Delaunay Refinement for Piecewise Smooth Complexes *

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

Abstract

We present a Delaunay refinement algorithm for meshing a piecewise smooth complex in three dimensions. The algorithm protects edges with weighted points to avoid the difficulty posed by small angles between adjacent input elements. These weights are chosen to mimic the local feature size and to satisfy a Lipschitz-like property. A Delaunay refinement algorithm using the weighted Voronoi diagram is shown to terminate with the recovery of the topology of the input. Guaranteed bounds on the aspect ratios, normal variation and dihedral angles are also provided. To this end, we present new concepts and results including a new definition of local feature size and a proof for a generalized topological ball property.

1 Introduction

Delaunay refinement for meshing domains has been around for more than a decade now [15, 32, 33]. However, an algorithm that handles input as general as piecewise smooth complexes (PSC) is still lacking. There is a need in solid modeling to represent objects that are dimensionally or materially inhomogeneous [30]. For example, in a part consisting of different materials, the internal boundaries may meet to form a non-manifold. A boundary representation is essentially a graph in which the nodes represent the vertices, edges, and faces in the object, and the links represent the incidence relation [25]. Such a representation is a PSC when the curved segments called edges and the smooth surface patches called faces meet properly. The boundary of any solid with smooth faces and sharp junctions is a special case where the PSC is a 2-manifold.

As opposed to the polyhedral domains, topology recovery becomes a non-trivial issue for domains with curved elements since the output cannot conform to the input exactly. Some recent works address this issue by developing a Delaunay refinement strategy whose design and analysis are driven by topological property violations [8, 14, 18]. These algorithms work on the assumption that the input is a smooth surface or a surface that approximates a smooth surface.

To handle PSC one has to deal with non-smoothness. On top of it, possible small angles between the curved elements add further difficulty. The menace of small angles in Delaunay meshing is well known [17, 27, 34]. Although solutions for polyhedral domains have been obtained recently [11, 13, 29], the case for curved domains remains mostly open except for a recent work by Boissonnat and Oudot [7]. They showed that a class of surfaces called Lipschitz surfaces can be meshed with a Delaunay refinement strategy as long as the input angles are

*Research supported by Research Grant Council, Hong Kong, China (HKUST 6181/04E) and NSF, USA (CCF-0430735, CCF-0635008, and CCF-0635366).

†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

§Escuela de Matemáticas, Universidad Nacional de Colombia, Sede Medellín, Medellín, Colombia. Work performed in part while at the University of Illinois at Urbana-Champaign, Urbana, IL, USA. Email: earamosn@unalmed.edu.co

sufficiently large. Unfortunately, the derived angle bound allows only a very restricted class of non-smooth surfaces and the case of PSC which includes non-manifolds remains open. Recently, Rineau and Yvinec [31] proposed a method to mesh volumes bounded by piecewise smooth surfaces under the constraint that the tangents to the surface subtend large input angles at non-smooth regions.

In this paper we present an algorithm to mesh PSC with Delaunay refinement without any constraint on input angles. One notable aspect of our algorithm is that, unlike previous approaches [7, 18], it respects the input "non-smooth features" as it meshes the input curves and vertices with 1-complexes.

One difficulty in meshing PSC stems from the fact that the usual local feature size definitions fail to provide an upper bound on the number of repeated insertions triggered by topological ball property violations. We overcome this problem by a new definition of local feature size which does not vanish on the domain. It is reminiscent of the effort in [10] where a non-vanishing local feature size is defined for compact sets in the context of surface reconstruction. Here we combine the two classical definitions, one based on the medial axis [2] and the other based on the adjacency of elements [11, 32]. The non-smooth regions are then protected by balls whose sizes and placements are guided by a simulation of a Lipschitz property. The Delaunay refinement is run on the weighted Voronoi diagram where the protecting balls are treated as weighted points. The protected curves in PSC always appear as a union of edges of the restricted weighted Delaunay triangulation. This helps the triangulations, restricted to the individual smooth surface patches, match seamlessly. It also preserves non-smooth features in the output, a concern that is important in various applications. An extended topological ball property [20] for general topological spaces is proved to claim homeomorphism between the input and output. It is noteworthy that our weighted point approach never inserts a point inside the protecting balls. This eliminates the check for point insertions inside the protecting balls as well as the exception handler for such events.

1.1 Notation and domain

Throughout this paper, we assume a generic intersection property that a k-manifold $\sigma \subset \mathbb{R}^3$, $0 \leqslant k \leqslant 3$, and a j-manifold $\sigma' \subset \mathbb{R}^3$, $0 \leqslant j \leqslant 3$, intersect (if at all) in a $(k+j-3)$-manifold if $\sigma \not\subset \sigma'$ and $\sigma' \not\subset \sigma$. We will use both geometric and topological versions of closed balls. A *geometric* closed ball centered at point $x \in \mathbb{R}^3$ with radius $r > 0$, is denoted as $B(x,r)$. We use $\operatorname{int} \mathbb{X}$ and $\operatorname{bd} \mathbb{X}$ to denote the interior and boundary of a topological space $\mathbb{X}$, respectively. Given a complex $\mathbb{A}$, we use $|\mathbb{A}|$ to denote its underlying space.

The domain $\mathcal{D}$ is a PSC where each element is a compact subset of a C^3-smooth k-manifold, $0 \leqslant k \leqslant 2$. Each element is closed and hence contains its boundaries. We use $\mathcal{D}_k$ to denote the subset of all k-dimensional elements. $\mathcal{D}_0$ is a set of *vertices*; $\mathcal{D}_1$ is a set of curves called *1-faces*; $\mathcal{D}_2$ is a set of surface patches called *2-faces*. For $1 \leqslant k \leqslant 2$, we use $\mathcal{D}_{\leqslant k}$ to denote $\bigcup_{j=0}^{k} \mathcal{D}_j$. Two elements are *adjacent* if their intersection is non-empty. Two elements σ and τ are *incident* if $\tau \subseteq \operatorname{bd} \sigma$.

The domain $\mathcal{D}$ satisfies the usual requirements for being a complex: (i) interiors of the elements are pairwise disjoint; (ii) for any $\sigma \in \mathcal{D}$, $\operatorname{bd} \sigma$ is a union of elements in $\mathcal{D}$; (iii) for any $\sigma, \tau \in \mathcal{D}$, either $\sigma \cap \tau = \emptyset$ or $\sigma \cap \tau$ is a union of elements in $\mathcal{D}$.

For each 2-face σ, we assume that σ is part of a smooth compact surface Σ without boundary whose equation is given, and that σ is represented using the equation of Σ and some additional inequalities so that one can test whether a given point lies on σ. We also assume that $\operatorname{bd} \sigma$ consists of disjoint simple closed curves (no boundary points). For each 1-face σ, we assume that σ is an open curve (has boundary points) contained in the intersection of two surfaces Σ_1 and Σ_2 whose equations are also given, and that σ is represented using the equations of Σ_1 and Σ_2 as well as some additional inequalities so that one can test whether a given point lies on σ. We introduce the notation $mani(\sigma)$: if σ is a 1-face, $\text{mani}(\sigma)$ denotes $\Sigma_1 \cap \Sigma_2$; if σ is a 2-face, $\text{mani}(\sigma)$ denotes Σ.

For any 2-face σ and for any point $x \in \text{mani}(\sigma)$, we use $n_\sigma(x)$ to denote the unit outward normal to $\text{mani}(\sigma)$ at x. For any 1-face σ and for any point $x \in \text{mani}(\sigma)$, $n_\sigma(x)$ denotes the plane orthogonal to $\text{mani}(\sigma)$ at x.

Given two non-zero vectors $\vec{d_1}$ and $\vec{d_2}$, let $\angle \vec{d_1}, \vec{d_2}$ denote the angle between them (which lies in the range $[0, \pi]$). Given a 1- or 2-dimensional linear subset L in an Euclidean space and the vector $\vec{d}$, let $\angle \vec{d}, L$ denote the non-obtuse angle between $\vec{d}$ and the support line/plane of L. Similarly, given L_1 and L_2 where L_i is a 1- or 2-dimensional linear subset of an Euclidean space, let $\angle L_1, L_2$ denote the non-obtuse angle between the support line/plane of L_1 and the support line/plane of L_2.

Notice that our definition of PSC disallows certain kinds of surfaces such as cones or non-orientable surfaces since they cannot be extended to a smooth surface without boundary. Although we have not included three dimensional faces in the PSC definition, our algorithm can be extended to handle such a PSC using an approach similar to [28, 31]. We avoid them for saving further technicalities and analysis.

Various data structures have been proposed in the computer-aided design literature in order to capture the combinatorial structure of non-manifold solid models; for example, see [9, 21, 25]. Any of these data structures can be used to represent our input PSC with each element tagged with the corresponding equations. It is popular for surfaces to be represented parametrically by rational functions in computer aided design (e.g., NURBS). Given such a representation, we can first implicitize the surfaces using the algorithm by Manocha and Canny [24] and then run our meshing algorithm.

1.2 Overview

Our strategy is to run Delaunay refinement as long as the topology of the domain is not recovered. Such a strategy was used to mesh smooth surfaces [14] where each insertion was triggered by a violation of topological ball property [20] (TBP). First of all, we cannot use the TBP for manifolds here since we are dealing with PSC. It turns out that the extended topological ball property which can ensure topology recovery for PSC [20] requires TBP for each element in $\mathcal{D}$. One may drive the refinement with the violation of the extended TBP, but the possible presence of small angles in the input causes difficulty. As a remedy, we protect the vertices and 1-faces with some geometric balls so that no point is inserted in these protecting balls. Such a protection strategy made the Delaunay refinement possible for polyhedral domains with small angles [11, 13, 29]. However, placing the protecting balls of appropriate size seems to be far more difficult in the case of PSC. Once the balls are computed, we treat them as weighted points which brings the weighted Voronoi and weighted Delaunay diagrams into picture.

For a weighted point set $S \subset \mathbb{R}^3$, let $\text{Vor}\, S$ and $\text{Del}\, S$ denote the weighted Voronoi diagram and the weighted Delaunay triangulation of S, respectively. $\text{Vor}\, S$ and $\text{Del}\, S$ are cell complexes, where each k-face in $\text{Vor}\, S$ is a k-polytope and each k-face in $\text{Del}\, S$ is a k-simplex. Each k-face in $\text{Vor}\, S$ is dual to a $(3-k)$-face in $\text{Del}\, S$ and vice versa. We say that S has the TBP for an element in $\mathcal{D}_i$ if this element intersects any k-face in $\text{Vor}\, S$ in either an empty set or a closed topological ball of dimension $(i + k - 3)$.

If an element in the domain does not satisfy the TBP, our algorithm detects it and computes a point far away from all other existing points. At termination, which is guaranteed by a standard packing argument, S has the TBP for each element in $\mathcal{D}$.

The entire analysis and the crucial step of ball protection needs a feature size definition which is 1-Lipschitz and non-zero everywhere. We achieve it by a non-trivial combination of the local feature sizes defined for polyhedral domains and smooth surfaces.

1.3 Background results

We state a few technical results whose proofs are deferred to Appendix A.

Lemma 1.1 *Let E be the set of elements in $\mathcal{D}$ containing a point x. Let $B(x,r)$ be a ball such that for any $\sigma \in E$, $B(x,r) \cap \mathrm{mani}(\sigma)$ is a closed ball of dimension $\dim(\sigma)$. Then, for any $\sigma \in E$,*

(i) *if σ is a 1-face, $B(x,r) \cap \sigma$ is an open curve (topological interval);*

(ii) *if σ is a 2-face and every 1-face of σ intersecting $B(x,r)$ contains x, then $B(x,r) \cap \sigma$ is a topological disk.*

Lemma 1.2 *Let σ be a 2-face. Let $B(x,r)$ be a ball for some point $x \in \mathrm{mani}(\sigma)$ such that for any point $z \in B(x,r) \cap \mathrm{mani}(\sigma)$, $\angle n_\sigma(x), n_\sigma(z) < \theta$ for some $\theta \leqslant \pi/8$. Let $B \subset B(x,r)$ be a ball centered at a point $y \in B(x,r) \cap \mathrm{mani}(\sigma)$. For any plane H containing y such that $\angle n_\sigma(x), H < \theta$, $H \cap B \cap \mathrm{mani}(\sigma)$ is an open curve.*

Lemma 1.3 *Let σ be a 1- or 2-face. Let $B(x,r)$ be a ball for some point $x \in \mathrm{mani}(\sigma)$ such that for any point $z \in B(x,r) \cap \mathrm{mani}(\sigma)$, $\angle n_\sigma(x), n_\sigma(z) < \theta$ for some $\theta < \pi/2$.*

(i) *If σ is a 1-face and $B(x,r) \cap \mathrm{mani}(\sigma)$ is connected, then for any points $y, z \in B(x,r) \cap \mathrm{mani}(\sigma)$, $\angle n_\sigma(x), yz > \pi/2 - \theta$.*

(ii) *If σ is a 2-face, $\theta \leqslant \pi/8$ and $B(x, \frac{1}{3}r) \cap \mathrm{mani}(\sigma)$ is connected, then for any points $y, z \in B(x, \frac{1}{3}r) \cap \mathrm{mani}(\sigma)$, $\angle n_\sigma(x), yz > \pi/2 - \theta$.*

2 Feature size

We propose a definition of local feature size for PSC and establish several geometrical and topological properties. The local feature size and the properties will be useful in analyzing our meshing algorithm. As in the case of meshing smooth closed surfaces, it is hard to compute the local feature size at points in a PSC. Therefore, we propose a computable alternative for points on 1-faces. Our algorithm uses this computable alternative to construct balls to protect the vertices and 1-faces before running Delaunay refinement.

2.1 Local feature size

For any point x on an element $\sigma \in \mathcal{D}$, define $m_\sigma(x)$ to be the distance between x and the medial axis of $\mathrm{mani}(\sigma)$. Take $m_\sigma(x) = \infty$ if σ is a vertex. It is well-known that m_σ is 1-Lipschitz, i.e., for any points $x, y \in \sigma$, $m_\sigma(x) \leqslant m_\sigma(y) + \|x - y\|$. Define $m(x) = \min\{m_\sigma(x) : x \in \sigma\}$. Notice that m may not be continuous. For example, for a 2-face σ, only the medial axis of mani(σ) contributes to $m(x)$ when $x \in \operatorname{int}\sigma$, but more than one medial axis may contribute to $m(x)$ when $x \in \operatorname{bd}\sigma$. We also employ the *local gap size* $g(x)$ used in [11] to deal with small input angles: for any point $x \in |\mathcal{D}|$, $g(x)$ is the minimum $r > 0$ such that $B(x, r)$ intersects two elements of $\mathcal{D}_{\leqslant 2}$, one of which does not contain x.

