda)f(p)$.*

Proof. Because f is 1-Lipschitz, $f(x) \geq f(p) - \|p - x\| \geq (1 - \lambda)f(p)$. \square

Lemma 4.2 *Let p , x , and y be three points on Σ . If both $\|p - x\|$ and $\|p - y\|$ are at most $\lambda f(p)$, then $\|x - y\| \leq 2\lambda f(p) \leq 2\lambda f(x)/(1 - \lambda)$.*

Proof. By the triangle inequality, $\|x - y\| \leq \|p - x\| + \|p - y\| \leq 2\lambda f(p)$, which is at most $2\lambda f(x)/(1 - \lambda)$ by Lemma 4.1. \square

Lemma 4.3 *Let e be an edge of a Voronoi cell V_p . For any point $x \in e \cap \Sigma$, if $\|p - x\| \leq \lambda f(p)$, then $\angle e, n_x \leq \alpha(\lambda) + \beta(\lambda) < \pi/3$.*

Proof. By Lemma 2.1, $\angle n_p, n_x \leq \alpha(\lambda)$. The circumradius of the Delaunay triangle dual to e is at most $\|p - x\| \leq \lambda f(p)$. By Lemma 2.2, $\angle e, n_p \leq \beta(\lambda)$. Thus, $\angle e, n_x \leq \angle e, n_p + \angle n_p, n_x \leq \alpha(\lambda) + \beta(\lambda)$, which is less than $\pi/3$ as $\lambda \leq \lambda_0$. \square

Lemma 4.4 *Let F be a facet of a Voronoi cell V_p . Let Π be the plane of F . Suppose that $\Pi \cap \Sigma$ contains a point x such that $\|p - x\| \leq \lambda f(p)$.*

- (i) *The acute angle between Π and n_p is at most $\arcsin \lambda < \beta(\lambda)$.*
- (ii) *The acute angle between Π and n_x is at most $\alpha(\lambda) + \arcsin \lambda < \alpha(\lambda) + \beta(\lambda) < \pi/3$.*

Proof. Let pq be the Delaunay edge dual to F . We have $\|p - q\| \leq \|p - x\| + \|q - x\| = 2\|p - x\| \leq 2\lambda f(p)$. It then follows from Lemma 2.3 that $\angle pq, n_p \geq \arccos \lambda$ which implies (i). By Lemma 2.1, $\angle n_p, n_x \leq \alpha(\lambda)$. By the triangle inequality, $\angle pq, n_x \geq \angle pq, n_p - \angle n_p, n_x \geq \arccos \lambda - \alpha(\lambda)$. So the acute angle between Π and n_x is at most $\alpha(\lambda) + \arcsin(\lambda) < \alpha(\lambda) + \beta(\lambda)$, which is less than $\pi/3$ as $\lambda \leq \lambda_0$. \square

Lemma 4.5 *Let F be a facet of a Voronoi cell V_p such that $F \cap \Sigma$ contains no tangential contact point. Let x be a point in $F \cap \Sigma$. Let L be a line in the plane of F through x and normal to $F \cap \Sigma$ at x . If $\|p - x\| \leq \lambda f(p)$, then $\angle L, n_x < \alpha(\lambda) + \beta(\lambda)$.*

Proof. Since L is normal to $F \cap \Sigma$ at x , n_x lies in the plane containing L perpendicular to F . This implies that $\angle L, n_x$ is the acute angle between F and n_x , which is less than $\alpha(\lambda) + \beta(\lambda)$ by Lemma 4.4(ii). \square

Lemma 4.6 *Let p be a sample point. Let x and y be two points on Σ . If $\angle xy, n_x \leq \alpha(\lambda) + \beta(\lambda)$, then $\|p - x\|$ or $\|p - y\|$ is greater than $\lambda f(p)$.*

Proof. Assume to the contrary that both $\|p - x\|$ and $\|p - y\|$ are at most $\lambda f(p)$. By Lemma 4.2, $\|x - y\| \leq 2\lambda f(x)/(1 - \lambda)$. On the other hand, by Lemma 2.3, $\|x - y\| \geq 2f(x) \cos(\angle xy, n_x) \geq 2f(x) \cos(\alpha(\lambda) + \beta(\lambda))$. The lower and upper bounds on $\|x - y\|$ together yield $\lambda/(1 - \lambda) \geq \cos(\alpha(\lambda) + \beta(\lambda))$, contradicting $\lambda \leq \lambda_0$. \square

4.2 Violation at Voronoi edges

If a Voronoi edge violates GIP or TBP, it intersects Σ in some point far away from existing sample points. The next lemma establishes this fact.

Lemma 4.7 *Let e be an edge of a Voronoi cell V_p .*

- (i) *If $e \cap \Sigma$ contains two points, the distance between p and the further one is at least $\lambda f(p)$.*
- (ii) *If the support line of e meets Σ tangentially at a point $x \in e \cap \Sigma$, then $\|p - x\| \geq \lambda f(p)$.*

Proof. Consider the case in which $e \cap \Sigma$ contains two points x and y (Figures 2(a) and 2(b)). If $\|x - y\| \geq \lambda f(p)$, we are done. If $\|x - y\| < \lambda f(p)$, $\angle xy, n_x = \angle e, n_x \leq \alpha(\lambda) + \beta(\lambda)$ by Lemma 4.3. Then, Lemma 4.6 implies that $\|p - y\| > \lambda f(p)$.

Consider the case in which the support line of e intersects Σ tangentially at a point $x \in e \cap \Sigma$ (Figure 2(c)). If $\|p - x\| < \lambda f(p)$, then $\angle e, n_x < \pi/3$ by Lemma 4.3. This is impossible because the support line of e meets Σ tangentially at x . \square

4.3 Violation at Voronoi facets

We first characterize the intersection between Σ and a Voronoi facet F . Let $\text{Int } \mathbb{X}$ denote the interior of a topological space \mathbb{X} . There are three possible configurations for any connected




Figure 2: The curves are subsets of Σ . A Voronoi edge (shown in bold) intersects Σ at two points or tangentially.

component I in $\Sigma \cap F$. First, I may contain a tangential contact point between Σ and the plane of F . Second, I may be a smooth closed curve. Third, I may be a smooth open curve (possibly degenerate). In the third possibility, if I is not a single point, $\text{Int } I \subseteq \text{Int } F$ and the endpoints of I lie on the boundary of F . In this case, we call I a *topological interval*. If I is a single point, Σ barely cuts F at a single vertex v (e.g., Σ cuts across F at its topmost vertex). Then I is just v and we call I a *degenerate topological interval*.

The GIP and TBP require that F intersects Σ transversally, if at all, and in at most one topological interval. Recall that transversal intersection permits F (but not its *affine plane*) to meet Σ tangentially at a Voronoi vertex. Also we let GIP and TBP be violated for Voronoi vertices allowing them to be on Σ . This means we would allow a facet F meeting Σ only at a single topological interval degenerate or not. In all other cases we show that a new point can be sampled far away from existing sample points. First, we show that any tangential contact point between Σ and a Voronoi facet is far away from existing sample points.

Lemma 4.8 *Let F be a facet of a Voronoi cell V_p . If the plane of F meets Σ tangentially at a point $x \in F \cap \Sigma$, then $\|p - x\| \geq \lambda f(p)$.*

Proof. Since x is a tangential contact point, the angle between F and n_x is equal to $\pi/2$. Then the contrapositive of Lemma 4.4(ii) implies that $\|p - x\| \geq \lambda f(p)$. \square

Because of Lemma 4.8, in the rest of this section, we focus on the case where a Voronoi facet F meets Σ transversally, i.e., $F \cap \Sigma$ is a collection of closed curves and/or (possibly degenerate) open curves.

Let L be a line in the plane of F which is normal to a curve in $F \cap \Sigma$. In the next lemma we establish that if L intersects $F \cap \Sigma$ at two or more points, we can find a point in $F \cap \Sigma$ that is far away from existing sample points. Notice that this result holds irrespective of whether $F \cap \Sigma$ is a closed curve or not. Although our main motivation is to get rid of closed curves in $F \cap \Sigma$, it is algorithmically easier to check if L intersects $F \cap \Sigma$ in at least two points instead of whether $F \cap \Sigma$ contains a closed curve.

Lemma 4.9 *Let F be a facet of a Voronoi cell V_p where F intersects Σ transversally. Let x be a point on a curve C (possibly closed) in $F \cap \Sigma$. Let L be the line in the plane of F that*

is normal to C at x . If L intersects $F \cap \Sigma$ at a point other than x , the distance from p to the furthest point in $L \cap F \cap \Sigma$ is at least $\lambda f(p)$.

Proof. Refer to Figure 3. By assumption, there is a point y other than x in $L \cap F \cap \Sigma$. If $\|p - x\| \geq \lambda f(p)$, we are done. If $\|p - x\| < \lambda f(p)$, $\angle xy, n_x = \angle L, n_x < \alpha(\lambda) + \beta(\lambda)$ by




Figure 3: The line L is normal to the curve at x and L intersects $F \cap \Sigma$ at another point y .

Lemma 4.5. Then, Lemma 4.6 implies that $\|p - y\| > \lambda f(p)$. \square

The next lemma considers the scenario of F intersecting Σ in two topological intervals, assuming that each edge of F intersects Σ in at most one point. We show that at least one endpoint of the two topological intervals is far away from existing sample points.

Lemma 4.10 *Let F be a facet of a Voronoi cell V_p where F intersects Σ transversally. Assume that each edge of F intersects Σ in at most one point. If $F \cap \Sigma$ contains two topological intervals I and I' , the distance between p and the furthest endpoint of I and I' is at least $\lambda f(p)$.*

Proof. Let u and v be the endpoints of I . Let x and y be the endpoints of I' . Since no edge of F intersects Σ in two or more points, the four edges of F containing u , v , x , and y are distinct. Let Q be the convex quadrilateral on the plane of F bounded by the support lines of these four edges. We denote the edges of Q by e_u , e_v , e_x , and e_y according to the interval endpoints that the edges contain. (In the degenerate case in which $I = u = v$, the two edges of F incident to u give rise to two edges of Q . We arbitrarily call one of them e_u and the other e_v . The degenerate case in which $I' = x = y$ is handled similarly.) Refer to Figure 4(a).