Definition 2.1 *For any point $x \in |\mathcal{D}|$, define $f(x) = \min\{m(x), g(x)\}$.*

Lemma 2.1 *For any 1- or 2-face σ, the following properties hold.*

(i) *For any point $x \in \sigma$,*

  (a) *if B is a ball tangent to* mani(σ) *at x and* radius(B) $< f(x)$*, then $\sigma \cap \operatorname{int} B = \emptyset$;*

  (b) *if $r < f(x)$, then $B(x, r) \cap \operatorname{mani}(\sigma)$ and $B(x, r) \cap \sigma$ are closed balls of dimension* $\dim(\sigma)$*.*

(ii) *For any points $x, y \in \operatorname{mani}(\sigma)$, if $\|x - y\| \leqslant \mu f(x)$ for some $\mu \leqslant 1/3$, then $\angle n_\sigma(x), n_\sigma(y) \leqslant \mu/(1 - \mu)$.*

Proof. Because $f(x) \leqslant m_\sigma(x)$, (i)(a) follows from standard results in the theory of ε-sampling [19]. Since $r < f(x) \leqslant m(x)$, it is known that $B(x, r) \cap \operatorname{mani}(\tau)$ is a closed ball of dimension $\dim(\tau)$ for all faces τ containing x [19]. Coupled with the fact that $r < f(x) \leqslant g(x)$, Lemma 1.1 can be invoked to imply (i)(b). Since $f(x) \leqslant m_\sigma(x)$, (ii) follows from the standard normal variation result; see [4]. ◻

The function f is not continuous because m and g may not be continuous. The value of $m(x)$ may change discontinuously when a point x moves from a 2-face to an adjacent 1-face because the distance to the medial axis of the 1-face suddenly enters into the minimization in the definition of $m(x)$. If a point x moves from a 1- or 2-face to an adjacent face of lower dimension, the value of $g(x)$ approaches zero by definition. This makes g discontinuous because $g(x)$ is strictly positive at the 0-faces. Still, f, g and m are 1-Lipschitz in a restricted sense.

Lemma 2.2 *For any $\sigma \in \mathcal{D}$, the functions f, g and m are 1-Lipschitz over* $\operatorname{int}\sigma$*.*

Proof. The lemma is trivial for the vertices of $\mathcal{D}$. Let σ be a 1- or 2-face. Take any two points $x, y \in \operatorname{int}\sigma$. Since $x, y \in \operatorname{int}\sigma$, they belong to the same set of elements of $\mathcal{D}$ which we denote as E. So $m(y) + \|x - y\| = \min_{\tau \in E}\{m_\tau(y) + \|x - y\|\} \geqslant \min_{\tau \in E} m_\tau(x) = m(x)$, proving that m is 1-Lipschitz over $\operatorname{int}\sigma$. By definition, $B(y, g(y))$ intersects another element τ not containing y. Since $B(y, g(y))$ is contained inside $B(x, r)$, where $r = \|x - y\| + g(y)$, $B(x, r)$ intersects both σ and τ. As $x \in \operatorname{int}\sigma$, $x \notin \tau$. It follows that $g(x) \leqslant r = g(y) + \|x - y\|$, proving that g is 1-Lipschitz over $\operatorname{int}\sigma$. Finally, $f(y) + \|x - y\| = \min\{m(y) + \|x - y\|, g(y) + \|x - y\|\} \geqslant \min\{m(x), g(x)\} = f(x)$. So f is 1-Lipschitz over $\operatorname{int}\sigma$. ◻

The function f is close to being a local feature size function for $\mathcal{D}$. It blends the definition of local feature sizes of closed surfaces, which are used in meshing smooth surfaces [8, 14], and the definition of local gap sizes of polyhedral domains, which are used in polyhedral meshing [11, 17] in the presence of small angles. The blending enables us to take care of both types of geometries in $\mathcal{D}$: smooth surfaces meeting sharply at vertices and curves. The problem is that f is only 1-Lipschitz within the interior of an element in $\mathcal{D}$. As a remedy, we propagate the value of f from $\operatorname{bd}|\mathcal{D}_i|$ to $\operatorname{int}|\mathcal{D}_i|$ in order to define our local feature size function. We introduce a parameter $\delta \in (0, 1)$ to control the extent of the propagation. The local feature size function $\operatorname{lfs}_\delta : |\mathcal{D}| \to \mathbb{R}$ is defined as follows.

Definition 2.2 *For any* $\delta \in (0,1]$ *and for any point* $x \in |\mathcal{D}|$,

$$\mathrm{lfs}_\delta(x) = \begin{cases} \delta f(x) & \text{if } x \in |\mathcal{D}_0|, \\ \max\{\delta f(x), \max_{y \in \mathrm{bd}\,\sigma}\{\delta f(y) - \|x-y\|\}\} & \text{if } x \in \mathrm{int}\,\sigma \text{ for some } \sigma \in \mathcal{D}_1 \cup \mathcal{D}_2. \end{cases}$$

Roughly speaking, we design $\mathrm{lfs}_\delta(x)$ so that if x lies on a face σ and is sufficiently far away from $\mathrm{bd}\,\sigma$, $\mathrm{lfs}_\delta(x)$ is equal to $\delta f(x)$. Indeed, if $\|x-y\| < \delta f(y)$ for some point $y \in \mathrm{bd}\,\sigma$, then $\delta f(y) - \|x-y\| > 0$ and so it is possible for $\mathrm{lfs}_\delta(x)$ to be equal to $\max_{y\in \mathrm{bd}\,\sigma}\{\delta f(y) - \|x-y\|\}$. On the other hand, if $\|x-y\| \geqslant \delta f(y)$ for any point $y \in \mathrm{bd}\,\sigma$, $\max_{y\in \mathrm{bd}\,\sigma}\{\delta f(y) - \|x-y\|\}$ is non-positive and so $\mathrm{lfs}_\delta(x) = \delta f(x)$. The maximization $\max_{y\in \mathrm{bd}\,\sigma}\{\delta f(y) - \|x-y\|\}\}$ is employed in the definition to help making lfs_δ 1-Lipschitz. Similar maximizations and minimizations have been used before in obtaining Lipschitz functions [1, 22, 36]. The next result shows that lfs_δ is indeed 1-Lipschitz.

Lemma 2.3 lfs_δ *is 1-Lipschitz.*

Proof. Take any points $x, y \in |\mathcal{D}|$. Let σ_x and σ_y be the elements of $\mathcal{D}$ such that $x \in \mathrm{int}\,\sigma_x$ and $y \in \mathrm{int}\,\sigma_y$. (When x is a vertex, $x = \sigma_x$. The same is true for y.)

Suppose that $y \in \sigma_x$. So $\sigma_y = \sigma_x$ or $\sigma_y \subset \mathrm{bd}\,\sigma_x$. We claim that $\mathrm{lfs}_\delta(x) \geqslant \delta f(y) - \|x-y\|$: if $y \in \mathrm{bd}\,\sigma_x$, $\mathrm{lfs}_\delta(x) \geqslant \delta f(y) - \|x-y\|$ by definition; if $y \in \mathrm{int}\,\sigma_x$, $\mathrm{lfs}_\delta(x) \geqslant \delta f(x) > \delta f(y) - \|x-y\|$ by Lemma 2.2. By our claim, if $\mathrm{lfs}_\delta(y) = \delta f(y)$, we have $\mathrm{lfs}_\delta(x) \geqslant \mathrm{lfs}_\delta(y) - \|x-y\|$. Otherwise, $\mathrm{lfs}_\delta(y) = \delta f(z) - \|y-z\|$ for some point $z \in \mathrm{bd}\,\sigma_y \subseteq \mathrm{bd}\,\sigma_x$. By definition, $\mathrm{lfs}_\delta(x) \geqslant \delta f(z) - \|x-z\| \geqslant \delta f(z) - \|y-z\| - \|x-y\| = \mathrm{lfs}_\delta(y) - \|x-y\|$.

Suppose that $y \notin \sigma_x$. By definition, $\mathrm{lfs}_\delta(y) = \delta f(z) - \|y-z\|$ for some point $z \in \sigma_y$. If $z \notin \sigma_x$, $\|x-z\| \geqslant g(z)$ and so $\mathrm{lfs}_\delta(y) \leqslant g(z) - \|y-z\| \leqslant \|x-z\| - \|y-z\| \leqslant \|x-y\| < \mathrm{lfs}_\delta(x) + \|x-y\|$. If $z \in \sigma_x$, we have proved in the previous paragraph that $\mathrm{lfs}_\delta(x) \geqslant \mathrm{lfs}_\delta(z) - \|x-z\|$. So $\mathrm{lfs}_\delta(y) = \delta f(z) - \|y-z\| \leqslant \mathrm{lfs}_\delta(z) - \|y-z\| \leqslant \mathrm{lfs}_\delta(x) + \|x-z\| - \|y-z\| \leqslant \mathrm{lfs}_\delta(x) + \|x-y\|$. ◻

The next result shows that the properties in Lemma 2.1 hold for lfs_δ as well.

Lemma 2.4 *Lemma 2.1 holds with* f *replaced by* lfs_δ.

Proof. Take a point x on a 1- or 2-face σ. By definition, $\mathrm{lfs}_\delta(x) = \delta f(y) - \|x-y\|$ for some $y \in \sigma$. So $\mathrm{lfs}_\delta(x) \leqslant \delta m_\sigma(y) - \|x-y\| \leqslant m_\sigma(x)$. Then, standard results in ε-sampling imply that Lemma 2.1 (i)(a) and (ii) hold for lfs_δ. Take any $r < \mathrm{lfs}_\delta(x)$. Since $\mathrm{lfs}_\delta(x) \leqslant m_\sigma(x)$, then $B(x,r) \cap \mathrm{mani}(\sigma)$ is a closed ball of dimension $\dim(\sigma)$ [6, 19]. It remains to show that $B(x,r) \cap \sigma$ is also a closed ball of dimension $\dim(\sigma)$.

If σ is a 1-face, the endpoints of σ split $B(x,r) \cap \mathrm{mani}(\sigma)$ into at most three subcurves, one of which is $B(x,r) \cap \sigma$.

Suppose that σ is a 2-face. If $\mathrm{int}\,B(x,r)$ avoids $\mathrm{bd}\,\sigma$, $B(x,r) \cap \sigma = B(x,r) \cap \mathrm{mani}(\sigma)$ and we are done. Suppose that $\mathrm{int}\,B(x,r)$ intersects $\mathrm{bd}\,\sigma$. Then, $\mathrm{lfs}_\delta(x) = \delta f(y) - \|x-y\|$ for some point $y \in \mathrm{bd}\,\sigma$ because if $x \in \mathrm{int}\,\sigma$ and $\mathrm{lfs}_\delta(x) = \delta f(x)$, $B(x,r)$ would avoid $\mathrm{bd}\,\sigma$ as $f(x) \leqslant g(x)$. So $B(x,r) \subset B(y, \delta f(y))$. Let E be the set of 1-faces in $\mathrm{bd}\,\sigma$ containing y. We have $|E| = 1$ when y is in the interior of a 1-face and $|E| = 2$ when y is the common endpoint of two 1-faces in $\mathrm{bd}\,\sigma$. Since $f(y) \leqslant g(y)$ and $f(y) \leqslant m(y)$, we know that:

(1) $B(y, \delta f(y))$ avoids any 1-face in $\mathrm{bd}\,\sigma$ that does not belong to E; so does $B(x,r)$.

(2) $B(y, \delta f(y))$ avoids the medial axis of $\mathrm{mani}(\tau)$ for each $\tau \in E$; so does $B(x,r)$.

By (2), it is known that $B(x,r) \cap \mathrm{mani}(\tau)$ is an open curve for each $\tau \in E$.

If E contains only one 1-face τ, then $y \in \mathrm{int}\,\tau$. Since $f(y) \leqslant g(y)$, $B(y, \delta f(y))$ avoids the endpoints of τ. So does $B(x,r)$. Thus, $B(x,r) \cap \tau = B(x,r) \cap \mathrm{mani}(\tau)$, i.e., an open curve, and it splits $B(x,r) \cap \mathrm{mani}(\sigma)$ into two topological disks, one of which is $B(x,r) \cap \sigma$.

If E contains two 1-faces τ_1 and τ_2, y is a common endpoint of τ_1 and τ_2. Again, both $B(y, \delta f(y))$ and $B(x,r)$ avoid the other endpoints of τ_1 and τ_2. If $y \in B(x,r)$, $B(x,r) \cap \tau_1$ concatenates with $B(x,r) \cap \tau_2$ to form an open curve. This open curve splits $B(x,r) \cap \mathrm{mani}(\sigma)$ into two topological disks, one of which is $B(x,r) \cap \sigma$. If $y \notin B(x,r)$, $B(x,r) \cap \tau_i = B(x,r) \cap \mathrm{mani}(\tau_i)$ for $i \in \{1,2\}$. These two open curves are disjoint and they split $B(x,r) \cap \mathrm{mani}(\sigma)$ into three topological disks, one of which is $B(x,r) \cap \sigma$. ◻

Notice that lfs_δ is actually δ-Lipschitz at points that are away from the element boundaries in $\mathcal{D}$. However, when a point x is close to some element boundary, $\mathrm{lfs}_\delta(x)$ becomes 1-Lipschitz by construction. The parameterized propagation of the value of f from boundaries matches well with our algorithmic strategy to protect the vertices and 1-faces with balls. We will see that δ is related to the radii of these protecting balls.

2.2 A computable alternative

It would be ideal if we can compute $\mathrm{lfs}_\delta(x)$ for any point $x \in |\mathcal{D}|$. Doing so requires computing $m(x)$ which in turn requires computing the medial axis of smooth surfaces and curves. This is a challenging problem by itself. We face a greater problem because of the maximization over $\mathrm{bd}\,|\mathcal{D}_i|$ used in defining lfs_δ.