Assume to the contrary that the distances $\|p - u\|$, $\|p - v\|$, $\|p - x\|$, and $\|p - y\|$ are less than $\lambda f(p)$. Consider the Delaunay triangles dual to the edges of F containing u , v , x , and y . Their circumradii are less than $\lambda f(p)$. By Lemma 2.2, the angles $\angle e_u, n_p$, $\angle e_v, n_p$, $\angle e_x, n_p$, and $\angle e_y, n_p$ are at most $\beta(\lambda)$.

By Lemma 4.4(i), the acute angle between the plane of F and n_p is less than $\beta(\lambda)$.

Let d be the projection of n_p onto the plane of F . So $\angle n_p, d < \beta(\lambda)$. It follows that $\angle e_u, d \leq \angle e_u, n_p + \angle n_p, d < 2\beta(\lambda)$. Similarly, the angles $\angle e_v, d$, $\angle e_x, d$, and $\angle e_y, d$ are less than $2\beta(\lambda)$. Then the convexity of Q implies that one of its interior angles must be greater than $\pi - 4\beta(\lambda)$, say the interior angle between e_v and e_x . In this case, a line parallel to d may either cut through or be tangent to the corner of Q between e_v and e_x . Figures 4(b) and 4(c) show the two possibilities. In both configurations, $\angle vx, e_v < 2\beta(\lambda)$ or $\angle vx, e_x < 2\beta(\lambda)$. Assume that $\angle vx, e_x < 2\beta(\lambda)$.




Figure 4: A Voronoi facet (bounded by solid line segments) intersects Σ in two topological intervals (shown as curves). The convex quadrilateral Q is shown in dashed line segments. In (b) and (c), the direction d is the projection of n_p onto the plane of F .

By Lemma 4.3, $\angle e_x, n_x \leq \alpha(\lambda) + \beta(\lambda)$. We conclude that $\angle vx, n_x \leq \angle vx, e_x + \angle e_x, n_x < \alpha(\lambda) + 3\beta(\lambda)$. On the other hand, $\|v - x\| \leq 2\lambda f(x)/(1 - \lambda)$ by Lemma 4.2. Then Lemma 2.3 implies that $\angle vx, n_x \geq \arccos(\lambda/(1 - \lambda))$. Combining this with the previous upper bound on $\angle vx, n_x$ yields $\arccos(\lambda/(1 - \lambda)) < \alpha(\lambda) + 3\beta(\lambda)$ and hence $\lambda/(1 - \lambda) > \cos(\alpha(\lambda) + 3\beta(\lambda))$. But this contradicts the fact that $\lambda \leq \lambda_0$. \square

4.4 Violation at Voronoi cells

The TBP requires that $V_p \cap \Sigma$ is a topological disk. There are several possibilities when this condition is violated. We assume that Σ meets the edges or facets of V_p transversally because tangential contacts are already handled by Lemma 4.7 and Lemma 4.8.

- $V_p \cap \Sigma$ has more than one boundary cycle.
- Some connected component of $V_p \cap \Sigma$ is a surface without boundary.
- $V_p \cap \Sigma$ has non-zero genus.

Notice that $V_p \cap \Sigma$ is orientable because Σ is orientable. We deal with the first possibility in Section 4.4.1. The second possibility is eliminated because the initialization in our algorithm inserts all critical points of Σ in the x_3 -direction. Any component of Σ would contain two such critical points which would violate the emptiness of V_p . It turns out that the last possibility is hard to detect. In Section 4.4.2, we propose to enforce a stronger condition based on the notion of *silhouette*. Detecting the violation of this stronger condition is easier and a new point can be sampled readily in case of a violation.

Since we allow Voronoi vertices to lie on Σ , we give a definition of boundary cycles of $V_p \cap \Sigma$ that capture degenerate cases as well. Let $\text{Bd } \mathbb{X}$ denote the boundary of a topological space \mathbb{X} . Consider a connected component C in $\Sigma \cap \text{Bd } V_p$. The non-tangential contact assumption implies that C is either a non-degenerate closed curve or a vertex of V_p . If C is a non-degenerate closed curve, C is clearly a boundary cycle of $V_p \cap \Sigma$. The case of C being a vertex v of V_p happens when Σ barely cuts V_p at v . We consider v as a *degenerate boundary cycle* of $V_p \cap \Sigma$.

4.4.1 Two or more boundary cycles

We first prove a technical result as stated in Lemma 4.11. This lemma says that the distance from p to any point in a topological interval is dominated by its distances to the interval endpoints. (Recall that a topological interval is allowed to degenerate to a vertex of V_p .)

Lemma 4.11 *Let F be a facet of a Voronoi cell V_p . Suppose that Σ meets edges of F transversally and $F \cap \Sigma$ contains a topological interval I . If the distances from p to the endpoints of I are less than $\lambda f(p)$, the distance from p to any point in I is less than $\lambda f(p)$.*

Proof. The lemma is trivially true if I degenerates to a vertex of V_p . So we can assume that I has two distinct endpoints.

Let B be the ball centered at p with radius $\lambda f(p)$. Let Π denote the plane of F . Since the distances from p to the endpoints of I are less than $\lambda f(p)$, $B \cap \Pi$ is a disk D and the endpoints of I lie in $\text{Int } D$. To prove the lemma, it suffices to show that $I \subseteq \text{Int } D$. Assume to the contrary that $I \not\subseteq \text{Int } D$.

Let C be the connected component of $\Pi \cap \Sigma$ containing I . Notice that $C \not\subseteq \text{Int } D$ because $I \subset C$ and $I \not\subseteq \text{Int } D$ by assumption. For any point $x \in \Pi \cap \Sigma \cap D$, since $\|p - x\| \leq \lambda f(p)$, x cannot be a tangential contact point between Π and Σ by Lemma 4.4(ii). Thus, $C \cap \text{Int } D$ is a collection of disjoint simple curves (open or closed).

We claim that $C \cap \text{Int } D$ is not connected. Assume to the contrary that $C \cap \text{Int } D$ is connected, i.e., $C \cap \text{Int } D$ is a single curve. Recall that $I \subset C$, $I \not\subseteq \text{Int } D$, and the endpoints of I lie in $C \cap \text{Int } D$. It follows that $(C \cap \text{Int } D) \cup I$ is a closed curve. Take an edge e of F that contains an endpoint x of I . Since Σ meets e transversally by assumption, the support line ℓ of e crosses C . Because ℓ intersects I only at x , ℓ must intersect $C \cap \text{Int } D$ at x and at least one other point y . Since $x, y \in D$, both $\|p - x\|$ and $\|p - y\|$ are at most $\lambda f(p)$. By Lemma 4.3, $\angle xy, n_x = \angle e, n_x \leq \alpha(\lambda) + \beta(\lambda)$. But then Lemma 4.6 implies that $\|p - x\|$ or $\|p - y\|$ is greater than $\lambda f(p)$, a contradiction.

So we can assume that $C \cap \text{Int } D$ consists of at least two disjoint curves. Then, D can be shrunk to a smaller disk D' as follows so that D' meets C tangentially at two points and $C \cap \text{Int } D' = \emptyset$. First, shrink D radially until it touches C at some point a . Refer to Figure 6(a). It follows that $\|p - a\| < \lambda f(p)$. If this shrunk D does not meet the requirement of D' yet, we shrink it further by moving its center towards a until we obtain the disk D' as required. Refer to Figure 5(b). Notice that a is one of the contact points between D' and C .

The plane Π intersects the two medial balls of Σ at a in two disks. Among these two disks, let D'' be the one that intersects D' . Let B'' be the medial ball so that $D'' = B'' \cap \Pi$. The boundary circles of D' and D'' meet tangentially at a . So either $D'' \subseteq D'$ (Figure 6(a)) or $D' \subset D''$ (Figure 6(b)).

We claim that $D'' \subseteq D'$ and $\text{radius}(D'') > \lambda f(p)$. Suppose that $D' \subset D''$. By construction, D' meets Σ tangentially at two points. So one of these contact points must lie in $\text{Int } D''$. This is a contradiction because $D'' = B'' \cap \Pi$ and $\text{Int } B'' \cap \Sigma = \emptyset$ as B'' is a medial ball. This shows that $D'' \subseteq D'$. By Lemma 4.4(ii), the acute angle between Π and n_a is less than $\alpha(\lambda) + \beta(\lambda)$. Observe that the angle between the diametric segments of B'' and D'' incident to a is equal to the angle between n_a and Π . Therefore, $\text{radius}(D'') > \text{radius}(B'') \cdot \cos(\alpha(\lambda) + \beta(\lambda)) \geq f(a) \cdot \cos(\alpha(\lambda) + \beta(\lambda)) > \lambda f(a)/(1 - \lambda)$ as $\lambda \leq \lambda_0$. It follows from Lemma 4.1 that $\text{radius}(D'') > \lambda f(p)$. This completes the proof of our claim.




Figure 5: Shrinking D to D' : (a) shows the radial shrinking and (b) shows the shrinking by moving the center.




Figure 6: D' and D'' .

By our claim, $\text{radius}(D') \geq \text{radius}(D'') > \lambda f(p)$. But D' is obtained by shrinking $D = B \cap \Pi$ and $\text{radius}(B) = \lambda f(p)$, a contradiction. In all, the contrapositive assumption that $I \not\subseteq \text{Int } D$ cannot hold. It follows that the distance from p to any point in I is less than $\lambda f(p)$. \square

Lemma 4.11 is used in proving Lemma 4.12, which says that if $V_p \cap \Sigma$ has more than one boundary cycle, some edge of V_p intersects Σ in a point far away from existing sample points. It is convenient to distinguish between different types of cycles in $V_p \cap \Sigma$. A boundary cycle is of *type 1* if it is degenerate or a concatenation of topological intervals in the intersections between Σ and the facets of V_p . A boundary cycle is of *type 2* if it is non-degenerate and contained in a facet of V_p .

Lemma 4.12 *Let p be a sample point. Assume that the following conditions hold.*

- Σ meets edges or facets of V_p transversally.
- $V_p \cap \Sigma$ contains at least two boundary cycles of type 1.

Then the distance from p to the furthest intersection point between Σ and the edges of V_p is at least $\lambda f(p)$.

Proof. Assume to the contrary that the distances from p to the intersection points between Σ and the edges of V_p are less than $\lambda f(p)$. This will lead to contradictions thereby proving the lemma. We first prove a technical result that will be used later.

Claim 1 *Let $x \in \Sigma$ be a vertex of V_p such that $\|p - x\| \leq \lambda f(p)$. Let E be a set of edges of V_p incident to x that point towards the same side of Σ . The smallest cone enclosing E with apex x and axis n_x cannot contain a point $y \in \Sigma$ other than x such that $\|p - y\| \leq \lambda f(p)$.*

Proof. By Lemma 2.2, each edge in E makes an angle at most $\beta(\lambda)$ with n_p . By Lemma 2.1, $\angle n_x, n_p \leq \alpha(\lambda)$. Therefore, each edge in E makes an angle at most $\alpha(\lambda) + \beta(\lambda)$ with n_x . So the aperture of the smallest cone enclosing E with apex x and axis n_x is at most $2\alpha(\lambda) + 2\beta(\lambda)$. Let $y \in \Sigma$ be any point other than x in this cone. Thus, $\angle n_x, xy \leq \alpha(\lambda) + \beta(\lambda) < \pi/3$ as $\lambda \leq \lambda_0$. Then Lemmas 2.3 and 4.1 imply that $\|x - y\| \geq 2f(x) \cos(\angle n_x, xy) \geq f(x) \geq (1 - \lambda)f(p)$. Since $\|p - y\| \geq \|x - y\| - \|p - x\| \geq (1 - 2\lambda)f(p)$, $\|p - y\| \geq \lambda f(p)$ as $\lambda \leq \lambda_0 = 0.06$. \square

Lemma 4.11 implies that the boundary cycles of type 1 lie strictly inside a closed ball B centered at p with radius $\lambda f(p)$. By a result of Boissonnat and Cazals [5], any closed ball centered at p with radius less than $f(p)$ intersects Σ in a topological disk. Thus $B \cap \Sigma$ is a topological disk. It follows that each non-degenerate boundary cycle of type 1 bounds exactly one topological disk in $B \cap \Sigma$ (strictly inside B). Since the boundary cycles are disjoint, the topological disks bounded by them are either disjoint or are nested. This means there exists one such topological disk that does not contain any cycle of type 1 inside. The next claim proves some properties of such a topological disk.

Claim 2 *Let C be a boundary cycle of type 1 bounding a topological disk D which does not contain any other boundary cycle of type 1. Then, (i) C is a non-degenerate boundary cycle, (ii.a) D does not contain any other boundary cycles and (ii.b) D lies in V_p .*

Proof. Consider (i). If not, C is an isolated vertex v of V_p in the intersection $V_p \cap \Sigma$. The edges of V_p incident to v must point towards the same side of Σ . Indeed, otherwise, Σ intersects the interior of V_p in a small neighborhood of v , contradicting the assumption that v is an isolated point in the intersection $V_p \cap \Sigma$. On the other hand, Claim 1 is contradicted because $\|p - v\| < \lambda f(p)$ and p lies inside the smallest cone with apex v and axis n_v that encloses the edges of V_p incident to v . This proves (i).

Consider (ii.a). By the definition of C , D does not contain other boundary cycles of type 1. Assume to the contrary that D contains a boundary cycle C' of type 2. So C' is contained in a facet of V_p . Since D lies strictly inside B , C' lies strictly inside B . But by applying Lemma 4.9 to C' , the distance from p to some point in C' is at least $\lambda f(p) = \text{radius}(B)$, a contradiction.

Consider (ii.b). It is sufficient to show that $\text{Int } D$ lies in V_p . Suppose not. Then $\text{Int } D$ lies completely outside V_p . Otherwise, $\text{Int } D$ would contain a boundary cycle which is prohibited by (ii.a). Refer to Figure 7. Take an edge e of V_p that intersects C . Let x be an intersection point between e and C . By Lemma 4.3, $\angle e, n_x < \pi/3$.




Figure 7: Proof of Claim 2: disk D with the opening C is like a ‘sack’ that contains a part of V_p inside.

Thus, the support line of e intersects Σ transversally at x . Let ℓ be a line outside V_p that is parallel to and arbitrarily close to the support line of e . Then ℓ must intersect $\text{Int } D$ transversally at a point x_1 arbitrarily close to x .

Since $\text{Bd } V_p$ is a topological sphere, it is divided by C into two topological disks. Let T denote one of them. Then $T \cup D$ is a topological sphere. Since ℓ intersects $\text{Int } D$ at x_1 , ℓ must intersect $T \cup D$ at another point $x_2 \neq x_1$.

The point x_2 must lie on D because $T \subseteq \text{Bd } V_p$ and ℓ lies outside V_p . By Lemma 2.3, $\|x_2 - x_1\| \geq 2f(x_1) \cos(\angle \ell, n_{x_1})$. Note that x_1 is arbitrarily close to x and $\angle \ell, n_x = \angle e, n_x < \pi/3$. Thus, $\angle \ell, n_{x_1} < \pi/3$ and so $\|x_2 - x_1\| \geq f(x) \geq (1 - \lambda)f(p)$ by Lemma 4.1. But this is a contradiction because D lies inside B whose diameter is $2\lambda f(p) < (1 - \lambda)f(p)$. \square

Let C be a boundary cycle as stated in Claim 2. Let D be the topological disk in $B \cap \Sigma$ bounded by C . Let \mathcal{F} be the set of facets of V_p that intersect C . Each facet in \mathcal{F} bounds a half-space containing p . The intersection of these half-spaces is a convex polytope P containing V_p . Since D lies in V_p by Claim 2, D lies in P too.

Let C' be another boundary cycle of type 1 which must exist by the assumption of the lemma. Recall that both C and C' lie inside $B \cap \Sigma$ by Lemma 4.11. Let ρ be a curve in $(B \cap \Sigma) \setminus \text{Int } V_p$ such that ρ connects C with C' . Since $D \subset P$ and the contact between D and $\text{Bd } P$ is non-tangential, ρ leaves P when ρ leaves D . Since $C' \subset V_p \subseteq P$, ρ must return to some facet of P in order to meet C' eventually. Let G be a facet of P that ρ intersects after leaving D . Let y be a point in $\rho \cap G$. Let F be the facet of V_p contained in G . By the definition of P , C must intersect F .

If $C \cap F$ consists of two or more topological intervals, Lemma 4.10 is applicable and we are done. So we assume in the rest of the proof that $C \cap F$ is a single topological interval I . Refer to Figure 8.

Claim 3 *Each edge of F that contains an endpoint of I is contained in some edge of G .*

Proof. Consider an endpoint z of I . The point z lies on the boundary of F , which means that the other facet(s) of V_p that share z with F are intersected by C . So




Figure 8: (a) Two cycles C and C' drawn schematically on the patch $B \cap \Sigma$. The path ρ starting from C goes outside V_p and then has to reach V_p again to reach C' . (b) A different view with the polyhedron P . The lower bold curve denotes C and its intersection with the shaded facet G is a topological interval I . The curved patch shown is part of $(B \cap \Sigma) \setminus \text{Int } V_p$. The curved path on it is ρ .

the planes of these facets also bound P . It follows that the edges of F containing z are contained in some edges of G . □

It follows from Claim 3 that the endpoints of I lie on the boundary of G . Let x be the closest point to y on I . Notice that $x \in I \subset C \subset B \cap \Sigma$ and $y \in \rho \subset B \cap \Sigma$. So the distances from p to x and y are at most $\text{radius}(B) = \lambda f(p)$. Let L be the line passing through x and y . There are three cases to consider.

- Case 1: I is a degenerate interval. So $I = x$ is a vertex of F . By Claim 3, the two edges of F incident to x are contained in edges of G . So x is also a vertex of G . Let e_1 and e_2 be the two edges of G incident to x (which contain the edges of F incident to x). Since I is a degenerate interval, e_1 and e_2 must point towards the same side of Σ . But then Claim 1 is contradicted because y lies inside the smallest cone with apex x and axis n_x that encloses e_1 and e_2 .
- Case 2: I is non-degenerate and x lies in the interior of G . Then L intersects I at x at right angle. By Lemma 4.5, $\angle L, n_x < \alpha(\lambda) + \beta(\lambda)$. But then $\|p - x\|$ or $\|p - y\|$ is greater than $\lambda f(p)$ by Lemma 4.6, a contradiction.
- Case 3: I is non-degenerate and x lies on the boundary of G . Let e be an edge of G containing x . By Claim 3, e contains an edge of F which contains x . Let d be the projection of n_x onto the plane of G . Refer to Figure 9(a). Since x is the closest point on I to y , $\angle L, d \leq \angle e, d$. Next, observe that for any line ℓ in the plane of G , the angle $\angle \ell, n_x$ increases as the angle $\angle \ell, d$ increases. Refer to Figure 9(b). We conclude that $\angle L, n_x \leq \angle e, n_x$, which is at most $\alpha(\lambda) + \beta(\lambda)$ by Lemma 4.3. But then $\|p - x\|$ or $\|p - y\|$ is greater than $\lambda f(p)$ by Lemma 4.6, a contradiction.