We propose to compute a function $\mathrm{f}_\omega : |\mathcal{D}_{\leqslant 1}| \to \mathbb{R}$ for a parameter $\omega \in (0,1)$. The value $\mathrm{f}_\omega(x)$ enjoys several nice features of $\mathrm{lfs}_\delta(x)$ and it is good enough for our purpose. For example, we will see that if $\delta \leqslant \omega$, $\mathrm{lfs}_\delta(x) \leqslant \mathrm{f}_\omega(x)$ for the points $x \in |\mathcal{D}_{\leqslant 1}|$ that we are interested in. This makes f_ω appropriate for defining the radii of protecting balls covering $|\mathcal{D}_{\leqslant 1}|$, i.e., points outside the protecting balls are at distances at least local feature sizes away from $|\mathcal{D}_{\leqslant 1}|$.

Let x be a point in some 1-face. First, for any 1- or 2-face σ containing x, define $\sigma_{x,\omega} = \{y \in \mathrm{mani}(\sigma) : \cos(\angle n_\sigma(x), n_\sigma(y)) = \cos\omega\}$. The distance between x and $\sigma_{x,\omega}$ is the ω-deviation radius of x with respect to σ. We define $d_\omega(x)$ to be the minimum of the ω-deviation radii of x over all faces containing x. Second, we define $b(x)$ to be the largest value such that for any $r < b(x)$, $B(x,r) \cap \mathrm{mani}(\sigma)$ is a closed ball for any 1- or 2-face σ containing x.

Definition 2.3 *For any point $x \in |\mathcal{D}_{\leqslant 1}|$, define $\mathrm{f}_\omega(x) = \min\{d_\omega(x), g(x), b(x)\}$.*

The function f_ω satisfies the following properties.

Lemma 2.5 *Let $\omega \in (0, \frac{\pi}{8})$.*

(i) *For any point $x \in |\mathcal{D}_1|$, $\mathrm{f}_\omega(x) \in [\,\frac{\omega}{2} f(x)\,,\, g(x)\,]$.*

(ii) *For any 1- or 2-face σ,*

   (a) *for any points $x, y \in \mathrm{mani}(\sigma)$ such that $\|x-y\| < \mathrm{f}_\omega(x)$, $\angle n_\sigma(x), n_\sigma(y) < \omega$;*

   (b) *for any point $x \in \sigma$ and for any $r < \mathrm{f}_\omega(x)$, $B(x,r) \cap \mathrm{mani}(\sigma)$ and $B(x,r) \cap \sigma$ are closed balls of dimension $\dim(\sigma)$.*

(iii) *For any 1-face σ and for any points $x, y, z \in \mathrm{mani}(\sigma)$ such that $y, z \in B(x,r)$ for some $r < \mathrm{f}_\omega(x)$, one has $\angle n_\sigma(x), yz > \pi/2 - \omega$.*

(iv) *For any 2-face* σ *and for any points* $x, y, z \in \mathrm{mani}(\sigma)$ *such that* $y, z \in B(x, r)$ *for some* $r < \frac{1}{3}\mathrm{f}_\omega(x)$, *one has* $\angle n_\sigma(x), yz > \frac{\pi}{2} - \omega$.

Proof. By the standard normal variation result, for any 1- or 2-face σ and for any points $x, y \in \mathrm{mani}(\sigma)$, if $\|x - y\| \leqslant \frac{\omega}{2} m_\sigma(x)$, $\angle n_\sigma(x), n_\sigma(y) \leqslant \frac{\omega/2}{1-\omega/2}$. By the definition of $d_\omega(x)$, $\angle n_\tau(x), n_\tau(z) = \omega$ for some face τ containing x and some point $z \in \mathrm{mani}(\tau)$ such that $\|x - z\| = d_\omega(x)$. Because $\omega < 1$, $\frac{\omega/2}{1-\omega/2} < \omega$, which implies that $d_\omega(x) > \frac{\omega}{2} m_\tau(x)$. Therefore, $d_\omega(x) > \frac{\omega}{2} m(x) \geqslant \frac{\omega}{2} f(x)$. By definition, $g(x) \geqslant f(x)$. It follows from Lemma 2.1(i)(b) that $b(x) \geqslant f(x)$. Thus, $g(x) \geqslant \min\{d_\omega(x), g(x), b(x)\} = \mathrm{f}_\omega(x) \geqslant \frac{\omega}{2} f(x)$. This proves (i). Since $\mathrm{f}_\omega(x) \leqslant d_\omega(x)$, (ii)(a) is enforced by construction. Since $r < \mathrm{f}_\omega(x) \leqslant b(x)$, for any 1- or 2-face σ containing x, $B(x, r) \cap \mathrm{mani}(\sigma)$ is enforced to be a closed ball. Then, since $r < \mathrm{f}_\omega(x) \leqslant \omega g(x)$, $B(x, r) \cap \sigma$ is also a closed ball by Lemma 1.1. This proves (ii)(b). Since $\omega < \pi/8$ by assumption, the correctness of (iii) and (iv) follows from (ii) and Lemma 1.3. ◻

By Definition 2.3, $\mathrm{f}_\omega(x)$ is obtained by computing $d_\omega(x)$, $g(x)$, and $b(x)$. The computations of $g(x)$ and $b(x)$ are easily translated into solving systems of equations. So is the computation of $d_\omega(x)$ provided that for any 2-face σ containing x, the set $\sigma_{x,\omega} = \{y \in \mathrm{mani}(\sigma) : \cos(\angle n_\sigma(x), n_\sigma(y)) = \cos\omega\}$ is a collection of disjoint smooth closed curves. When $\sigma_{x,\omega}$ is not a collection of disjoint smooth closed curves, $\cos\omega$ is a *critical value* of the function $y \mapsto \cos(\angle n_\sigma(x), n_\sigma(y))$ by well known results in Morse theory [26]. We call ω *critical* if $\cos\omega$ is a critical value. We first detect whether ω is critical by testing the solvability of a system of equations. If ω is critical, we perturb ω to a non-critical value $\bar{\omega} < \omega$ and compute $\mathrm{f}_{\bar{\omega}}(x)$ instead. Moreover, $\bar{\omega}$ can be made as close to ω as one wishes. We call $\bar{\omega}$ a *non-critical value for* x. The details of the perturbation of ω and the computations of $d_\omega(x)$, $g(x)$ and $b(x)$ are given in Appendix B.

3 Protecting vertices and curves

Our algorithm starts with constructing protecting balls centered at judiciously chosen locations in $|\mathcal{D}_{\leqslant 1}|$. The radii of the protecting balls are controlled by f_ω (We may need to perturb ω if it happens to be critical). The value of f_ω may fluctuate greatly as it is not 1-Lipschitz. So we enforce a Lipschitz-like property on the fly as we construct the balls. In the end each protecting ball is turned into an equivalent weighted point.

We choose $\lambda \leqslant 0.01$ and $\omega \leqslant 0.01$. First, we put a protecting ball B_u centered at each vertex $u \in \mathcal{D}_0$ with $\mathrm{radius}(B_u) = \lambda \mathrm{f}_{\omega_u}(u)$, where ω_u is a non-critical value for u in the range $[0.99\omega, \omega]$. Clearly, $\mathrm{f}_{\omega_u}(u) \leqslant \mathrm{f}_\omega(u)$. Let σ be a 1-face. Let u and v be the endpoints of σ. We will incrementally put a certain number of points in $\mathrm{int}\,\sigma$ from u to v. For notational convenience, we denote this number by m although m is unknown at the beginning.

Define $x_1 = u$ and $r_1 = \lambda \mathrm{f}_{\omega_1}(u)$, where $\omega_1 = \omega_u$. We compute the intersection points $x_2 = \sigma \cap \mathrm{bd}\, B_u$ and $x_0 = \sigma \cap \mathrm{bd}\, B_v$. Define $r_2 = \lambda \mathrm{f}_{\omega_2}(x_2)$ and $r_0 = \lambda \mathrm{f}_{\omega_0}(x_0)$, where ω_0 and ω_2 are some non-critical values for x_0 and x_2, respectively, in the range $[0.99\omega, \omega]$. The protecting ball at x_2 is $B_{x_2} = B(x_2, r_2)$. The protecting ball at x_0 will be constructed last. We march from B_{x_2} toward x_0 to construct more protecting balls. For $k \geqslant 3$, we compute the two intersection points between σ and the boundary of $B(x_{k-1}, \frac{6}{5} r_{k-1})$. Among these two points let x_k be the point such that $\angle x_{k-2} x_{k-1} x_k > \pi/2$. One can show that x_k is well-defined and x_k lies between

x_{k-1} and v along σ. Define

$$r_k = \max\left\{ \frac{1}{2}\|x_{k-1} - x_k\|,\ \min_{0\leqslant j\leqslant k}\{\, \lambda \mathrm{f}_{\omega_j}(x_j) + \lambda\|x_j - x_k\|\,\}\right\}$$

where ω_j is some non-critical value for x_j in the range $[0.99\omega, \omega]$ and

$$r_{0k} = \min_{0\leqslant j\leqslant k}\{\, r_j + \lambda\|x_j - x_0\|\,\}.$$

If $B(x_k, r_k) \cap B(x_0, r_{0k}) = \emptyset$, the protecting ball at x_k is

$$B_{x_k} = B(x_k, r_k).$$

(Although we do not perform such a test for $B(x_2, r_2)$, we can define $r_{02} = \min_{0\leqslant j\leqslant 2}\{\, r_j + \lambda\|x_j - x_0\|\,\}$ and prove that $B(x_2, r_2) \cap B(x_0, r_{02}) = \emptyset$.)

Figure 1 shows an example of the construction of B_{x_k}. We force $r_k \geqslant \frac{1}{2}\|x_{k-1} - x_k\|$ so that B_{x_k} overlaps significantly with $B_{x_{k-1}}$.

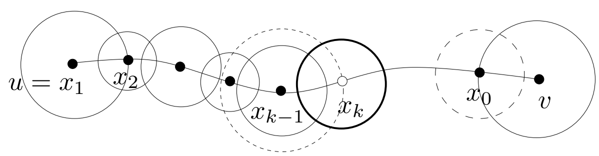

Figure 1: The balls $B(x_{k-1}, \frac{6}{5}r_{k-1})$ and $B(x_0, r_{0k})$ are drawn with dashed circles. The bold circle denotes B_{x_k}.

We continue to march toward x_0 and construct protecting balls until the candidate ball $B(x_m, r_m)$ that we want to put down overlaps with $B(x_0, r_{0m})$. In this case, we reject x_m and $B(x_m, r_m)$. We set the protecting ball at x_0 to be $B_{x_0} = B(x_0, \frac{4}{5}r_{0\,m-1})$. We use $\frac{4}{5}r_{0\,m-1}$ as the radius of B_{x_0} so as to keep some distance between $B_{x_{m-1}}$ and B_{x_0} This will be useful in showing that non-consecutive protecting balls are "far apart". On the other hand, radius(B_{x_0}) cannot be too small so as to keep some distance between B_v and the protecting ball to be placed to cover the gap between $B_{x_{m-1}}$ and B_{x_0}.

To cover the gap between $B_{x_{m-1}}$ and B_{x_0}, a convenient choice is a ball orthogonal to $B_{x_{m-1}}$ and B_{x_0}. The details are as follows. Compute the bisector plane of $B_{x_{m-1}}$ and B_{x_0} with respect to the weighted distance [5] and then intersect it with σ. Select the intersection point y_m that is within a distance of $4r_0$ from x_0. We can show that y_m is uniquely defined and y_m lies between x_{m-1} and x_0 along σ. The protecting ball B_{y_m} is the ball centered at y_m orthogonal to $B_{x_{m-1}}$ and B_{x_0}, i.e., radius$(B_{y_m})^2 = \|x_i - y_m\|^2 -$ radius$(B_{x_i})^2$ for $i \in \{0, m-1\}$. Notice that radius(B_{x_0}) may not be r_0, y_m may not be x_m, and radius(B_{y_m}) may not be r_m. Also, y_m lies outside $B_{x_{m-1}}$ and B_{x_0}; x_{m-1} and x_0 lie outside B_{y_m}.

We turn each protecting ball B_p into a weighted point p with weight $w_p =$ radius$(B_p)^2$. Although our algorithm uses only the weighted vertices, we will refer to both (p, w_p) and the protecting ball B_p in the analysis. Two protecting balls are *consecutive* if their centers are adjacent along σ. Lemma 3.1 below summarizes the properties of the protecting balls. The proof involves some detailed calculations and are left in Appendix C.

Lemma 3.1 *Let σ be a 1-face. The construction of protecting balls for σ terminates and the following properties hold:*

(i) *Two consecutive protecting balls overlap.*

(ii) *For each weighted point p on σ,* $\mathrm{radius}(B_p) \in \big(\frac{\lambda\omega}{21} f(p), 5\lambda g(p)\big)$.

(iii) *Let τ be σ or any 2-face incident to σ. Let p be a weighted point on σ. For any* $r \leqslant 20 \cdot \mathrm{radius}(B_p)$,

  (a) *for any point $z \in B(p,r) \cap \mathrm{mani}(\tau)$, $\angle n_\tau(p), n_\tau(z) < 2\omega$;*

  (b) *$B(p,r) \cap \mathrm{mani}(\tau)$ and $B(p,r) \cap \tau$ are closed balls of dimension* $\dim(\tau)$.

(iv) *For any two consecutive protecting balls B_p and B_q,*

  (a) $q \notin \mathrm{int}\, B_p$,

  (b) *for any point b on the circle* $\mathrm{bd}\, B_p \cap \mathrm{bd}\, B_q$, $\angle pbq < 170°$,

  (c) *$B_p \cup B_q$ contains the subcurve $\sigma(p,q)$ of σ between p and q, and*

  (d) *for any point $z \in \sigma(p,q)$, $B_p \cup B_q$ contains the ball centered at z with radius* $\frac{\lambda\omega}{2500} f(z)$.