Figure 9: In (a), the curve denotes $C \cap F$ and $\angle L, d \leq \angle e, d$. In (b), the angle $\angle l, n_x$ is an increasing function of $\angle l, d$.

□

4.4.2 Silhouette

Consider the possibility of some connected component of $V_p \cap \Sigma$ being a closed surface. This closed surface must contain two critical points of Σ in the x_3 -direction. Therefore, this possibility is eliminated by our algorithm because we start by including all critical points of Σ in the x_3 -direction as sample points.

The remaining possible violation of TBP is that $V_p \cap \Sigma$ is connected, has exactly one boundary cycle, and has positive genus. Intuitively, $V_p \cap \Sigma$ contains a handle and Morse theory says that there should be a critical point of Σ in the x_3 -direction. However, it is not guaranteed that this critical point lies inside V_p . (If it were, the initialization in our algorithm would eliminate this possibility.)




Figure 10: The closed dashed curve cuts a topological disk from the torus which contains all four critical points of the torus: a maximum (black), a minimum (gray), and two saddles (white).

Figure 10 illustrates this possibility. A closed dashed curve cuts a topological disk from the torus. The parts of the topological disk in front are shown shaded and the rest of the topological disk is at the back. All the critical points of the torus lie on the topological disk.

Suppose that the facets of a Voronoi cell V_p intersects a torus at this closed dashed curve such that the bounded topological disk lies outside V_p and the rest of the torus lies inside V_p . Then V_p violates TBP although the surface clipped within V_p is connected, has a single boundary cycle, and has no critical point. Apparently it may seem that a convex polytope cannot intersect the standard torus in a closed curve as shown in the figure. However, one may imagine deforming the torus to admit a closed curve with the stated property.

We propose a stronger condition to ensure that $V_p \cap \Sigma$ is a topological disk. The violation of this stronger condition can be detected and a new point can be sampled readily.

Lemma 4.13 below states that if a connected component M of $V_p \cap \Sigma$ is connected, has exactly one boundary cycle, and avoids J_d , then M is a topological disk.

Lemma 4.13 *Let M be a connected component of $V_p \cap \Sigma$ such that the boundary of M is a single boundary cycle. If the boundary cycle of M is degenerate or there is a direction d such that $M \cap J_d = \emptyset$, M is a topological disk.*

Proof. If the boundary cycle of M is degenerate, M is a vertex of V_p and so M is a degenerate topological disk. Otherwise, M is a connected compact 2-manifold and its boundary is a simple closed curve by the assumption of the lemma. Let d be a direction satisfying the condition of the lemma. Let H be a plane perpendicular to d . Consider the map $\varphi : M \rightarrow H$ that projects each point of M orthogonally to H . Since M is connected and compact and M has a single boundary cycle, it suffices to prove that φ is injective.

Assume to the contrary that φ is not injective. Then there is a line L parallel to d that intersects M in two or more points. Let x and y be two consecutive intersection points along L . Let M' be the connected component of Σ containing M . By the convexity of V_p , x and y are consecutive intersection points in $L \cap M'$ too.

Since $M \cap J_d = \emptyset$, neither x nor y belongs to J_d , which means that neither n_x nor n_y is orthogonal to d . Because x and y are consecutive in $L \cap M'$, n_x and n_y are oppositely oriented in the sense that the inner products $\langle n_x, d \rangle$ and $\langle n_y, d \rangle$ have opposite signs. Since M is connected, there is a smooth curve ρ in M connecting x and y . The normal to Σ changes smoothly from n_x to n_y along ρ . By the mean-value theorem, there is a point $z \in \rho$ such that n_z is orthogonal to d . But then $z \in M \cap J_d$, contradicting the emptiness of $M \cap J_d$. \square

By Lemma 4.13, when we are left with the case that $V_p \cap \Sigma$ is connected and has a single boundary cycle, it suffices to check whether $V_p \cap \Sigma$ intersects J_d where $d = n_p$. If not, $V_p \cap \Sigma$ is a topological disk. Otherwise, the following result says that any point in $V_p \cap J_d$ can be inserted as a new sample point.

Lemma 4.14 *Let p be a sample point. Let $d = n_p$. If x is a point in $V_p \cap J_d$, $\|p - x\| \geq \lambda f(p)$.*

Proof. If $\|p - x\| < \lambda f(p)$, Lemma 2.1 would imply that $\angle n_p, n_x \leq \alpha(\lambda)$. But this cannot be the case because n_x is orthogonal to n_p by definition. \square

5 Topology recovery

In this section, we present an algorithm to sample a point set S from Σ . We discuss how to handle Voronoi vertices lying on Σ in Section 5.1, which is the remaining violation of GIP that

we did not address in the previous section. When some Voronoi vertices lie on Σ , $\text{Del } S|_{\Sigma}$ contains their dual Delaunay tetrahedron. The handling of this case boils down to extracting a triangulated surface $\text{Tri } S|_{\Sigma}$ from $\text{Del } S|_{\Sigma}$ to approximate Σ . The definition of $\text{Tri } S|_{\Sigma}$ is given in Section 5.1. In Section 5.2, we present some numerical primitives needed by our algorithm. Section 5.3 describes several subroutines. These subroutines sample points from Σ based on the results in the Section 4. In Section 5.4, we put these subroutines together to form the topology recovering algorithm. Section 5.5 describes the analysis of the algorithm.

5.1 Handling Voronoi vertices on the surface

Our idea is to conceptually perturb Σ to obtain another surface Σ' so that no Voronoi vertex lies on Σ' . The perturbation is kept small so that Σ remains homeomorphic to Σ' .

We elaborate on the conceptual perturbation. Let S be the current set of sample points on Σ . Let U be the subset of Voronoi vertices of $\text{Vor } S$ that lie on Σ . Let $B_{v,\delta}$ denote the ball centered at $v \in U$ with radius δ . Refer to Figure 11(a). First, choose δ small enough so that the following conditions hold for each $v \in U$:

- C1: $\Sigma \cap B_{v,\delta}$ is a topological disk.
- C2: $B_{v,\delta}$ does not intersect $B_{w,\delta}$ for any other Voronoi vertex $w \in U$.
- C3: $B_{v,\delta}$ intersects only the Voronoi edges and facets incident to v .
- C4: Within $B_{v,\delta}$, the Voronoi edges incident to v intersect Σ only at v .

Condition C1 guarantees that for each $v \in U$, $\Sigma \cap B_{v,\delta}$ separates $B_{v,\delta}$ into two regions, one on each side of Σ . Let R_v denote the region inside Σ . Let D_v denote the portion of the boundary of $B_{v,\delta}$ inside R_v . Notice that D_v is a topological disk.

We obtain a new surface by replacing $\Sigma \cap B_{v,\delta}$ with D_v for every Voronoi vertex $v \in U$. Refer to Figure 11(b). The new surface needs to be smoothed at the sharp boundary of D_v . The smoothing can be restricted in an arbitrarily small neighborhood of the boundary of D_v by introducing arbitrarily high curvature. Hence, it has no effect on the conceptual perturbation. Let Σ' denote the resulting surface.

Condition C2 guarantees that the above disk replacements can be performed simultaneously for all vertices in U . Condition C3 guarantees that only Voronoi edges, facets, and cells incident to v are affected by the perturbation.

The idea is to use $\text{Del } S|_{\Sigma'}$ instead of $\text{Del } S|_{\Sigma}$ as the triangulation approximating Σ . We do not actually perturb Σ to obtain a new surface Σ' . Instead, we perform a simulation. This allows us to exclude some triangles from $\text{Del } S|_{\Sigma}$ to obtain $\text{Del } S|_{\Sigma'}$. The details of the simulation are as follows.

Let uv be a Voronoi edge incident to v . Since our topology recovery algorithm will first check for tangential contacts between Σ and the Voronoi edges, we can assume that uv is not orthogonal to n_v . Condition C4 means that uv does *not* intersect Σ' if $\langle u - v, n_v \rangle > 0$. (Recall that n_v is the unit outward normal at v .) In this case, if uv does not intersect Σ in a point other than v , we exclude the dual Delaunay triangle of uv . On the other hand, if uv intersects Σ in a point other than v , we keep its dual Delaunay triangle.

Repeating the above for every $v \in U$ yields $\text{Del } S|_{\Sigma'}$. Since we compute with Σ instead of Σ' , it is more concrete to have a notation for $\text{Del } S|_{\Sigma'}$ using Σ . We use $\text{Tri } S|_{\Sigma}$ to denote




Figure 11: The solid curve denotes Σ . The shaded area denotes the inside of Σ . The bold curve on the right denotes the topological disk D_v . The Voronoi edge e' is treated as not intersecting Σ at v , while the Voronoi edge e is treated as intersecting Σ at v .

$\text{Del } S|_{\Sigma'}$. The surface Σ' cannot introduce new tangential intersections. In the rest of this section, we argue that if $\text{Vor } S$ violates TBP with Σ' , it violates TBP with Σ too. Therefore, when Σ intersects the edges and facets of $\text{Vor } S$ transversally and $\text{Vor } S$ satisfies TBP with Σ , $\text{Vor } S$ satisfies GIP and TBP with Σ' . Applying Theorem 2.1 we obtain $\text{Tri } S|_{\Sigma} = \text{Del } S|_{\Sigma'}$ is homeomorphic to Σ .

Take a Voronoi edge e incident to $v \in U$. Let v' be the point near v in which Σ' intersects e after perturbation. Suppose that e violates TBP with respect to Σ' . The violation means that e intersects Σ' in v' and at least one other point. It follows that e intersects Σ in v and at least one other point. So e violates TBP with respect to Σ too.