(v) *For any non-consecutive protecting balls B_p and B_q, the minimum distance between B_p and B_q is at least* $0.06 \cdot \min\{\mathrm{radius}(B_p), \mathrm{radius}(B_q)\}$*. The points p and q are allowed to lie on different 1-faces.*

4 Meshing a PSC

Our algorithm grows an *admissible* point set S throughout the meshing process. We call a point set admissible if:

- The set contains all weighted points placed in $|\mathcal{D}_{\leqslant 1}|$ during the protection phase.
- Other points in the set are unweighted and they lie outside the protecting balls.

The point set S is initialized to contain the weighted points on the 1-faces. Then, we incrementally insert unweighted vertices in regions not covered by the protecting balls using the Delaunay refinement paradigm. In Section 4.1 we describe the subroutines used by our meshing algorithm. We present the meshing algorithm in Section 4.2. The following notation will be needed.

For each simplex $\eta \in \mathrm{Del}\, S$, V_η denotes the Voronoi cell in $\mathrm{Vor}\, S$ dual to η. For any 1- or 2-face σ and for any simplex $\eta \in \mathrm{Del}\, S$, $V_\eta|_\sigma$ denotes $V_\eta \cap \sigma$. We use $\mathrm{Del}\, S|_\sigma$ to denote the Delaunay subcomplex restricted to σ, i.e., the set of simplices $\eta \in \mathrm{Del}\, S$ such that $V_\eta|_\sigma \neq \emptyset$. We extend the above definition to $\mathcal{D}_i$ for $0 \leqslant i \leqslant 2$ and $\mathcal{D}$:

$$\mathrm{Del}\, S|_{\mathcal{D}_i} = \bigcup_{\sigma \in \mathcal{D}_i} \mathrm{Del}\, S|_\sigma, \qquad \mathrm{Del}\, S|_{\mathcal{D}} = \bigcup_{\sigma \in \mathcal{D}} \mathrm{Del}\, S|_\sigma.$$

For any triangle $t \in \mathrm{Del}\, S|_\sigma$, define $\mathrm{size}(t,\sigma)$ to be the maximum weighted distance between the vertices of t and points in $V_t|_\sigma$. For each point $p \in S \cap \sigma$, we use $\mathrm{star}(p,\sigma)$ to denote the set of edges and triangles in $\mathrm{Del}\, S|_\sigma$ incident to p.

4.1 Subroutines

We describe the subroutines used by our algorithm. Two of them, Infringed and SurfaceNormal, require some preconditions and we show later that the preconditions are satisfied throughout the algorithm.

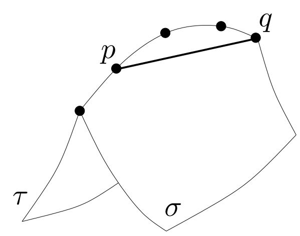

Figure 2: It is dangerous to allow an edge between non-adjacent weighted points. If pq belongs to $\mathrm{Del}\,S|_\tau$ and $\mathrm{Del}\,S|_\sigma$, $\mathrm{Del}\,S|_{\mathcal{D}}$ is pinched at pq.

Multiple intersection. Let p be a point in $S \cap \sigma$ for some 2-face σ. If V_t intersects σ only once for any triangle $t \in \mathrm{star}(p, \sigma)$, $\mathsf{MultiIntersection}(p, \sigma)$ returns null; otherwise, $\mathsf{MultiIntersection}(p, \sigma)$ chooses a triangle $t \in \mathrm{star}(p, \sigma)$ such that V_t intersects σ more than once and returns the point $x \in V_t|_\sigma$ that achieves $\mathrm{size}(t, \sigma)$.

Infringement. Let p be a point in $S \cap \sigma$ for some 2-face σ. We say that (p, σ) is *infringed* if:

- $pq \in \mathrm{Del}\,S|_\sigma$ for some point $q \notin \sigma$, or
- $p \in \mathrm{bd}\,\sigma$ and $pq \in \mathrm{Del}\,S|_\sigma$ for some point $q \in \mathrm{bd}\,\sigma$ not adjacent to p. (See Figure 2.)

We say that pq *certifies* that (p, σ) is infringed. Given (p, σ), the subroutine $\mathsf{Infringed}(p, \sigma)$ requires a precondition that $V_{pq} \cap \mathrm{bd}\,\sigma = \emptyset$ for any edge pq that certifies that (p, σ) is infringed. Assuming this precondition, if (p, σ) is not infringed, $\mathsf{Infringed}(p, \sigma)$ returns null; otherwise, it returns a point in $V_{pq} \cap \sigma$ for some certifying edge pq. The implementation is given below.

Let pq be any edge that certifies that (p, σ) is infringed. $V_{pq} \cap \mathrm{mani}(\sigma)$ is a collection of smooth curves (open or closed). By the precondition, each curve in $V_{pq} \cap \mathrm{mani}(\sigma)$ is either contained in $V_{pq} \cap \sigma$ or disjoint from $V_{pq} \cap \sigma$. Notice that $V_{pq} \cap \sigma \neq \emptyset$ because $pq \in \mathrm{Del}\,S|_\sigma$. We first solve for the intersections between $\mathrm{bd}\,V_{pq}$ and $\mathrm{mani}(\sigma)$. If any intersection obtained belongs to σ, return it and we are done. Otherwise, we need to check the closed curves in $V_{pq} \cap \mathrm{mani}(\sigma)$. Let $E(x) = 0$ be the equation of $\mathrm{mani}(\sigma)$. Choose two unit vectors $\vec{d}_1$ and $\vec{d}_2$ orthogonal and parallel to the plane of V_{pq}, respectively. Let us denote the standard inner product between two vectors $\vec{u}$ and $\vec{v}$ by $\langle \vec{u}, \vec{v} \rangle$. We solve the following system for x: $E(x) = 0$, $\langle \nabla E(x) \times \vec{d}_1, \vec{d}_2 \rangle = 0$. The solutions are the critical points of curves in $V_{pq} \cap \mathrm{mani}(\sigma)$ in direction $\vec{d}_2$. Return any critical point that belongs to σ.

Surface normal variation. Let p be a point in $S \cap \sigma$ for some 2-face σ. Let ϕ_s be a constant chosen from $(4\omega, \frac{\pi}{16})$. Recall that $\sigma_{p,\phi_s} = \{x \in \mathrm{mani}(\sigma) : \angle n_\sigma(p), n_\sigma(x) = \phi_s\}$. The subroutine $\mathsf{SurfaceNormal}(p, \sigma, \phi_s)$ requires a precondition that $\sigma_{p,\phi_s} \cap V_p \cap \mathrm{bd}\,\sigma = \emptyset$.

Given p, σ, and ϕ_s that satisfy the precondition, if $\angle n_\sigma(p), n_\sigma(z) < \phi_s$ for any point $z \in V_p|_\sigma$, $\mathsf{SurfaceNormal}(p, \sigma, \phi_s)$ returns null; otherwise, it returns a point $z \in V_p|_\sigma$ such that $\angle n_\sigma(p), n_\sigma(z) = \phi_s$. We describe the implementation of $\mathsf{SurfaceNormal}(p, \sigma, \phi_s)$ below, which is reminiscent of the computation of $d_\omega(x)$ in setting $\mathrm{f}_\omega(x)$.

First, we can assume that ϕ_s is not a critical value for p, i.e., σ_{p,ϕ_s} is a collection of disjoint smooth closed curves. This can be enforced by perturbing ϕ_s as described in Appendix B. Let $E(x) = 0$ be the equation of $\mathrm{mani}(\sigma)$. Define $G(x) = \langle \nabla E(p), \nabla E(x) \rangle - \|\nabla E(p)\| \cdot \|\nabla E(x)\| \cdot \cos\phi_s$. The set σ_{p,ϕ_s} is described by the system: $E(x) = 0$, $G(x) = 0$. We solve for the set of intersections between σ_{p,ϕ_s} and the support planes of V_p. If any intersection obtained belongs to $V_p|_\sigma$, return it. If no intersection in $V_p|_\sigma$ is obtained, the open curves in $\sigma_{p,\phi_s} \cap V_p$ are disjoint from $V_p|_\sigma$ by the precondition. We still need to check the closed curves in σ_{p,ϕ_s} inside V_p. By the precondition, any closed curve in σ_{p,ϕ_s} inside V_p is either contained in $V_p|_\sigma$ or disjoint from $V_p|_\sigma$. We solve the system: $E(x) = 0$, $G(x) = 0$, $\langle \nabla G(x) \times \nabla E(x), (p - x) \rangle = 0$. The solutions are the tangential contact points with σ_{p,ϕ_s} and balls centered at p. If any tangential contact point belongs to $V_p|_\sigma$, return it. If no contact point is found in $V_p|_\sigma$, σ_{p,ϕ_s} avoids $V_p|_\sigma$ and we return null.

Curve normal variation. Let p be a point in $S \cap \sigma$ for some 2-face σ. Let ϕ_c be a constant chosen from $(\phi_s + 8\omega, \frac{\pi}{16} + 8\omega)$. CurveNormal$(p, \sigma, \phi_c)$ returns a point $y \in V_e|_\sigma$ for some edge $e \in \text{star}(p, \sigma)$ such that $\angle \vec{d_1}, \vec{d_2} = \phi_c$, where $\vec{d_1}$ and $\vec{d_2}$ are the projections of $n_\sigma(p)$ and $n_\sigma(y)$ onto the plane of V_e, respectively. CurveNormal(p, σ, ϕ_c) returns null if such a point does not exist.

The implementation works as follows. Let $E(x) = 0$ be the equation of mani(σ). For each edge $e \in \text{star}(p, \sigma)$, let q_e denote a point on the support plane of V_e; let n_e denote a unit vector orthogonal to the support plane of V_e; and define the function $G_e(x) = \nabla E(x) - \langle \nabla E(x), n_e \rangle \cdot n_e$. Notice that $G_e(x)$ is the projection of the vector $\nabla E(x)$ onto a plane orthogonal to n_e. We solve the system: $E(x) = 0$, $\langle (x - q_e), n_e \rangle = 0$, $\langle G_e(x), G_e(p) \rangle = \|G_e(x)\| \cdot \|G_e(p)\| \cdot \cos \phi_c$. If any solution obtained belongs to $V_e|_\sigma$, return it. Return null if no solution can be returned for any edge in $\text{star}(p, \sigma)$.

Disk neighborhood. Let p be a point in $S \cap \sigma$ for some 2-face σ. NoDisk(p, σ) returns null if the union of triangles in $\text{star}(p, \sigma)$ is a topological disk. Otherwise, NoDisk(p, σ) chooses the triangle $t \in \text{star}(p, \sigma)$ with maximum size(t, σ) and returns the point $x \in V_t|_\sigma$ that achieves size(t, σ). Notice that when $p \in \text{bd}\,\sigma$, p belongs to the boundary of the union of triangles in $\text{star}(p, \sigma)$.

4.2 Algorithm

The following pseudocodes summarize our algorithm. The procedure FindViolation is responsible for the incremental point insertion. The meshing of the input PSC $\mathcal{D}$ is initiated by calling MeshPSC$(\mathcal{D})$. The parameters λ and ω used in protecting the input vertices and 1-faces are picked from $(0, 0.01]$. The angle parameters ϕ_s and ϕ_c are picked from $\left(4\omega, \frac{\pi}{16}\right)$ and $\left(\phi_s + 4\omega, \frac{\pi}{16} + 4\omega\right)$, respectively. A noteworthy feature of our algorithm is that FindViolation never inserts a point inside the protecting balls. This eliminates the check for point insertions inside the protecting balls as well as the exception handler for such events.

FindViolation(p, σ)

1. Call MultiIntersection(p, σ), Infringed(p, σ), SurfaceNormal(p, σ, ϕ_s), CurveNormal(p, σ, ϕ_c), and NoDisk(p, σ) in this order. If a point x is returned in some call, stop and return x.
2. Return null.

MeshPSC$(\mathcal{D})$

1. Protect the vertices and 1-faces in $\mathcal{D}$ with weighted points. Insert a point in each 2-face in $\mathcal{D}$ outside the protecting balls. Let S be the resulting admissible point set.
2. For every (p, σ) where σ is a 2-face and $p \in S \cap \sigma$,
   (a) Compute $v := \mathsf{FindViolation}(p, \sigma)$.
   (b) If v is non-null, insert v into S, and update $\mathrm{Del}\, S$ and $\mathrm{Vor}\, S$.
3. If S has grown in the last execution of step 2, repeat step 2.
4. Return $\mathrm{Del}\, S|_{\mathcal{D}}$.

Step 1 of MeshPSC requires further explanation because a point has to be inserted in each 2-face outside the protecting balls. Let σ be a 2-face. We arbitrarily pick two consecutive protecting balls B_p and B_q, where p and q belong to $\mathrm{bd}\, \sigma$. Lemma 5.3 in Section 5.1 implies the following:

- B_p and B_q induce a Voronoi facet V_{pq} in $\mathrm{Vor}\, S$.
- V_{pq} intersects the subcurve $\sigma(p, q)$ of $\mathrm{bd}\, \sigma$ between p and q exactly once and V_{pq} does not intersect $\mathrm{bd}\, \sigma$ elsewhere.

It follows that we can trace a curve in $V_{pq} \cap \sigma$ starting from the intersection point between V_{pq} and $\sigma(p, q)$. The tracing of this intersection curve either ends at a Voronoi edge in the boundary of V_{pq} or at a tangential contact point between σ and the interior of V_{pq}. We claim that any intersection point between the boundary of V_{pq} and σ and any tangential contact point between V_{pq} and σ is a point in σ outside all protecting balls.

By our claim, a point in σ outside all protecting balls can be computed as follows. First, compute the intersection points between the edges of V_{pq} and σ. If an intersection point is found, return it and we are done. Otherwise, let $H(x) = 0$ be the equation of the support plane of V_{pq} and we compute the tangential contact points between this plane and $\mathrm{mani}(\sigma)$ by solving the system: $E(x) = 0$, $H(x) = 0$, $\nabla H(x) \times \nabla E(x) = 0$. Among the contact points obtained, return one that lies on V_{pq} and σ.

We prove the correctness of our claim. Let x be an intersection point between some edge of V_{pq} and σ. The Voronoi edge containing x is induced by B_p, B_q and a third protecting ball B_s. Observe that B_p and B_s are non-consecutive or B_q and B_s are non-consecutive. So Lemma 3.1(v) implies that the weighted distance from x to B_p, B_q and B_s is positive. Notice that B_p, B_q and B_s are the protecting balls with minimum weighted distance to x. Hence, x lies outside all protecting balls. Let x be a tangential contact point between V_{pq} and σ. Without loss of generality, assume that $\mathrm{radius}(B_p) \geqslant \mathrm{radius}(B_q)$. Because B_p and B_q overlap, $B_q \subseteq B(p, \frac{1}{3}r)$ where $r = 9\, \mathrm{radius}(B_p)$. It follows from Lemma 3.1(iii) and Lemma 1.3(ii) that $\angle n_\sigma(p), pq > \pi/2 - 2\omega$. Since x is a tangential contact point between V_{pq} and σ, $n_\sigma(x)$ is parallel to pq (i.e., orthogonal to V_{pq}), which implies that $\angle n_\sigma(p), n_\sigma(x) > \pi/2 - 2\omega$. Therefore, by Lemma 3.1(iii)(a), x lies outside $B(p, r)$ which contains both B_p and B_q. Hence, the weighted distance from x to B_p and B_q is positive. Notice that B_p and B_q are the protecting balls with minimum weighted distance to x. This implies that x is outside all protecting balls.

5 Analysis of the Meshing Algorithm

We will establish several facts about MeshPSC:

- The algorithm conforms to the 1-faces.
- The preconditions of Infringed and SurfaceNormal are always satisfied.
- The algorithm maintains the admissibility of the point set S.
- The algorithm terminates.
- The extended TBP holds for $\mathcal{D}$ at the termination of the algorithm.

The conformity, preconditions of Infringed and SurfaceNormal, and the maintenance of admissibility is proved in Section 5.1. The termination of MeshPSC is proved in Sections 5.2 and 5.3. The extended TBP is shown to hold in Section 5.4.

Throughout this section, ε denotes a number picked from $(0, \lambda\omega/2500)$ and $\mathrm{lfs}_\varepsilon^*$ denotes $\min\{\mathrm{lfs}_\varepsilon(x) : x \in |\mathcal{D}|\}$. For each point p in an admissible point set, define the *range* of p as $\mathrm{range}(p) = \mathrm{radius}(B_p)$ if p is weighted, and $\mathrm{range}(p) = \mathrm{lfs}_\varepsilon(p)$ otherwise. The *orthoradius* of a triangle $t \in \mathrm{Del}\, S$ is the radius of the smallest ball that has zero weighted distance to the vertices (possibly weighted) of t. We use $\mathrm{ortho}(t)$ to denote the orthoradius of t. Note that if all vertices of t are unweighted, $\mathrm{ortho}(t)$ is just the circumradius of t.

The next lemma shows that $\mathrm{lfs}_\varepsilon(p) = \varepsilon f(p)$ for any point p in an admissible point set. Moreover, our choice of ε makes the protecting ball radii bounded from below by the local feature sizes at the ball centers.

Lemma 5.1 *For any point p in an admissible point set, $\mathrm{lfs}_\varepsilon(p) = \varepsilon f(p)$ and if p is weighted, $\mathrm{range}(p) = \mathrm{radius}(B_p) > \mathrm{lfs}_\varepsilon(p)$.*

Proof. If $p \in \mathcal{D}_0$, $\mathrm{lfs}_\varepsilon(p) = \varepsilon f(p)$ by definition and $\mathrm{radius}(B_p) = \lambda \mathrm{f}_{\omega_p}(p) \geqslant \frac{\lambda\omega_p}{2} f(p)$ by Lemma 2.5(i). As $\omega_p \geqslant 0.99\omega$ by our choice, we get $\mathrm{radius}(B_p) \geqslant \frac{99}{200}\lambda\omega f(p) > \varepsilon f(p) = \mathrm{lfs}_\varepsilon(p)$. Suppose that $p \in \mathrm{int}\,\sigma$ for some 1-face σ. For any endpoint q of σ, $q \in \mathcal{D}_0$ and we have just shown that $\mathrm{radius}(B_q) > \mathrm{lfs}_\varepsilon(q) = \varepsilon f(q)$. By Lemma 3.1(iv)(a) and (v), $p \notin \mathrm{int}\, B_q$ and so $\varepsilon f(q) - \|p-q\| < \mathrm{radius}(B_q) - \|p-q\| \leqslant 0$. Then, it follows from the definition of lfs_ε that $\mathrm{lfs}_\varepsilon(p) = \varepsilon f(p)$. By Lemma 3.1(ii), $\mathrm{radius}(B_p) > \frac{\lambda\omega}{21} f(p) > \varepsilon f(p)$. Suppose that $p \in \mathrm{int}\,\sigma$ for some 2-face σ. Then, p is unweighted. By Lemma 3.1(iv)(d), for any point $q \in \mathrm{bd}\,\sigma$, $B(q, \varepsilon f(q))$ is contained in the union of protecting balls whose centers belong to $\mathrm{bd}\,\sigma$. Thus, $\varepsilon f(q) - \|p-q\| \leqslant 0$ because p lies outside the protecting balls. Thus, the definition of lfs_ε implies that $\mathrm{lfs}_\varepsilon(p) = \varepsilon f(p)$. ◻

The next lemma gives some topological and geometric properties of the neighborhood of a point p in S.

Lemma 5.2 *Let p be a point in $S \cap \sigma$ for some admissible point set S and for some 1- or 2-face σ.*

(i) *For any point $y \in B(p, 20\,\mathrm{range}(p)) \cap \mathrm{mani}(\sigma)$, $\angle n_\sigma(p), n_\sigma(y) < 2\omega$.*

(ii) *For any $r \leqslant 20\,\mathrm{range}(p)$, $B(p,r) \cap \mathrm{mani}(\sigma)$ and $B(p,r) \cap \sigma$ are topological balls of dimension $\dim(\sigma)$.*

(iii) *For any two points $y, z \in B(p, 6\,\mathrm{range}(p)) \cap \mathrm{mani}(\sigma)$, $\angle n_\sigma(p), yz > \pi/2 - 2\omega$.*

Proof. If p is weighted, (i) follows from Lemma 3.1(iii)(a), (ii) follows from Lemma 3.1(iii)(b), and (iii) follows from Lemma 3.1(iii) and Lemma 1.3. Suppose that p is unweighted. By Lemma 5.1, $\mathrm{range}(p) = \mathrm{lfs}_\varepsilon(p) = \varepsilon f(p)$. Then, (i) and (ii) follow from Lemma 2.1, and (iii) is obtained by applying Lemma 1.3. ◻

5.1 Conformity, preconditions and admissibility

Lemma 5.3 below states that the edges between adjacent weighted points are always restricted weighted Delaunay.

Lemma 5.3 *Let S be an admissible point set. Let p and q be two adjacent weighted points on some 1-face σ. Let $\sigma(p,q)$ denote the subcurve between p and q. V_{pq} is the only Voronoi facet in* $\mathrm{Vor}\,S$ *that intersects $\sigma(p,q)$ and V_{pq} intersects $\sigma(p,q)$ exactly once.*

Proof. Let H_{pq} be the bisector of B_p and B_q with respect to the weighted distance. By Lemma 3.1(iv)(a), $p \notin \mathrm{int}\, B_q$ and $q \notin \mathrm{int}\, B_p$. So H_{pq} lies between p and q which means that V_{pq} can potentially intersect $\sigma(p,q)$. We prove that this definitely happens by showing that no other Voronoi facet can intersect $\sigma(p,q)$. Pick an arbitrary point $z \in \sigma(p,q)$. Since $\sigma(p,q) \subset B_p \cup B_q$ by Lemma 3.1(iv)(c), the weighted distance of z from B_p or B_q is negative.

For any unweighted point $s \in S$, since $s \notin B_p \cup B_q$, $\|s - z\| > 0$. This implies that z does not belong to any Voronoi facet partly defined by s. Take any weighted point $s \in S$ other than p and q. Assume without loss of generality that $z \in B_p$. If s is not adjacent to p, $B_p \cap B_s = \emptyset$ by Lemma 3.1(v), implying that $z \notin B_s$. Suppose that s is adjacent to p. Since $B_s \cap \sigma$ is an open curve by Lemma 3.1(iii)(b), $z \notin \mathrm{int}\, B_s$; otherwise, p would also belong to $\mathrm{int}\, B_s$ which is forbidden by Lemma 3.1(iv)(a). We conclude that the weighted distance of z from B_s is non-negative. Thus, z does not lie in any Voronoi facet partly defined by s.

Hence, V_{pq} is the only Voronoi facet in $\mathrm{Vor}\,S$ that intersects $\sigma(p,q)$. Without loss of generality, assume that $\mathrm{radius}(B_p) \geqslant \mathrm{radius}(B_q)$. So $q \in B(p, 2\,\mathrm{range}(p))$. We apply Lemma 5.2(iii) to $B(p, 2\,\mathrm{range}(p)) \cap \sigma$. We conclude that $\angle n_\sigma(p), pq > \pi/2 - 2\omega$ and for any points $y, z \in B(p, 2\,\mathrm{range}(p)) \cap \sigma$, $\angle n_\sigma(p), yz > \pi/2 - 2\omega$. This implies that $\angle pq, yz < 4\omega$. By Lemma 5.2(ii), $B(p, 2\,\mathrm{range}(p)) \cap \sigma$ is an open curve, implying that $\sigma(p,q) \subset B(p, 2\,\mathrm{range}(p)) \cap \sigma$. If $\sigma(p,q)$ intersects V_{pq} at two points y and z, then $\angle pq, yz = \pi/2$, an impossibility. ◻

The next lemma shows that the preconditions of Infringed and SurfaceNormal are always satisfied.

Lemma 5.4 *The preconditions of* Infringed *and* SurfaceNormal *are satisfied when S is an admissible point set.*

Proof. Consider the call Infringed(p, σ). Let pq be any edge that certifies that (p, σ) is infringed. By Lemma 5.3, $V_{pq} \cap \mathrm{bd}\,\sigma = \emptyset$ because either $q \notin \sigma$ or $p \in \mathrm{bd}\,\sigma$ but q is not an adjacent weighted point in $\mathrm{bd}\,\sigma$.

Consider the call SurfaceNormal(p, σ, ϕ_s). If $p \in \mathrm{int}\,\sigma$, by Lemma 5.3, $V_p \cap \mathrm{bd}\,\sigma$ is empty. So is $\sigma_{p,\phi_s} \cap V_p \cap \mathrm{bd}\,\sigma$. Suppose that $p \in \mathrm{bd}\,\sigma$, i.e., p is weighted. Let q be a weighted point in $\mathrm{bd}\,\sigma$ adjacent to p. Let $B(a, \mathrm{range}(a))$ denote the larger ball among B_p and B_q. It follows that $p, q \in B(a, 2\,\mathrm{range}(a))$. By Lemma 5.2(ii), $B(a, 2\,\mathrm{range}(a)) \cap \mathrm{bd}\,\sigma$ contains the subcurve ξ in $\mathrm{bd}\,\sigma$ between p and q. By Lemma 5.2(i), for any point $y \in \xi$, $\angle n_\sigma(p), n_\sigma(y) \leqslant \angle n_\sigma(p), n_\sigma(a) + \angle n_\sigma(y), n_\sigma(a) < 4\omega$. Let s be the other weighted point in $\mathrm{bd}\,\sigma$ adjacent to p.

Let ξ' be the subcurve in $\mathrm{bd}\,\sigma$ between p and s. We can similarly show that $\angle n_\sigma(p), n_\sigma(z) < 4\omega$ for any point $z \in \xi'$. By Lemma 5.3, $V_p \cap \mathrm{bd}\,\sigma$ contains only portions of ξ and ξ'. This implies that $\sigma_{p,\phi_s} \cap V_p \cap \mathrm{bd}\,\sigma = \emptyset$ as $\phi_s > 4\omega$. ◻

The next result shows that MeshPSC maintains the admissibility of S throughout its execution. The proof depends on the usage of weighted Delaunay triangulation.

Lemma 5.5 MeshPSC *never attempts to insert a point in any protecting ball. Hence, the point set S is admissible throughout the execution of* MeshPSC.

Proof. Each point inserted by MultiIntersection and NoDisk lies in the intersection of some Voronoi edge and $|\mathcal{D}|$. Since no three protecting balls intersect by Lemma 3.1(v), all points on a Voronoi edge have positive distance from all vertices, weighted or not. This means that no Voronoi edge intersects a protecting ball. Therefore, the points inserted by MultiIntersection and NoDisk must lie outside all protecting balls.

Assume to the contrary that some call SurfaceNormal(p, σ, ϕ_s) returns a point x that lies inside some protecting ball. So the weighted distance between x and some weighted point is negative. The weighted distance between x and p must also be negative because $x \in V_p|_\sigma$. It follows that p is weighted and $x \in B_p$. Then, by Lemma 5.2(i), $\angle n_\sigma(p), n_\sigma(x) < 2\omega < \phi_s$. This is a contradiction because $\angle n_\sigma(p), n_\sigma(x) = \phi_s$ for x to be returned by SurfaceNormal(p, σ, ϕ_s).

Consider a call Infringed(p, σ) that returns a point $x \in V_{pq} \cap \sigma$ for some edge $pq \in \mathrm{star}(p, \sigma)$. Assume to the contrary that x lies inside some protecting ball. It follows as in the previous paragraph that p and q are weighted and $x \in B_p \cap B_q$. In order to trigger Infringed to return a point, B_p and B_q must be non-consecutive, but then $B_p \cap B_q = \emptyset$ by Lemma 3.1(v), contradicting the fact that $x \in B_p \cap B_q$.