Take a Voronoi facet F incident to $v \in U$. If Σ cuts through F at v , the topology of $F \cap \Sigma$ does not change after perturbation at v . The interesting case is that Σ' intersect both edges of F incident to v at points near v after perturbation. In this case, $F \cap \Sigma'$ is not just a geometric perturbation of $F \cap \Sigma$ because $F \cap \Sigma'$ and $F \cap \Sigma$ have different topologies. Specifically, $F \cap \Sigma'$ contains a short topological interval I' near v . If I' participates in a violation of TBP with respect to Σ' , $F \cap \Sigma'$ must contain a connected component other than I' . Then, F intersects Σ in at least two components, one of them being the degenerate topological interval v . So F violates TBP with respect to Σ too.

Similarly, consider a Voronoi cell V_p incident to $v \in U$. The interesting case is that Σ' intersects all edges of V_p incident to v at points near v after perturbation. So $V_p \cap \Sigma'$ contains a small topological disk D' near v . If D' participates in a violation of TBP with respect to Σ' , $V_p \cap \Sigma'$ must have a connected component other than D' . It follows that $V_p \cap \Sigma$ has at least two connected components, one of them being the degenerate boundary cycle v . So V_p violates TBP with respect to Σ too.

5.2 Numerical primitives

We introduce six numerical primitives. They are responsible for computing the intersections between Σ and lines and computing the critical points of certain height functions on Σ . We assume a numerical or symbolic solver for solving a system of equations (e.g. [26]).

- **SURFACECRITICAL**(surface Σ): Let d denote the x_3 -direction. This primitive returns the critical points of Σ in the x_3 -direction by returning the solutions of the system of

equations: $E(x) = 0, \nabla E(x) \times d = 0$.

- SURFACELINE(surface Σ , line ℓ): This primitive returns the intersection points between Σ and the line ℓ . We assume that ℓ is given by its direction and a point on it. Using this information, we can compute two planes $H_1(x) = 0$ and $H_2(x) = 0$ whose intersection is equal to ℓ . The intersection points are the solutions of the system of equations: $E(x) = 0, H_1(x) = 0, H_2(x) = 0$.
- SURFACEPLANECONTACT(surface Σ , plane Π): This primitive returns the tangential contact points between Σ and the plane Π . Let $H(x)$ be the equation of Π . Let d be the normal to Π . The contact points are the solutions of the system of equations: $E(x) = 0, H(x) = 0, \nabla E(x) \times d = 0$.
- CURVECRITICAL(surface Σ , plane Π , vector d): We assume that d is parallel to Π . This primitive returns the critical points of $\Sigma \cap \Pi$ in direction d and the tangential contact points in $\Sigma \cap \Pi$. Let d' be a normal of Π . The tangent to a point $x \in \Sigma \cap \Pi$ is parallel to $\nabla E(x) \times d'$. So x is a critical point in direction d if and only if $\langle \nabla E(x) \times d', d \rangle = 0$. Let $H(x)$ be the equation of Π . The critical points desired satisfy the system of equations: $E(x) = 0, H(x) = 0, \langle \nabla E(x) \times d', d \rangle = 0$. The same system captures all tangential contact points $x \in \Sigma \cap \Pi$ too because $\nabla E(x) \times d' = 0$ in this case.
- SILHOUETTEPLANE(surface Σ , plane Π , vector d): Let J_d be the silhouette of Σ in direction d . Let $H(x)$ be the equation of Π . This primitive returns the intersection points between J_d and Π , which are the solutions of the system of equations: $E(x) = 0, H(x) = 0, \langle \nabla E(x), d \rangle = 0$.
- SILHOUETTECRITICAL(surface Σ , vector d , vector d'): Let J_d be the silhouette of Σ in direction d . This primitive returns the critical points of J_d in direction d' , i.e., points on J_d whose tangents are orthogonal to d' . Let $G(x) = \langle \nabla E(x), d \rangle$. The silhouette J_d is the intersection of the isosurface $G(x) = 0$ and the input surface $E(x) = 0$. Therefore, the tangent to a point $x \in J_d$ is parallel to $\nabla G(x) \times \nabla E(x)$. So x is a critical point if and only if $\langle \nabla G(x) \times \nabla E(x), d' \rangle = 0$. The critical points desired are the solutions of the system of equations: $E(x) = 0, G(x) = 0, \langle \nabla G(x) \times \nabla E(x), d' \rangle = 0$.

5.3 Subroutines

We describe five subroutines VOREEDGE, TOPODISK, FACETCONTACT, FACETCYCLE, and SILHOUETTE that implement the results in Section 4. They check for any violation of GIP and TBP at Voronoi edges and facets. In case of violation, they sample new points on Σ . We use S to denote the set of sample points maintained by our algorithm.

The first subroutine VOREEDGE checks the GIP and TBP for a Voronoi edge. In case of violation, it returns a point as stated in Lemma 4.7.

VOREEDGE(Voronoi edge e)

1. Compute $X := \text{SURFACELINE}(\Sigma, \ell)$, where ℓ is the support line of e .
2. Let V_p be a Voronoi cell incident to e .
3. If ℓ meets Σ tangentially at some point x on e , return x .
4. If $|e \cap X| \geq 2$, return the point in $e \cap X$ furthest from p .

5. Return null.

The next subroutine FACETCONTACT detects any violation GIP at a Voronoi facet (Lemma 4.8).

FACETCONTACT(Voronoi facet F)

1. Compute $X := \text{SURFACEPLANECONTACT}(\Sigma, \Pi)$, where Π is the plane of F .
2. If some point in X lies on F , return it. Otherwise, return null.

In the next subroutine TOPODISK, we need to check whether the triangles in $\text{Tri } S|_{\Sigma}$ incident to a sample point p form a topological disk. It is assumed that VOREEDGE has been called and returns null for all Voronoi edges. This implies that Σ meets Voronoi edges transversally, which allows the extraction of triangles of $\text{Tri } S|_{\Sigma}$ using the conceptual perturbation.

TOPODISK performs the checking in two steps. Let T_p be the set of triangles in $\text{Tri } S|_{\Sigma}$ incident to p . First, check if every triangle edge in T_p is incident to exactly two triangles in T_p . Second, check if T_p forms exactly one cycle of triangles around p . If both tests are passed, T_p forms a topological disk; otherwise, it does not.

In the subsequent proof of correctness, we will see that TOPODISK handles two possible violations of TBP. First, a facet of V_p intersects Σ in two topological intervals (Lemma 4.10). Second, $V_p \cap \Sigma$ contains two or more boundary cycles (Lemma 4.12).

TOPODISK(sample point p)

1. If the triangles in $\text{Tri } S|_{\Sigma}$ incident to p form a topological disk, return null.
2. Otherwise, find the intersection point x between Σ and the edges of V_p that is furthest from p . Return x .

The next subroutine FACETCYCLE guards against the possibility of a Voronoi facet F intersecting Σ in a cycle. It assumes that FACETCONTACT(F) has been called and returns null. This implies that Σ intersects F transversally. It returns a point as stated in Lemma 4.9.

FACETCYCLE(Voronoi facet F)

1. Compute $X := \text{CURVECRITICAL}(\Sigma, \Pi, d)$, where Π is the plane of F and d is a direction parallel to Π .
2. If no point in X lies on F , return null.
3. Since FACETCONTACT(F) returned null, $F \cap \Sigma$ is a collection of disjoint simple curves (open or closed) and $X \cap F$ is the set of critical points of these curves in direction d . Let V_p be a Voronoi cell incident to F . For each $x \in X \cap F$,
 - (a) Compute the line ℓ_x in Π through x parallel to d . Notice that ℓ_x is normal to $F \cap \Sigma$ at x .
 - (b) Compute $X' := \text{SURFACELINE}(\Sigma, \ell_x)$. If $|X' \cap F| \geq 2$, return the point in $X' \cap F$ furthest from p .
4. Return null.

The next subroutine SILHOUETTE checks if a Voronoi cell V_p intersects the silhouette J_d , where $d = n_p$ (Lemma 4.14). If $V_p \cap J_d \neq \emptyset$, either J_d intersects some facets of V_p or V_p completely contains a component of J_d . The second possibility can be detected by checking if V_p contains any critical point of J_d in a direction orthogonal to d .

SILHOUETTE(sample point p)

1. Choose a direction d' orthogonal to n_p .
2. Compute $X := \text{SILHOUETTECRITICAL}(\Sigma, n_p, d')$.
3. If X contains a point inside V_p , return it.
4. Otherwise, for each facet F of V_p ,
 - (a) Compute $X' := \text{SILHOUETTEPLANE}(\Sigma, \Pi, n_p)$, where Π is the plane of F .
 - (b) If X' contains a point in F , return it.
5. Return null.

5.4 Topology recovery algorithm

Algorithm SAMPLETOPOLOGY samples a set of points S on Σ so that $\text{Tri}S|_\Sigma$ is homeomorphic to Σ . It begins with initializing S to contain the critical points of Σ in the x_3 -direction. We denote it by S_0 , the *seeds*. Then it calls a procedure TOPOLOGY that repeatedly invokes the subroutines presented in the last subsection in case of any violation of GIP or TBP. Upon the return of TOPOLOGY, $\text{Tri}S|_\Sigma$ is homeomorphic to Σ . However, it is possible that some seeds are too close together, which means that the surface triangulation may be denser than necessary around the seeds. We fix this problem by deleting the seeds incrementally. One may observe that we may start the algorithm with a single initial point and then let SILHOUETTE generate more points on each component of Σ instead of computing the seed set S_0 . However, to avoid the more expensive computation of critical points of the silhouette in a given direction, we recommend starting with the seed set S_0 .

SAMPLETOPOLOGY(surface Σ)

1. Compute $S_0 := \text{SURFACECRITICAL}(\Sigma)$.
2. Compute $S := \text{TOPOLOGY}(S_0)$.
3. While there is a seed $p \in S$, delete p from S and compute $S := \text{TOPOLOGY}(S)$.
4. Return S .

TOPOLOGY(sample set S)

1. Perform steps (a)–(e) in order. Terminate the current step as soon as the returned x is non-null; skip the following steps; and go to step 2.
 - (a) For every edge e of $\text{Vor} S$, compute $x := \text{VOREEDGE}(e)$.
 - (b) For every facet F of $\text{Vor} S$, compute $x := \text{FACETCONTACT}(F)$.
 - (c) For every $p \in S$, compute $x := \text{TOPODISK}(p)$.
 - (d) For every facet F of $\text{Vor} S$, compute $x := \text{FACETCYCLE}(F)$.
 - (e) For every $p \in S$, compute $x := \text{SILHOUETTE}(p)$.
2. If x is non-null, insert x into S , update $\text{Vor} S$, and go to step 1. Otherwise, return S .

5.5 Analysis

We first prove that any point x inserted in step 2 of TOPOLOGY is at distance $\lambda f(p)$ or more from its closest sample p in S . This shows that TOPOLOGY terminates after introducing only finitely many points on Σ .

Lemma 5.1 *Let x be a point inserted by TOPOLOGY into S . The distance from x to its closest sample p in the current S is at least $\lambda f(p)$.*

Proof. If x is returned by VOREEDGE(e) for some edge e of a Voronoi cell V_p , Lemma 4.7 implies that $\|p - x\| \geq \lambda f(p)$.

If x is returned by FACETCONTACT(F) (resp. FACETCYCLE(F)) for some facet F of a Voronoi cell V_p , then $\|p - x\| \geq \lambda f(p)$ by Lemma 4.8 (resp. Lemma 4.9).

If x is returned by TOPODISK(p) for some sample $p \in S$, the set T_p of triangles in $\text{Tri} S|_\Sigma$ incident to p do not form a topological disk. There are three possibilities.

Case 1: Some triangle edge in T_p is not shared by another triangle in T_p . Let $F \subset V_p$ be the dual Voronoi facet of this triangle edge. This case happens because Σ intersects only one boundary edge e of F . Since Σ has no boundary, the endpoint(s) of $F \cap \Sigma$ must lie on e . Thus, e intersects Σ in more than one point or e intersects Σ tangentially. But this is impossible because VOREEDGE(e) did not return any point.

Case 2: Three or more triangles in T_p share an edge. Let $F \subset V_p$ be the dual Voronoi facet of this common edge. This case happens because Σ intersects three or more distinct boundary edges of F . Since VOREEDGE did not return any point for all edges of $\text{Vor} S$, every edge of F intersects Σ in at most one point (non-tangentially). Also, Σ meets F transversally as FACETCONTACT did not return any point. It follows that $F \cap \Sigma$ contains two or more topological intervals. Thus, Lemma 4.10 applies to imply that $\|p - x\| \geq \lambda f(p)$.

Case 3: T_p contains two or more cycles of triangles around p . It implies that $V_p \cap \Sigma$ has at least two boundary cycles C_1 and C_2 , and if C_i is non-degenerate, C_i is not contained in a facet of V_p . Thus, Lemma 4.12 applies to conclude that $\|p - x\| \geq \lambda f(p)$.

If x is returned by SILHOUETTE(p) for some $p \in S$, then $\|p - x\| \geq \lambda f(p)$ by Lemma 4.14. \square

The next result shows a useful property of the Voronoi diagram enforced by step 1 of SAMPLETOPOLOGY.

Lemma 5.2 *If a sample set S includes all critical points of Σ in the x_3 -direction, no connected component of Σ is contained in a Voronoi cell in $\text{Vor} S$.*

Proof. The critical points of any connected component M in the x_3 -direction belong to S by assumption. Each connected component has at least two critical points. Thus, if M is contained inside a Voronoi cell V_p , V_p must contain two seed points. This is a contradiction to the emptiness of V_p even though p may be one of the seed points. \square

We are ready to show that at the end of step 2 of SAMPLETOPOLOGY, GIP and TBP are satisfied.

Lemma 5.3 *Let S_0 be a sample set such that no connected component of Σ is contained in a Voronoi cell in $\text{Vor } S_0$. Let Σ' be the surface obtained by perturbing Σ as presented in Section 5.1. Then TOPOLOGY(S_0) returns a set S such that $\text{Vor } S$ satisfies GIP and TBP with respect to Σ' .*

Proof. By Lemma 5.1 TOPOLOGY(S_0) maintains a positive lower bound on inter-point distances for the points it inserts. This means it can insert only finitely many points on a compact surface Σ . Therefore, it must terminate and return a sample set S . VOREEDGE and FACETCONTACT make sure that Σ intersects all edges and facets of $\text{Vor } S$ transversally. Any vertex of $\text{Vor } S$ on Σ is perturbed by the method presented in Section 5.1. Thus, $\text{Vor } S$ satisfies GIP for Σ' .

VOREEDGE guarantees that Σ does not intersect any Voronoi edge in $\text{Vor } S$ in more than one point at the end of TOPOLOGY. Neither does Σ' .

Consider a Voronoi facet F in $\text{Vor } S$. The intersection $F \cap \Sigma'$ cannot contain more than two or more topological intervals. (Because of the perturbation, any topological interval in $F \cap \Sigma'$ is non-degenerate.) Otherwise, Σ' must intersect at least four boundary edges of F because Σ' intersects any edge of F in at most one point. This means that the dual Delaunay edge of F is incident to more than two triangles in $\text{Tri } S|_{\Sigma}$. But then TOPODISK should have detected this, a contradiction. The intersection $F \cap \Sigma'$ cannot contain any cycle. Otherwise, $F \cap \Sigma$ would contain the same cycle and TOPOLOGY would not have terminated because FACETCYCLE(F) would have returned a point. In all, if $F \cap \Sigma'$ is non-empty, it is a single topological interval.

Consider a Voronoi cell V_p in $\text{Vor } S$. The intersection $V_p \cap \Sigma'$ is a manifold possibly with boundary. By assumption, no Voronoi cell of $\text{Vor } S_0$ contains a connected component of Σ . Hence, no Voronoi cell of $\text{Vor } S_0$ contains a connected component of Σ' too. Since $S_0 \subseteq S$, no connected component of Σ' is contained in any Voronoi cell in $\text{Vor } S$ either.

Can $V_p \cap \Sigma'$ have two or more boundary cycles? If so, $V_p \cap \Sigma$ also has two or more boundary cycles. (Some may degenerate to a vertex of V_p .) VOREEDGE guarantees that the boundary of $V_p \cap \Sigma$ does not intersect the same edge of V_p twice. FACETCYCLE guarantees that no non-degenerate boundary cycle of $V_p \cap \Sigma$ is contained in a facet of V_p . Therefore, the boundary cycles of $V_p \cap \Sigma$ must induce at least two cycles of triangles in $\text{Tri } S|_{\Sigma}$ around p . But then TOPODISK(p) should have detected this, a contradiction.

We conclude that $V_p \cap \Sigma'$ must be connected and have a single boundary cycle. Assume to the contrary that $V_p \cap \Sigma'$ is not a topological disk, i.e., its genus is positive. The perturbation scheme in Section 5.1 does not create any component inside V_p with positive genus. Thus, $V_p \cap \Sigma$ has a connected component M with positive genus. So M is not a single vertex of V_p , i.e., the boundary cycle of M is non-degenerate. Lemma 4.13 implies that M must intersect the silhouette of Σ with respect to direction n_p . But then TOPOLOGY would not have terminated because SILHOUETTE(p) would have returned a point, a contradiction. \square

Corollary 5.1 *Let S_0 be a sample set such that no connected component of Σ is contained in any Voronoi cell in $\text{Vor } S_0$. Then TOPOLOGY(S_0) returns a set S where the underlying space of $\text{Tri } S|_{\Sigma}$ is homeomorphic to Σ .*

Proof. Let Σ' be the surface obtained by perturbing Σ as presented in Section 5.1. By Lemma 5.3, $\text{Tri } S|_{\Sigma} = \text{Del } S|_{\Sigma'}$ is homeomorphic to Σ' and hence to Σ . \square

The lower bound in Lemma 5.1 on the distances from any new point inserted by `TOPOLOGY` to existing samples is instrumental to the analysis of the size of the final triangulation. However, the seeds inserted in step 1 of `SAMPLETOPOLOGY` can be arbitrarily close together (unlikely in practice though). Therefore, we delete the seeds one by one in step 3 of `SAMPLETOPOLOGY`. The GIP or TBP may become invalid after the deletion of a seed. Thanks to the next lemma, we can restore it by running `TOPOLOGY`.

Lemma 5.4 *Suppose that no Voronoi cell in $\text{Vor } S$ contains a connected component of Σ and no Voronoi facet in $\text{Vor } S$ intersects Σ in a cycle. For any $p \in S$, no connected component of Σ is contained in a Voronoi cell in $\text{Vor}(S \setminus \{p\})$.*

Proof. Assume to the contrary that a connected component M of Σ is contained in the Voronoi cell of a sample q in $\text{Vor}(S \setminus \{p\})$. Consider the insertion of p into $S \setminus \{p\}$ and the corresponding update of the Voronoi diagram. Let H be the bisecting plane of p and q . If H does not intersect M , the Voronoi cell of p or q in $\text{Vor } S$ would contain M , contradicting our assumption. If H intersects M , $H \cap M$ would contain a cycle in the facet between the Voronoi cells of p and q in $\text{Vor } S$, a contradiction again. \square

Lemma 5.4 enables us to invoke Corollary 5.1 after the deletion of one seed. This shows that $\text{Tri } S|_{\Sigma}$ will be homeomorphic to Σ after running `TOPOLOGY`.

Corollary 5.2 *Following each deletion of a seed in step 3 of `SAMPLETOPOLOGY`, the invocation of `TOPOLOGY` guarantees that the underlying space of $\text{Tri } S|_{\Sigma}$ is homeomorphic to Σ afterward.*

6 Meshing algorithm

Our meshing algorithm `DELMESH` first invokes `SAMPLETOPOLOGY` to capture the topology of Σ . Topological guarantee alone is not sufficient for many applications. In finite element methods, it is important that the surface triangles have bounded aspect ratio. Also, the output approximation should be smooth enough as it approximates a smooth surface. We introduce two procedures to address these issues in Section 6.1. Notice that these two procedures help capturing the geometry of Σ but cannot guarantee a non-trivial upper bound on the Hausdorff distances between input and output relative to the local feature sizes. It seems that such a guarantee would require computing the local feature sizes explicitly which we want to avoid. In Section 6.2, we give the complete description of `DELMESH` and its analysis.