Consider a call CurveNormal(p, σ, ϕ_c) that returns a point $x \in V_{pq} \cap \sigma$ for some edge $pq \in \mathrm{star}(p, \sigma)$. Assume to the contrary that x lies inside some protecting ball. It follows as before that p and q must be weighted and consecutive and $x \in B_p \cap B_q$. Since $x \in B_p$, by Lemma 5.2(i), $\angle n_\sigma(p), n_\sigma(x) < 2\omega$. Let $B(a, \mathrm{range}(a))$ denote the larger of B_p and B_q. So $p, q \in B(a, 2\,\mathrm{range}(a))$. Then, Lemma 5.2(i) and (iii) imply that $\angle n_\sigma(p), pq \geqslant \angle n_\sigma(a), pq - \angle n_\sigma(a), n_\sigma(p) > \pi/2 - 4\omega$. It follows that $\angle n_\sigma(p), V_{pq} < 4\omega$. Let $\vec{d_1}$ and $\vec{d_2}$ denote the projections of $n_\sigma(p)$ and $n_\sigma(x)$ onto the plane of V_{pq}. We have $\angle \vec{d_1}, \vec{d_2} \leqslant \angle n_\sigma(p), n_\sigma(x) + \angle n_\sigma(p), V_{pq} + \angle n_\sigma(x), V_{pq} \leqslant 2 \cdot \angle n_\sigma(p), V_{pq} + 2 \cdot \angle n_\sigma(p), n_\sigma(x) < 12\omega$. However, in order that CurveNormal returns x, we have $\angle \vec{d_1}, \vec{d_2} = \phi_c > \phi_s + 8\omega > 12\omega$, a contradiction. ◻

5.2 Distance lower bounds

MeshPSC grows the point set S incrementally, while maintaining its admissibility by Lemma 5.5. Recall that ε is a number picked from $(0, \lambda\omega/2500)$ and $\mathrm{lfs}^*_\varepsilon$ denotes $\min\{\mathrm{lfs}_\varepsilon(x) : x \in |\mathcal{D}|\}$. We show that MeshPSC maintains inductively a lower bound of $0.03\,\mathrm{lfs}^*_\varepsilon$ on the distances among the points in S. This will allow us to apply the standard packing argument to claim that MeshPSC terminates.

5.2.1 Technical results

We first prove several technical results. The next result states that if the distances among the points in S are $\Omega(\mathrm{lfs}^*_\varepsilon)$ and the orthoradius of a triangle t is small with respect to $\mathrm{lfs}^*_\varepsilon$, the largest angle of t is bounded.

Lemma 5.6 *Let S be an admissible point set. Assume that $\|a-b\| \geqslant 0.03\,\mathrm{lfs}^*_\varepsilon$ for any $a, b \in S$. Let t be a triangle in $\mathrm{Del}\, S|_\sigma$ for some 2-face σ. If $\mathrm{ortho}(t) \leqslant 0.03\,\mathrm{lfs}^*_\varepsilon$, the largest angle of t is less than $170°$.*

Proof. Let p, q and s denote the vertices of t. Since $\mathrm{ortho}(t) \leqslant 0.03\,\mathrm{lfs}^*_\varepsilon$, by Lemma 3.1(v), t has at most two weighted vertices and they must be adjacent along some 1-face. Let $\angle psq$ be the largest angle in t. Assume that $\angle psq > \pi/2$; otherwise, we are done.

Suppose that s is weighted. So p or q is unweighted, say p. Let H_{ps} be the bisector plane of p and B_s. Since p lies outside B_s, H_{ps} lies between p and s and H_{ps} is further from s than p. This implies that the distance between p and the orthocenter of t is at least $\frac{1}{2}\|p-s\|\tan(\angle psq - \pi/2)$. Thus, $0.03\,\mathrm{lfs}^*_\varepsilon \geqslant \mathrm{ortho}(t) \geqslant \frac{1}{2}\|p-s\|\tan(\angle psq-\pi/2) \geqslant 0.015\,\mathrm{lfs}^*_\varepsilon \tan(\angle psq-\pi/2)$, which implies that $\angle psq \leqslant \pi/2 + \arctan(2) < 170°$.

Suppose that s is unweighted. If both p and q are weighted (i.e., p and q are adjacent), $\angle psq$ is maximized when s lies on the circle $\mathrm{bd}\, B_q \cap \mathrm{bd}\, B_q$. By Lemma 3.1(iv)(b), $\angle psq < 170°$ in this case. If p or q is unweighted, say p, we can repeat the argument in the previous paragraph to show that $\angle psq < 170°$. ∎

Under conditions similar to those in Lemma 5.6, the next result states that V_t makes an $O(\omega)$ angle with the surface normals at the vertices of t.

Lemma 5.7 *Let S be an admissible point set. Assume that $\|a-b\| \geqslant 0.03\,\mathrm{lfs}^*_\varepsilon$ for any $a, b \in S$. Let t be a triangle in $\mathrm{Del}\, S|_\sigma$ for some 2-face σ. If $\mathrm{ortho}(t) \leqslant 0.03\,\mathrm{lfs}^*_\varepsilon$, $\angle n_\sigma(x), V_t < 26\omega$ for any vertex x of t.*

Proof. Let p, q and s be the vertices of t. Let $\angle psq$ be the largest angle of t. Lemma 5.6 implies that $60° \leqslant \angle psq < 170°$. Let C be the double cone whose axis is the support line of $n_\sigma(s)$ and whose angular aperture is $\pi - 4\omega$.

We use α to denote $\angle n_\sigma(s), V_t$. We first show that $\alpha < 24\omega$. Assume that $\alpha > 2\omega$; otherwise, we are done. Let H denote the support plane of pqs. Since $\alpha > 2\omega$, $H \cap C$ is a double wedge. Because $\mathrm{range}(s) \geqslant \mathrm{lfs}^*_\varepsilon$ by Lemma 5.1 and $\mathrm{ortho}(t) \leqslant 0.03\,\mathrm{lfs}^*_\varepsilon$, we have $t \subset B(s,r)$, where $r = \mathrm{range}(s) + 0.06\,\mathrm{lfs}^*_\varepsilon < 2\,\mathrm{range}(s)$. By Lemma 5.2(iii), $B(s,r) \cap \sigma$ intersects the double cone C only at s. It follows that p and q do not belong to $H \cap C$. By elementary trigonometry, the angular aperture θ of $H \cap C$ satisfies the equation:

$$\tan\frac{\theta}{2} = \frac{\sqrt{\sin(\alpha - 2\omega)\sin(\alpha + 2\omega)}}{\sin(2\omega)}.$$

If $\alpha \geqslant 24\omega$, the above equation yields $\theta > 170°$. So the angular aperture of the open double wedge $H \setminus C$ is less than $10°$. Since $\angle psq \geqslant 60°$, p and q do not lie in the same wedge of $H \setminus C$. It follows that $\angle psq$ is greater than the angular aperture of $H \cap C$ (i.e., $> 170°$). But this is impossible as $\angle psq < 170°$. Hence, $\alpha < 24\omega$.

Since $t \subset B(s,r)$, by Lemma 5.2(i), for any vertex x of t, $\angle n_\sigma(x), V_t \leqslant \angle n_\sigma(s), V_t + \angle n_\sigma(s), n_\sigma(x) < \alpha + 2\omega < 26\omega$. ∎

Lemma 5.8 *Let S be an admissible point set. Let x be a point in $V_p|_\sigma$ for some 2-face σ. Let D be the minimum distance between x and the points in S. Let W be the minimum weighted distance between x and the points in S. If D or W is less than $\mathrm{lfs}^*_\varepsilon$, then $\|p-x\| < \sqrt{2}\,\mathrm{range}(p)$.*

Proof. We have $W < \mathrm{lfs}^*_\varepsilon$ if D or W is less than $\mathrm{lfs}^*_\varepsilon$. If p is unweighted, $\|p-x\| < \mathrm{lfs}^*_\varepsilon \leqslant \mathrm{lfs}_\varepsilon(p) = \mathrm{range}(p)$. If p is weighted, we have $\|p-x\|^2 - \mathrm{radius}(B_p)^2 < (\mathrm{lfs}^*_\varepsilon)^2$ because $x \in V_p$. Then, Lemma 5.1 implies that $\|p-x\| < (\,\mathrm{radius}(B_p)^2 + (\mathrm{lfs}^*_\varepsilon)^2\,)^{1/2} < \sqrt{2}\,\mathrm{radius}(B_p) = \sqrt{2}\,\mathrm{range}(p)$. ◻

Lemma 5.9 *Let S be an admissible point set. Let σ be a 2-face. Let pq be an edge in $\mathrm{star}(p,\sigma)$ for some point $p \in S \cap \sigma$. If there is a point $x \in V_{pq}|_\sigma$ such that the minimum distance between x and the points in S is less than $\mathrm{lfs}^*_\varepsilon$, then $\angle n_\sigma(p), V_{pq} < 4\omega$.*

Proof. By Lemma 5.8, $\|p-x\| < \sqrt{2}\,\mathrm{range}(p)$ and $\|q-x\| < \sqrt{2}\,\mathrm{range}(q)$. So $\|p-q\| < 2\sqrt{2}\cdot\max\{\mathrm{range}(p),\mathrm{range}(q)\}$. Then, we can apply Lemma 5.2(i) and (iii) to p or q whichever has a larger range and conclude that $\angle n_\sigma(p), pq > \pi/2 - 4\omega$. It follows that $\angle n_\sigma(p), V_{pq} < 4\omega$. ◻

5.2.2 Lower bounds

In this section we give the proofs that MeshPSC maintains inductively a lower bound of $0.03\,\mathrm{lfs}^*_\varepsilon$ on the distances among the points in S. The next lemma shows that the lower bound holds after step 1 of MeshPSC, i.e., protecting the vertices and 1-faces and inserting a point in each 2-face outside the protecting balls.

Lemma 5.10 *Let S be the admissible point set obtained at the end of step 1 of* MeshPSC. *For any $p,q \in S$, $\|p-q\| \geqslant \mathrm{lfs}^*_\varepsilon$.*

Proof. If p and q belong to different 2-faces, $\|p-q\| \geqslant \max\{g(p),g(q)\} \geqslant \max\{f(p),f(q)\}$ which is greater than $\mathrm{lfs}^*_\varepsilon$ by Lemma 5.1. Suppose that p and q belong to the same 2-face. So p or q is weighted. By Lemma 5.1, the radii of protecting balls are greater than $\mathrm{lfs}^*_\varepsilon$. If p or q is unweighted, say q, then q lies outside B_p and so $\|p-q\| \geqslant \mathrm{radius}(B_p) > \mathrm{lfs}^*_\varepsilon$. If both are weighted and p and q are non-adjacent, Lemma 3.1(v) implies that $\|p-q\| > \mathrm{lfs}^*_\varepsilon$. If p and q are weighted and adjacent, by Lemma 3.1(iv)(a), $q \notin \mathrm{int}\,B_p$ and so $\|p-q\| \geqslant \mathrm{radius}(B_p) > \mathrm{lfs}^*_\varepsilon$. ◻

Next, we show that any of the procedure calls within FindViolation(p,σ) preserves the lower bound of $0.03\,\mathrm{lfs}^*_\varepsilon$. MultiIntersection and NoDisk do so only when the lower bound has been preserved so far. Infringed, SurfaceNormal and CurveNormal preserve the lower bound without requiring it as a precondition.

Lemma 5.11 *Let S be the current admissible point set. If $\|a-b\| \geqslant 0.03\,\mathrm{lfs}^*_\varepsilon$ for any $a,b \in S$ and* MultiIntersection(p,σ) *returns a point x, the distance between x and any point in S is at least $0.03\,\mathrm{lfs}^*_\varepsilon$.*

Proof. The point x belongs to $V_t|_\sigma$ for some triangle $t \in \mathrm{star}(p,\sigma)$ such that V_t intersects σ more than once. Let φ be the function that returns the weighted distance between p and points on V_t. By the working of MultiIntersection(p,σ), $\varphi(x)$ is the largest among all points in $V_t|_\sigma$. Assume to the contrary that the minimum distance between x and the points in S is less than $0.03\,\mathrm{lfs}^*_\varepsilon$. So $\varphi(x) < 0.03\,\mathrm{lfs}^*_\varepsilon$ too. Let y be any point in $V_t|_\sigma$ other than x. Because $\varphi(y) \leqslant \varphi(x) < 0.03\,\mathrm{lfs}^*_\varepsilon$, Lemma 5.8 implies that $\max\{\|p-x\|,\|p-y\|\} < \sqrt{2}\,\mathrm{range}(p)$. Then Lemma 5.2(iii) implies that $\angle n_\sigma(p), xy > \pi/2 - 2\omega$. However, since $\mathrm{ortho}(t) \leqslant \varphi(x) < 0.03\,\mathrm{lfs}^*_\varepsilon$, by Lemma 5.7, $\angle n_\sigma(p), xy = \angle n_\sigma(p), V_t < 26\omega < \pi/2 - 2\omega$, a contradiction. ◻

Lemma 5.12 *Let S be the current admissible point set. If* $\mathsf{Infringed}(p, \sigma)$ *returns a point x, the distance between x and any point in S is at least* $0.03\,\mathrm{lfs}^*_\varepsilon$.

Proof. The point x belongs to V_{pq} for some edge $pq \in \mathrm{star}(p, \sigma)$ that certifies that (p, σ) is infringed. In other words, there is a ball B centered at x orthogonal to p and q, and the weighted distances of B from all weighted and unweighted vertices is non-negative. In particular, it means that B does not contain any point in S, weighted or unweighted. Therefore, it suffices to show that $\mathrm{radius}(B) \geqslant 0.03\,\mathrm{lfs}^*_\varepsilon$.

If $q \in \sigma$, p and q must be non-adjacent weighted points in $\mathrm{bd}\,\sigma$. Then, by Lemma 3.1(v) and Lemma 5.1, the distance between B_p and B_q is at least $0.06\,\mathrm{lfs}^*_\varepsilon$. So $\mathrm{radius}(B) \geqslant 0.03\,\mathrm{lfs}^*_\varepsilon$. The other case is that $q \notin \sigma$ and so $\|p - q\| \geqslant \max\{g(p), g(q)\}$. If p is weighted, by Lemma 3.1(ii), $\mathrm{radius}(B_p) < 5\lambda g(p) < g(p)/4 \leqslant \|p - q\|/4$. Similarly, if q is weighted, $\mathrm{radius}(B_q) < \|p - q\|/4$. Therefore, irrespective of whether any of p and q is weighted, B must fill a gap of width at least $\|p - q\|/2 \geqslant g(p)/2$. It follows that $\mathrm{radius}(B) \geqslant g(p)/4 \geqslant f(p)/4$, which is greater than $\mathrm{lfs}^*_\varepsilon$ by Lemma 5.1. ◻

Lemma 5.13 *Let S be the current admissible point set. If* $\mathsf{SurfaceNormal}(p, \sigma, \phi_s)$ *returns a point x, the distance between x and any point in S is at least* $\mathrm{lfs}^*_\varepsilon$.