6.1 Geometry sampling

Given a triangle t , define $\rho(t)$ to be the ratio of the circumradius of t to shortest side length of t . It is well-known that t has bounded aspect ratio if $\rho(t)$ is bounded from above by some constant. Following Chew [14], if there is a triangle t in $\text{Tri } S|_{\Sigma}$ with $\rho(t) > 1 + \lambda$, the procedure `QUALITY` below inserts into S the intersection point between Σ and the dual Voronoi edge of t .

QUALITY(sample set S)

1. While there is a triangle t in $\text{Tri } S|_{\Sigma}$ with $\rho(t) > 1 + \lambda$,
 - (a) Compute an intersection point x between Σ and the dual Voronoi edge of t . (Arbitrarily choose one if there are more than one.)
 - (b) Insert x into S and update $\text{Vor } S$.
2. Return S .

Assuming that $\text{Tri } S|_{\Sigma}$ is an orientable 2-manifold without boundary, we measure its smoothness using the dihedral angles at the edges. Specifically, for each edge e in $\text{Tri } S|_{\Sigma}$, we define the *roughness* of e , denoted by $g(e)$, to be π minus the dihedral angle at e . The procedure SMOOTH below samples a point from Σ if the roughness of some edge exceeds $2\beta(\lambda)$.

SMOOTH(sample set S)

1. If there is an edge pq in $\text{Tri } S|_{\Sigma}$ such that $g(pq) > 2\beta(\lambda)$
 - (a) Compute the intersections between Σ and the dual Voronoi edges of the triangles in $\text{Tri } S|_{\Sigma}$ incident to pq .
 - (b) Pick the furthest intersection point x from p .
 - (c) Insert x into S and update $\text{Vor } S$.
2. Return S .

QUALITY enforces that the angles of every triangle are no less than $\arcsin\left(\frac{1}{2+2\lambda}\right)$. SMOOTH enforces the dihedral angles are no less than $\pi - O(\lambda)$. Thus, we can improve the triangle shape and smoothness by decreasing λ . However, as explained in Theorem 6.1, the mesh size increases linearly in $\frac{1}{\lambda^2}$.

6.2 Finale

We give the pseudo-code of DELMESH below. DELMESH maintains the sample set S and $\text{Tri } S|_{\Sigma}$ throughout its execution. The final triangulation $\text{Tri } S|_{\Sigma}$ is the output surface mesh desired.

DELMESH(surface Σ)

1. Compute $S := \text{SAMPLETOPOLOGY}(\Sigma)$.
2. Compute $S := \text{QUALITY}(S)$. If QUALITY inserted some point(s) into S , compute $S := \text{TOPOLOGY}(S)$ and repeat step 2.
3. Compute $S := \text{SMOOTH}(S)$. If SMOOTH inserted a point into S , compute $S := \text{TOPOLOGY}(S)$ and go to step 2.
4. Output $\text{Tri } S|_{\Sigma}$.

Notice that after QUALITY or SMOOTH inserts new sample point(s), we call TOPOLOGY again because the new sample point(s) may disturb the topology of $\text{Tri } S|_{\Sigma}$. It is worthwhile to note that one does not need to search the entire $\text{Vor } S$ for possible topology violation. Instead, since the insertion of a new point changes $\text{Vor } S$ locally, a local search suffices.

We need the following technical result to analyze DELMESH.

Lemma 6.1 *Let x be a point whose distance to the nearest sample $p \in S$ is at least $\lambda f(p)$. For any point $q \in S$, (i) $\|q - x\| \geq \lambda f(x)/(1 + \lambda)$ and (ii) $\|q - x\| \geq \lambda f(q)/(1 + 2\lambda)$.*

Proof. By the Lipschitz condition, $f(x) \leq f(p) + \|p - x\| \leq (1 + \lambda) \|p - x\|/\lambda$. Since p is the nearest sample to x , for any $q \in S$, $\|q - x\| \geq \|p - x\| \geq \lambda f(x)/(1 + \lambda)$. By the Lipschitz condition again, $f(q) \leq f(x) + \|q - x\| \leq (1 + 2\lambda) \|q - x\|/\lambda$. \square

We are ready to prove the performance guarantees of DELMESH.

Theorem 6.1 *Let $\lambda \leq 0.06$ be a constant chosen a priori. Given a smooth closed surface Σ , DELMESH(Σ) outputs a mesh $\text{Tri}S|_\Sigma$ consisting of Delaunay triangles such that:*

- (i) *The underlying space of $\text{Tri}S|_\Sigma$ is homeomorphic to Σ .*
- (ii) *The radius-edge ratio of every triangle in $\text{Tri}S|_\Sigma$ is at most $1 + \lambda$ and the roughness of every edge of $\text{Tri}S|_\Sigma$ is at most $2\beta(\lambda)$.*
- (iii) *For any $\varepsilon < 1/5$, the size of S is within a factor $O(\frac{\varepsilon^2}{\lambda^2})$ of the size of an ε -sample of Σ .*

Proof. Take a non-seed point x inserted by DELMESH. We prove by induction that $\|p - x\| \geq \lambda f(p)$, where p is a nearest sample in S to x at the time of insertion.

If x is a non-seed point inserted by SAMPLETOPOLOGY, then $\|p - x\| \geq \lambda f(p)$ by Lemma 5.1.

Suppose that x is inserted by QUALITY. So x is an intersection point between Σ and the dual Voronoi edge of some triangle t in $\text{Tri}S|_\Sigma$. Let pq be the shortest edge of t . Since S does not contain any seed point at the end of SAMPLETOPOLOGY, p and q are non-seed points inserted some time in the past. Without loss of generality, assume that p was inserted after q . Then $\|p - q\| \geq \lambda f(p)/(1 + \lambda)$ by induction assumption and Lemma 6.1(i). Observe that p and the other vertices of t are the nearest samples to x when x is inserted. Since $\rho(t) > 1 + \lambda$, $\|p - x\| > (1 + \lambda) \|p - q\| > \lambda f(p)$.

Suppose that x is inserted by SMOOTH, triggered by an edge pq in $\text{Tri}S|_\Sigma$ whose roughness is greater than $2\beta(\lambda)$. Let pqr and pqs be the two triangles in $\text{Tri}S|_\Sigma$ incident to pq . So among the intersection points between Σ and the dual Voronoi edges of pqr and pqs , x is the furthest one from p . Observe that p is a nearest sample in S to x . Assume to the contrary that $\|p - x\| < \lambda f(p)$. This implies that the circumradii of pqr and pqs are less than $\lambda f(p)$. Let d and d' be the outward normals of pqr and pqs respectively. By Lemma 2.2, $\angle d, n_p \leq \beta(\lambda)$ and $\angle d', n_p \leq \beta(\lambda)$. Therefore, $\angle d, d' \leq 2\beta(\lambda)$ which means that the roughness of pq is at most $2\beta(\lambda)$, contradicting the insertion of x . Hence, $\|p - x\| \geq \lambda f(p)$.

In all, when x is inserted by DELMESH, the distance from x to its nearest sample p in S is at least $\lambda f(p)$. Therefore, DELMESH terminates by a packing argument.

By Lemma 5.2, Corollary 5.1, and Corollary 5.2, $\text{Tri}S|_\Sigma$ is homeomorphic to Σ and no Voronoi cell in $\text{Vor}S$ contains a connected component of Σ at the end of the first call to SAMPLETOPOLOGY. We claim that the subsequent insertion of new point(s) to repair the geometry preserves the property that no connected component of Σ is contained in any Voronoi cell. This can be seen as follows. Suppose that we add a new point p to repair the geometry after SAMPLETOPOLOGY. Any existing Voronoi cell can only shrink and the shrunk Voronoi cell cannot contain a component. If the new Voronoi cell V_p contains a component M , either p is the only vertex sampled from M or no point is sampled from M at all. But this is

impossible because SAMPLETOPOLOGY guarantees that $\text{Tri}S|_\Sigma$ is homeomorphic to Σ which means $\text{Tri}S|_\Sigma$ has a vertex from every connected component of Σ before we insert any point to repair the geometry. By our claim, one can argue inductively that no Voronoi cell in $\text{Vor}S$ contains a connected component of Σ throughout the execution of DELMESH. Then, (i) follows from Corollary 5.1. The correctness of (ii) follows from the termination of DELMESH.

For the rest of the proof, S denotes the final sample set obtained by DELMESH. Erickson [23] proved that for any $\varepsilon < 1/5$, the size of an ε -sampling of Σ is $\Omega(\int_\Sigma \frac{1}{\varepsilon^2 f(x)^2} dx)$. Therefore, to prove (iii), it suffices to show that the size of S is $O(\frac{1}{\lambda^2}) \cdot \int_\Sigma \frac{1}{f(x)^2} dx$.