Proof. By definition, $x \in V_p|_\sigma$ and $\angle n_\sigma(p), n_\sigma(x) = \phi_s > 4\omega$. If the minimum distance between x and the points in S is less than $\mathrm{lfs}^*_\varepsilon$, then $\|p - x\| < \sqrt{2}\,\mathrm{range}(p)$ by Lemma 5.8. But then Lemma 5.2(i) implies that $\angle n_\sigma(p), n_\sigma(x) < 2\omega$, a contradiction. ◻

Lemma 5.14 *Let S be the current admissible point set. If* $\mathsf{Infringed}(p, \sigma)$ *and* $\mathsf{SurfaceNormal}(p, \sigma, \phi_s)$ *return null and* $\mathsf{CurveNormal}(p, \sigma, \phi_c)$ *returns a point x, the distance between x and any point in S is at least* $\mathrm{lfs}^*_\varepsilon$.

Proof. By definition, $x \in V_{pq}|_\sigma$ for some edge $pq \in \mathrm{star}(p, \sigma)$. As $\mathsf{Infringed}(p, \sigma)$ returns null, $q \in \sigma$. Assume to the contrary that the minimum distance between x and the points in S is less than $\mathrm{lfs}^*_\varepsilon$. By Lemma 5.9, $\angle n_\sigma(p), \vec{d} < 4\omega$, where $\vec{d}$ is the projection of $n_\sigma(p)$ onto the plane of V_{pq}. Since $\mathsf{SurfaceNormal}(p, \sigma, \phi_s)$ returns null, $\angle n_\sigma(p), n_\sigma(x) < \phi_s$, which implies that $\angle n_\sigma(x), \vec{d} < \phi_s + 4\omega$. However, as $\mathsf{CurveNormal}(p, \sigma, \phi_c)$ returns x, $\angle n_\sigma(x), \vec{d} \geqslant \phi_c > \phi_s + 4\omega$, a contradiction. ◻

The case for $\mathsf{NoDisk}(p, \sigma)$ is more complicated. Let p be a point on a 2-face σ. We first show that $\mathrm{bd}\,(V_p \cap \sigma)$ is a closed curve under certain conditions. We proved a similar result earlier [14] in the context of meshing a smooth closed surface. That proof can be carried over after some modifications. The details can be found in Appendix D.

Lemma 5.15 *Let S be the current admissible point set. Assume that $\|a - b\| \geqslant 0.03\,\mathrm{lfs}^*_\varepsilon$ for any $a, b \in S$. Let p be a point in $S \cap \sigma$ for some 2-face σ. Suppose that* $\mathsf{Infringed}(p, \sigma)$ *and* $\mathsf{CurveNormal}(p, \sigma, \phi_c)$ *return null and* $\mathrm{size}(t, \sigma) < 0.03\,\mathrm{lfs}^*_\varepsilon$ *for any triangle $t \in \mathrm{star}(p, \sigma)$. Then,* $\mathrm{bd}\,(V_p \cap \sigma)$ *is a closed curve.*

We can then show that any point returned by $\mathsf{NoDisk}(p, \sigma)$ is far from existing vertices.

Lemma 5.16 *Let S be the current admissible point set. Assume that $\|a - b\| \geqslant 0.03\,\mathrm{lfs}^*_\varepsilon$ for any $a, b \in S$. Let p be a point in $S \cap \sigma$ for some 2-face σ. Suppose that* MultiIntersection(p, σ), Infringed(p, σ) *and* CurveNormal(p, σ, ϕ_c) *return null. If* NoDisk(p, σ) *returns a point x, the distance between x and any point in S is at least $0.03\,\mathrm{lfs}^*_\varepsilon$.*

Proof. If the lemma is false, the minimum distance between x and the points in S is less than $0.03\,\mathrm{lfs}^*_\varepsilon$. So is the minimum weighted distance. It follows from the working of NoDisk(p, σ) that $\mathrm{size}(t, \sigma) < 0.03\,\mathrm{lfs}^*_\varepsilon$ for any triangle $t \in \mathrm{star}(p, \sigma)$. Then, Lemma 5.15 implies that $\mathrm{bd}\, V_p|_\sigma$ is a closed curve. Furthermore, $\mathrm{bd}\, V_p|_\sigma$ is not contained in a single facet of V_p because CurveNormal(p, σ, ϕ_c) returns null. Then, because MultiIntersection(p, σ) returns null, the triangles in $\mathrm{star}(p, \sigma)$ form a topological disk by the duality between $\mathrm{bd}\, V_p|_\sigma$ and $\mathrm{star}(p, \sigma)$. But then NoDisk(p, σ) should have returned null, a contradiction. ◻

5.3 Termination

We use the distance lower bounds proved in the last section to show the termination of MeshPSC. Furthermore, we establish some topological properties of the resulting restricted Voronoi diagram.

Theorem 5.1 MeshPSC *terminates. The following properties hold for each point p in the final admissible point set.*

(i) *For any 1- or 2-face σ, V_p intersects σ only if $p \in \sigma$.*

(ii) *For any 2-face σ and for any triangle $t \in \mathrm{star}(p, \sigma)$, $V_t|_\sigma$ is a single point.*

(iii) *For any 1- or 2-face σ and for any edge $e \in \mathrm{star}(p, \sigma)$, $V_e|_\sigma$ is a topological ball of dimension $\dim(\sigma) - 1$;*

(iv) *For any 1- or 2-face σ containing p, $V_p|_\sigma$ is a topological ball of dimension $\dim(\sigma)$.*

Proof. By Lemma 5.10, any two mesh vertices are at distance $0.03\,\mathrm{lfs}^*_\varepsilon$ or more at the end of step 1 of MeshPSC. Afterwards, whenever MeshPSC inserts a new point, its distance from any existing mesh vertex is at least $0.03\,\mathrm{lfs}^*_\varepsilon$ by Lemma 5.11, Lemma 5.12, Lemma 5.13, Lemma 5.14, and Lemma 5.16. Recall that PSC is compact by definition. So a standard packing argument can be used to claim termination.

By Lemma 5.3, V_p intersects a 1-face only if p lies on it. Suppose that V_p intersects a 2-face σ that does not contain p. Recall that MeshPSC inserts an unweighted point in $\mathrm{int}\,\sigma$ during initialization. Walking from V_p to the Voronoi cell of this unweighted point, we must encounter two vertices $q \notin \sigma$ and $s \in \sigma$ such that the common facet V_{qs} intersects σ. But then Infringed(s, σ) should return a point, contradicting the termination of MeshPSC. This proves (i).

The correctness of (ii) follows from the fact that MultiIntersection is not applicable.

Let τ be a 1-face. We have just shown that $\mathrm{star}(p, \tau)$ is nonempty only if $p \in \tau$. By Lemma 5.3, a facet of V_p intersects τ only if the facet is between p and an adjacent weighted point on τ. Such a facet is V_e for some edge $e \in \mathrm{star}(p, \tau)$. Lemma 5.3 further implies that V_e intersects τ in one point. Let σ be a 2-face. By (i) again, $\mathrm{star}(p, \sigma)$ is nonempty only if $p \in \sigma$. For any edge $e \in \mathrm{star}(p, \sigma)$, $V_e|_\sigma$ is a collection of curve(s). There is no closed curve in $V_e|_\sigma$ because CurveNormal(p, σ, ϕ_c) returns null. If there are two open curves in $V_e|_\sigma$, at least three of their endpoints belong to $\text{bd}\,V_e$. Since MultiIntersection returns null, these curve endpoints belong to distinct boundary edges of V_e. Each such curve endpoint gives rise to a triangle in $\text{star}(p, \sigma)$ incident to e, implying that there are at least three triangles in $\text{star}(p, \sigma)$ incident to e. This is a contradiction because NoDisk(p, σ) returns null. Hence, $V_e|_\sigma$ is an open curve. This proves (iii).

Let τ be a 1-face containing p. Lemma 5.3 implies that $\text{bd}\,V_p$ intersects τ exactly twice if $p \in \text{int}\,\tau$, and $\text{bd}\,V_p$ intersects τ exactly once if p is an endpoint of τ. Therefore, $V_p \cap \tau$ is an open curve. Let σ be a 2-face containing p. Because NoDisk(p, σ) returns null, the union of the triangles in $\text{star}(p, \sigma)$ is a topological disk. Since MultiIntersection and CurveNormal return null, no two curves in $\text{bd}\,V_p|_\sigma$ intersect the same edge of V_p and no curve in $\text{bd}\,V_p|_\sigma$ lies within a facet of V_p. Then, the duality between $\text{star}(p, \sigma)$ and $\text{bd}\,V_p|_\sigma$ implies that $\text{bd}\,V_p|_\sigma$ is a closed curve. There is a direction $\vec{d}$ such that for any point $z \in V_p|_\sigma$, $\angle \vec{d}, n_\sigma(z) < \pi/2$. For example, since SurfaceNormal(p, σ, ϕ_s) returns null, we can choose $\vec{d} = n_\sigma(p)$. It has been proved in [14] that for such a 2-manifold $V_p|_\sigma$, its projection to a plane orthogonal to $\vec{d}$ is an injective map. Since the projected image is a planar bounded region with a closed boundary curve, it must be a topological disk. The injectivity of the projection makes it a homeomorphism between the projected image and $V_p|_\sigma$. Hence, $V_p|_\sigma$ is a topological disk. This proves (iv). □

5.4 Topology Preservation

A CW-complex $\mathcal{R}$ is a collection of closed (topological) balls whose interiors are pairwise disjoint and whose boundaries are union of other closed balls in $\mathcal{R}$. A finite set $S \subset |\mathcal{D}|$ has the *extended topological ball properties* (extended TBP) for $\mathcal{D}$ if there is a CW-complex $\mathcal{R}$ with $|\mathcal{R}| = |\mathcal{D}|$ that satisfies the following conditions for each Voronoi face $F \in \text{Vor}\,S$ intersecting $|\mathcal{D}|$:

- (C1) The restricted Voronoi face $F \cap |\mathcal{D}|$ is the underlying space of a CW-complex $\mathcal{R}_F \subseteq \mathcal{R}$.
- (C2) The (closed) balls in $\mathcal{R}_F$ that intersect $\text{int}\,F$ are incident to a unique (closed) ball $b_F \in \mathcal{R}_F$.
- (C3) $b_F \cap \text{bd}\,F$ is a sphere of dimension $\dim(b_F) - 1$.
- (C4) For each ℓ-ball $b \in \mathcal{R}_F \setminus \{b_F\}$ that intersects $\text{int}\,F$, $b \cap \text{bd}\,F$ is a $(\ell - 1)$-ball.

Figure 3 shows two examples of a Voronoi facet F that satisfy the above conditions.

The result of Edelsbrunner and Shah [20] says that if S has the extended TBP for $\mathcal{D}$, the underlying space of $\text{Del}\,S|_\mathcal{D}$ is homeomorphic to $|\mathcal{D}|$. Of course, to apply this result we would require a CW-complex with underlying space $|\mathcal{D}|$. We will see that $\text{Vor}\,S$ restricted to $\mathcal{D}$ provides such a CW-complex when our algorithm terminates.

We first show that the Voronoi faces of $\text{Vor}\,S$ satisfy three properties P1, P2 and P3 listed below at the termination of the algorithm. Let F be a k-face of $\text{Vor}\,S$ where $F = V_{p_1} \cap \ldots \cap V_{p_{(4-k)}}$.

- (P1) If F intersects an element $\sigma \in \mathcal{D}_j$, the intersection is a closed $(k + j - 3)$-ball.
- (P2) There is a unique element $\sigma_F \in \mathcal{D}$ such that the elements intersected by F are σ_F and those that contain σ_F in their boundaries.
- (P3) σ_F contains the vertices $p_1, ..., p_{(4-k)}$.

Property P1 follows from the fact that our algorithm enforces TBP for each 1-face and 2-face (Theorem 5.1). Properties P2 and P3 follow from the following result.

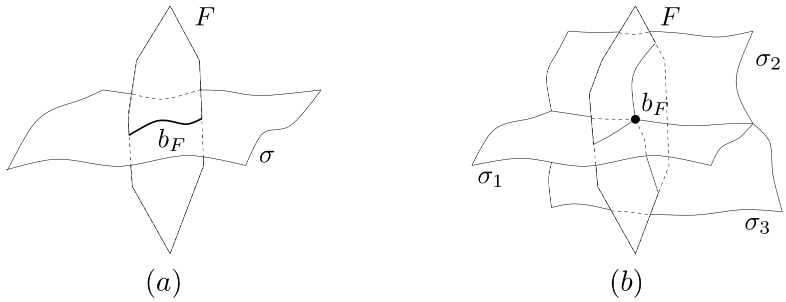

Figure 3: F is a Voronoi facet. In (a), F intersects a 2-face in a closed topological interval (1-ball) which is b_F. Here b_F intersects $\mathrm{bd}\, F$ at two points, a 0-sphere. In (b), F intersects the 1-face in a single point which is b_F, and for $1 \leqslant i \leqslant 3$, $F \cap \sigma_i$ are closed topological 1-balls incident to b_F. Here $b_F \cap \mathrm{bd}\, F = \emptyset$, a -1-sphere.

Lemma 5.17 *Let $F = V_{p_1} \cap \ldots \cap V_{p_{(4-k)}}$ be a k-face in* $\mathrm{Vor}\, S$ *when* MeshPSC *terminates. Let E_F be the set of elements in $\mathcal{D}$ that intersect F.*

(i) *For any $\sigma \in E_F$, σ contains p_i for $1 \leqslant i \leqslant 4-k$.*

(ii) *Let σ_F be an element in E_F of lowest dimension. For any $\sigma \in E_F \setminus \{\sigma_F\}$, $\sigma_F \subset \mathrm{bd}\, \sigma$.*

Proof. For any $\sigma \in E_F$, since σ intersects F, σ intersects V_{p_i} for $1 \leqslant i \leqslant 4-k$. So Theorem 5.1(i) implies that σ contains p_i for $1 \leqslant i \leqslant 4-k$. This proves (i). We conduct a case analysis depending on the dimension of F to prove (ii).