Let q and r be two points in S . Irrespective of whether DELMESH inserted q or r first, Lemma 6.1(i) and (ii) imply that $\|q - r\| \geq \lambda f(q)/(1 + 2\lambda)$. Therefore, if we put a ball B_q centered at q with radius $\lambda f(q)/(2 + 4\lambda)$ for each $q \in S$, the balls have disjoint interior. It follows that

$$\int_\Sigma \frac{1}{f(x)^2} dx \geq \sum_{q \in S} \int_{B_q \cap \Sigma} \frac{1}{f(x)^2} dx.$$

For each point $x \in B_q \cap \Sigma$, the Lipschitz condition implies that $f(x) \leq f(q) + \|q - x\| \leq (2 + 5\lambda)f(q)/(2 + 4\lambda)$. Hence,

$$\int_\Sigma \frac{1}{f(x)^2} dx \geq \sum_{q \in S} \int_{B_q \cap \Sigma} \frac{4(1 + 2\lambda)^2}{(2 + 5\lambda)^2 f(q)^2} dx.$$

By Lemma 2.3, $\angle qx, n_q \geq \arccos(\lambda/(4 + 8\lambda))$. Let x' be the orthogonal projection of x onto the tangent plane of Σ at q . It follows that $\|q - x'\| \geq \|q - x\| \cdot \cos(\arcsin(\lambda/(4 + 8\lambda)))$, which is greater than $\|q - x\|/\sqrt{2}$ for $\lambda \leq \lambda_0 = 0.06$. By a result of Boissonnat and Cazals [5], $B_q \cap \Sigma$ is a topological disk as $\text{radius}(B_q) < f(q)$. Moreover, it can be verified using Lemmas 2.3 and 2.1 that $B_q \cap \Sigma$ is monotone in direction n_q . We conclude that the orthogonal projection of $B_q \cap \Sigma$ onto the tangent plane of Σ at q covers a disk centered at q with radius $\frac{\lambda}{2\sqrt{2}(1+2\lambda)}f(q)$. Hence, $\text{area}(B_q \cap \Sigma) \geq \frac{\pi\lambda^2}{8(1+2\lambda)^2} f(q)^2$. Therefore,

$$\int_\Sigma \frac{1}{f(x)^2} dx \geq \sum_{q \in S} \int_{B_q \cap \Sigma} \frac{\pi\lambda^2}{2(2 + 5\lambda)^2 \text{area}(B_q \cap \Sigma)} dx = \frac{\pi\lambda^2 |S|}{2(2 + 5\lambda)^2}.$$

□

7 Discussions

We presented a provable algorithm for sampling and meshing a smooth surface without boundary. Implicit surfaces can be meshed with this algorithm, which offers guarantees on the topology, triangular shape, smoothness, and size of the output triangulation. The mesh is Delaunay. It is worthwhile to note that we also obtain a Delaunay meshing of the volume bounded by the output surface mesh.

We implemented a simplified version of DELMESH using CGAL [39]. We did not implement the FACETCONTACT, FACETCYCLE, and SILHOUETTE subroutines. Figure 12 shows the results of this implementation for some simple smooth surfaces. Although the theory applies to smooth surfaces, we experimented with some triangulated surfaces obtained by a surface reconstruction

software called TIGHT COCONE [38]. Although these surfaces already have sample points, we disregarded all these sample points for our experiments and considered the piecewise linear surface as input. For each surface, DELMESH generated a new set of sample points and the corresponding restricted Delaunay triangulation. Figure 13 shows these triangulations.

These examples show that DELMESH can be used for remeshing triangulated surfaces while guaranteeing bounded aspect ratio. An open question remains if the method or its variant can be proved to mesh non-smooth surfaces with guarantees. This question is partially addressed by Dey, Li, and Ray [19] and Boissonnat and Oudot [9].

The critical point computations are the most costly computations in DELMESH. Can we avoid them and under what circumstances?




Figure 12: Meshing of a smooth sphere, torus, and a metaball each of which is input with an implicit equation.

Acknowledgments

We thank Günter Rote immensely for detecting an error in an earlier draft which led us to introduce the concept of silhouette. We also thank anonymous referees whose detailed and thoughtful comments were very helpful in improving the paper.

References

- [1] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.*, 22 (1999), 481–504.
- [2] N. Amenta, S. Choi, T. K. Dey and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *Internat. J. Comput. Geom. Applications*, 12 (2002), 125–141.
- [3] C. Bajaj. Surface fitting with implicit algebraic surface patches. *Topics in Surface Modeling*, ed. H. Hagen, SIAM Publications (1992).
- [4] J. Bloomenthal. Polygonization of implicit surfaces. *Proc. Comput. Aided Geom. Design* **5** (1988), 341–355.
- [5] J.-D. Boissonnat and F. Cazals. Natural neighbor coordinates of points on a surface. *Comput. Geom. Theory Appl.*, 19 (2001), 87–120.




Figure 13: The first column shows the surfaces to be sampled. The second column shows the triangulations after capturing the topology and deleting the seeds. The third column shows the results after improving the triangular shape and smoothing.

- [6] J.-D. Boissonnat, D. Cohen-Steiner, B. Mourrain, G. Rote, and G. Vegter. Meshing of surfaces. Chapter 5 in *Effective Computational Geometry for Curves and Surfaces*, eds. J.-D. Boissonnat, M. Teillaud, Springer Verlag, 2006.
- [7] J.-D. Boissonnat, D. Cohen-Steiner and G. Vegter. Isotopic implicit surface meshing. *Proc. 36th ACM Sympos. Theory Comput.*, 2004, 301–309.
- [8] J.-D. Boissonnat and S. Oudot. Provably good sampling and meshing of surfaces. *Graphical Models*, 67 (2005), 405–451.
- [9] J.-D. Boissonnat and S. Oudot. Provably good sampling and meshing of Lipschitz surfaces. *Proc. 22nd Ann. Sympos. Comput. Geom.* (2006), 337–346.
- [10] H.-L. Cheng, T. K. Dey, H. Edelsbrunner and J. Sullivan. Dynamic skin triangulation. *Discrete Comput. Geom.*, 25 (2001), 525–568.
- [11] S.-W. Cheng and T. K. Dey. Quality meshing with weighted Delaunay refinement. *SIAM J. Computing*, 33 (2003), 69–93.
- [12] S.-W. Cheng, T. K. Dey, E. A. Ramos and T. Ray. Quality meshing for polyhedra with small angles. *Intl. J. Comput. Geom. Appl.*, 15 (2005), 421–461.
- [13] S.-W. Cheng and S.-H. Poon. Three-Dimensional Delaunay Mesh Generation. *Discrete Comput. Geom.*, 36 (3006), 419–456.

- [14] L. P. Chew. Guaranteed-quality mesh generation for curved surfaces. *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, 1993, 274–280.
- [15] D. Cohen-Steiner, É. Colin de Verdière and M. Yvinec. Conforming Delaunay triangulations in 3D. *Proc. 18th Annu. Sympos. Comput. Geom.*, 2002, 199–208.
- [16] J.C. Cuilliere. An adaptive method for the automatic triangulation of 3D parametric surfaces. *Computer-Aided Design*, 30 (1998), 139–149.
- [17] T. K. Dey and W. Zhao. Approximating the medial axis from the Voronoi diagram with a convergence guarantee. *Algorithmica*, 38 (2004), 179–200.
- [18] T. K. Dey. *Curve and surface reconstruction : Algorithms with mathematical analysis*. Cambridge University Press, New York, 2006.
- [19] T.K. Dey, G. Li, and T. Ray. Polygonal surface meshing with Delaunay refinement. *Proc. 14th International Meshing Roundtable*, 2005, 343–351.
- [20] H. Edelsbrunner. Deformable smooth surface design. *Discrete Comput. Geom.*, 21 (1999), 87–115.
- [21] H. Edelsbrunner and E. Mücke. Simulation of simplicity: a technique to cope with degenerate cases in geometric algorithms. *ACM Trans. Graphics*, 9 (1990), 66–104.
- [22] H. Edelsbrunner and N. Shah. Triangulating topological spaces. *Internat. J. Comput. Geom. Appl.*, 7 (1997), 365–378.
- [23] J. Erickson. Nice point sets can have nasty Delaunay triangulations. *Discrete Comput. Geom.*, 30 (2003), 109–132.
- [24] M. Golubitsky and V. Guillemin. *Stable Mappings and Their Singularities*. Springer-Verlag, New York, 1973.
- [25] V. Guillemin and A. Pollack. *Differential Topology*. Prentice-Hall, 1974.
- [26] T. G. Kolda, R. M. Lewis and V. Torczon. Optimization by direct search: New perspectives on some classical and modern methods. *SIAM Review*, 45 (2003), 385–482.
- [27] T.S. Lau and S.H. Lo. Finite element mesh generation over analytical curved surfaces. *Computers and Structures*, 59 (1996), 301–309.
- [28] W.E. Lorensen and H.E. Cline. Marching cubes: a high resolution 3D surface construction algorithm. *Proc. SIGGRAPH*, 1987, 163–169.
- [29] B. Mourrain and J.-P. Témour. Isotopic meshing of a real algebraic surface. Technical report RR-5508, Inria, Sophia-Antipolis, 2005.
- [30] S. Plantinga and G. Vegter. Isotopic approximation of implicit curves and surfaces. *Symposium on Geometry Processing*, 2004.
- [31] J. R. Shewchuk. Tetrahedral mesh generation by Delaunay refinement. *Proc. 14th Annu. ACM Sympos. Comput. Geom.*, 1998, 86–95.
- [32] J. R. Shewchuk. What is a good linear element? interpolation, conditioning and quality measures. *Proc. 11th Internat. Meshing Roundtable*, 2002, 115–126.

- [33] J. Snyder. *Generative Modeling for Computer Graphics and CAD*. Academic press, 1992.
- [34] B. T. Stander and J. C. Hart. Guaranteeing the topology of an implicit surface polygonalization for interactive modeling. *Proc. SIGGRAPH*, 1997, 279–286.
- [35] J.R. Tristano, S.J. Owen, and S.A. Canann. Advancing front surface mesh generation in parametric space using a Riemannian surface definition. *Proc. 7th International Meshing Roundtable* 1998.
- [36] A. H. Wallace. *Differential topology : first steps*. W. A. Benjamin, Inc., New York, 1968.
- [37] A. P. Witkin and P. S. Heckbert. Using particles to sample and control implicit surfaces. *Proc. SIGGRAPH*, 1994, 269–278.
- [38] <http://www.cis.ohio-state.edu/~tamaldehy/cocone.html>
- [39] <http://www.cgal.org>