Case 1: F is a Voronoi cell V_p. Take any element $\sigma \in E_F \setminus \{\sigma_F\}$. By (i), $p \in \sigma_F \cap \sigma$ which implies that $\sigma_F \cap \sigma$ contains an element σ' in E_F. If $\sigma_F \not\subset \mathrm{bd}\, \sigma$, $\dim(\sigma') \leqslant \dim(\sigma_F) - 1$. But this contradicts the definition of σ_F.

Case 2: F is a Voronoi facet V_{pq}. Since $p, q \in \sigma_F$ by (i), $\dim(\sigma_F) \geqslant 1$. If p or q lies in the interior of a 2-face, σ_F is this 2-face and E_F is the singleton set $\{\sigma_F\}$. So (ii) is trivially true. Suppose that $p, q \in |\mathcal{D}_{\leqslant 1}|$. We claim that σ_F is an 1-face. Assume to the contrary that σ_F is a 2-face. So $p, q \in \mathrm{bd}\, \sigma_F$ as $p, q \in |\mathcal{D}_{\leqslant 1}|$. The edge pq belongs to $\mathrm{Del}\, S|_{\sigma_F}$ as $V_{pq} \cap \sigma_F = F \cap \sigma_F \neq \emptyset$. The points p and q must be adjacent along $\mathrm{bd}\, \sigma_F$; otherwise, pq would trigger Infringed(p, σ_F) to insert a point, contradicting the termination of MeshPSC. As p and q are adjacent, by Lemma 5.3, $F = V_{pq}$ intersects the subcurve of $\mathrm{bd}\, \sigma_F$ between p and q. Thus, some 1-face in $\mathrm{bd}\, \sigma_F$ belongs to E_F. But this contradicts the definition of σ_F. This proves our claim.

The vertices p and q are adjacent along the 1-face σ_F. If not, for some 2-face σ' containing σ_F in its boundary, pq would trigger Infringed(p, σ') to insert a point. This contradicts the termination of MeshPSC.

Since p and q are adjacent along σ_F, by our protecting ball placement strategy, σ_F is the only 1-face that contains p and q. Moreover, σ_F belongs to the boundary of any 2-face containing p and q. For any $\sigma \in E_F \setminus \{\sigma_F\}$, σ is a 1- or 2-face containing p, q by (i). Consequently, since $\sigma \neq \sigma_F$, σ must be a 2-face that contains σ_F in its boundary.

Case 3: F is a Voronoi edge V_{pqs}. As in case 2, if p, q or s lies in the interior of a 2-face, σ_F is this 2-face and E_F is the singleton set $\{\sigma_F\}$. So (ii) is trivially true. We derive a contradiction for the case that $p, q, s \in |\mathcal{D}_{\leqslant 1}|$. In this case we can show as in case 2 that σ_F must be a 1-face. By (i), $p, q, s \in \sigma_F$. Since any three vertices on any 1-face cannot be mutually adjacent, two vertices in $\{p, q, s\}$ are non-adjacent along σ_F, say p and q. But then the edge pq would trigger Infringed(p, σ') to insert a point for some 2-face σ' whose boundary contains σ_F. This contradicts the termination of MeshPSC. ◻

We show that conditions C1–C4 follow from P1 and P2. The proof depends on our assumption of the generic intersection property (GIP): every Voronoi face in $\mathrm{Vor}\, S$ intersects an element of $\mathcal{D}$ transversally, if at all, at the termination of MeshPSC. The assumption of the GIP can be removed by employing a more elaborate analysis and symbolic perturbation in the algorithm; see [14] for details.

Theorem 5.2 *The output of* MeshPSC *is homeomorphic to* $|\mathcal{D}|$*.*

Proof. We prove that S has the extended TBP for $\mathcal{D}$ when MeshPSC terminates. Then, the result of Edelsbrunner and Shah implies the theorem. To apply the extended TBP we need a CW-complex whose underlying space is $|\mathcal{D}|$. Upon the termination of MeshPSC, we take $\mathcal{R}$ to be the complex formed by the intersection of the Voronoi faces in $\mathrm{Vor}\, S$ with the elements of $\mathcal{D}$. So $\mathcal{R}$ is a CW-complex according to P1. It immediately follows that a k-face F of $\mathrm{Vor}\, S$ intersects the CW-complex $\mathcal{R}$ in a CW-complex $\mathcal{R}_F \subseteq \mathcal{R}$ satisfying C1.

Define σ_F as in the statement of Lemma 5.17. We identify b_F as $\sigma_F \cap F$. By the assumption of GIP, σ_F intersects $\mathrm{int}\, F$ and so b_F intersects $\mathrm{int}\, F$. Condition C2 then follows from P2.

If $\mathrm{bd}\, b_F$ intersects $\mathrm{int}\, F$, $\mathrm{bd}\, \sigma_F$ intersects F too. But this would contradict P2 which states that any element other than b_F that intersects F contains σ_F in its boundary. Therefore, $\mathrm{bd}\, b_F$ must be contained in $\mathrm{bd}\, F$. The assumption of GIP allows us to conclude that $b_F \cap \mathrm{bd}\, F = \mathrm{bd}\, b_F$. Thus, $b \cap \mathrm{bd}\, F = \mathrm{bd}\, b_F$ is a sphere of dimension $\dim(b_F) - 1$. This establishes condition C3.

Take any ball $b \in \mathcal{R}_F \setminus \{b_F\}$ that intersects $\mathrm{int}\, F$. Let σ be the element in $\mathcal{D}$ such that $b = \sigma \cap F$. We claim that $\mathrm{bd}\, \sigma \cap \mathrm{int}\, F$ is an open $(\ell - 1)$-ball, where $\ell = \dim(F) + \dim(\sigma) - 3$. The proof involves a case analysis depending on $\dim(\sigma_F)$.

- If σ_F is a 2-face, σ does not exist by P2 because $\mathcal{D}$ does not contain any element with dimension higher than two. So there is nothing to prove.

- If σ_F is a 1-face, P2 implies that σ is a 2-face, F does not contain any vertex of σ, and F does not intersect any 1-face other than σ_F. It follows that $\mathrm{bd}\, \sigma \cap \mathrm{int}\, F = \sigma_F \cap \mathrm{int}\, F$ which by P1, is an open $(\ell - 1)$-ball as $\dim(\sigma_F) = \dim(\sigma) - 1$.

- If σ_F is a vertex in $\mathcal{D}_0$, F must be the Voronoi cell owned by σ_F. So F does not contain any vertex of σ other than σ_F. By P2, any 1-face intersected by F is incident to σ_F. If σ is a 1-face, $\mathrm{bd}\, \sigma \cap \mathrm{int}\, F = \sigma_F$ which is an open $(\ell - 1)$-ball as $\ell = 1$ in this case. If σ is a 2-face, $\mathrm{bd}\, \sigma \cap \mathrm{int}\, F$ is equal to $(\sigma_1 \cup \sigma_2) \cap \mathrm{int}\, F$, where σ_1 and σ_2 are the two 1-faces in $\mathrm{bd}\, \sigma$ incident to σ_F. Since the other endpoints of σ_1 and σ_2 are outside F, $(\sigma_1 \cup \sigma_2) \cap \mathrm{int}\, F$ is an open curve which is an open $(\ell - 1)$-ball as $\ell = 2$ in this case.

By P1, $b = \sigma \cap F$ is a closed ℓ-ball, where $\ell = \dim(F) + \dim(\sigma) - 3$. So $\mathrm{bd}\, b$ is a $(\ell - 1)$-sphere. Observe that $\mathrm{bd}\, b = (\mathrm{bd}\, b \cap \mathrm{int}\, F) \cup (\mathrm{bd}\, b \cap \mathrm{bd}\, F)$. Furthermore, $\mathrm{bd}\, b = (\mathrm{bd}\, b \cap \mathrm{int}\, F)$ and $\mathrm{bd}\, b \cap \mathrm{bd}\, F$ are disjoint, $\mathrm{bd}\, b \cap \mathrm{bd}\, F = b \cap \mathrm{bd}\, F$, and $\mathrm{bd}\, b \cap \mathrm{int}\, F = \mathrm{bd}\, \sigma \cap \mathrm{int}\, F$. We have shown in the above that $\mathrm{bd}\, \sigma \cap \mathrm{int}\, F$ is an open $(\ell - 1)$-ball. Hence, $b \cap \mathrm{bd}\, F$ must be a closed $(\ell - 1)$-ball in order to merge with the open $(\ell - 1)$-ball $\mathrm{bd}\, \sigma \cap \mathrm{int}\, F$ to form the $(\ell - 1)$-sphere $\mathrm{bd}\, b$. This proves condition C4. ◻

Non-smooth features. Following the proof of Edelsbrunner and Shah [20] one can construct a homeomorphism h between the underlying spaces of $\mathcal{D}$ and $\mathrm{Del}\,S|_{\mathcal{D}}$ such that h respects each stratum. That is, the restriction of h to $|\mathcal{D}_i|$ is a homeomorphism to the underlying space of $\mathrm{Del}\,S|_{\mathcal{D}_i}$. Property P3 ensures that $\mathrm{Del}\,S|_{\mathcal{D}_i}$ has all vertices in $\mathcal{D}_i$. This property ensures the preservation of non-smooth features. For example, a 1-face in $\mathcal{D}$ is meshed with a set of Delaunay edges connecting consecutive points sampled on it thereby preserving the 'non-smoothness' of the 1-face in the output.

6 Mesh quality improvements

We introduce two subroutines that enhance MeshPSC to offer bounds on the aspect ratios, normal variation, and dihedral angles in the output mesh. For a triangle t, the *radius-edge ratio* is defined to be the ratio of the circumradius to the shortest edge length of t. It is known that the radius-edge ratio of a triangle is bounded if and only if the aspect ratio is bounded. For a point p on a 2-face σ, the first subroutine Shape eliminates triangles in $\mathrm{star}(p,\sigma)$ with unweighted vertices and radius-edge ratios at least 1; the second subroutine TriangleNormal eliminates triangles in $\mathrm{star}(p,\sigma)$ whose normal deviates from $n_\sigma(p)$ by 26ω or more. Notice that we do not offer any guarantee on the shape of triangles incident to weighted vertices. The pseudocodes of Shape and TriangleNormal are given below.

Shape(p, σ)

1. If there is a triangle $t \in \mathrm{star}(p,\sigma)$ whose vertices are unweighted and radius-edge ratio is greater than or equal to 1, return $V_t|_\sigma$.
2. Otherwise, return null.

TriangleNormal(p, σ)

1. If there is a triangle $t \in \mathrm{star}(p,\sigma)$ such that $\angle n_\sigma(p), V_t \geqslant 26\omega$, return $V_t|_\sigma$.
2. Otherwise, return null.

We modify FindViolation(p,σ) to call the subroutines MultiIntersection(p,σ), Infringed(p,σ), SurfaceNormal(p,σ,ϕ_s), CurveNormal(p,σ,ϕ_c), Shape(p,σ), and TriangleNormal(p,σ) in this order. If any call returns a point, stop and return that point. Otherwise, return null. The enhanced algorithm works like MeshPSC by calling the enhanced FindViolation until the admissible point set stops growing. We call this enhanced algorithm QualMeshPSC.

By Theorem 5.1, when FindViolation(p,σ) calls Shape(p,σ) or TriangleNormal(p,σ), $V_t|_\sigma$ is a point for any triangle $t \in \mathrm{star}(p,\sigma)$. Thus, the subroutines Shape and TriangleNormal are well-defined. The next two lemmas show that termination is still guaranteed.

Lemma 6.1 *Let S be the current admissible point set. Assume that $\|a-b\| \geqslant 0.03\,\mathrm{lfs}^*_\varepsilon$ for any $a,b \in S$. If* FindViolation(p,σ) *calls* Shape(p,σ) *and* Shape(p,σ) *returns a point x, the distance between x and any point in S is at least $0.03\,\mathrm{lfs}^*_\varepsilon$.*

Proof. The point x returned is $V_t|_\sigma$ for some triangle $t \in \mathrm{star}(p,\sigma)$. Since the radius-edge ratio of t is at least 1, the distance from x to the vertices of t is at least the shortest edge length of t. The invocations of MultiIntersection, Infringed, SurfaceNormal, and CurveNormal have kept a distance lower bound of $0.03\,\mathrm{lfs}^*_\varepsilon$ among the vertices. So the distances from x to the vertices of t is at least $0.03\,\mathrm{lfs}^*_\varepsilon$. Because the vertices of t are unweighted, $0.03\,\mathrm{lfs}^*_\varepsilon$ is also a lower bound on the minimum weighted distance between x and the points in S. In turn, this implies that the distance between x and any point in S is at least $0.03\,\mathrm{lfs}^*_\varepsilon$. ◻

Lemma 6.2 *Let S be the current admissible point set. Assume that $\|a-b\| \geqslant 0.03\,\mathrm{lfs}^*_\varepsilon$ for any $a, b \in S$. If* FindViolation(p,σ) *calls* TriangleNormal(p,σ) *and* TriangleNormal(p,σ) *returns a point x, the distance between x and any point in S is greater than* $0.03\,\mathrm{lfs}^*_\varepsilon$.

Proof. Let t be the triangle in $\mathrm{star}(p,\sigma)$ such that $V_t|_\sigma = x$. The contrapositive of Lemma 5.7 implies that $\mathrm{ortho}(t) > 0.03\,\mathrm{lfs}^*_\varepsilon$. So the weighted distance between x and any point in S is greater than $0.03\,\mathrm{lfs}^*_\varepsilon$. So is the distance between x and any point in S. ◻

Because a distance lower bound is maintained by Shape and TriangleNormal, QualMeshPSC terminates. The guarantees in Theorem 5.1 and Theorem 5.2 are preserved as MultiIntersection, Infringed, SurfaceNormal, and CurveNormal are inapplicable at the termination of QualMeshPSC.

The inapplicability of Shape implies that all output triangles with unweighted vertices have radius-edge ratio less than 1. The angles of such triangles lie in the range $\left(\frac{\pi}{6}, \frac{2\pi}{3}\right)$. There is no shape guarantee for output triangles incident on weighted vertices.

TriangleNormal ensures that the normal of any triangle deviates from the surface normals at its vertices by less than 26ω. Furthermore, we prove below that for any 2-face σ, the dihedral angle between any pair of triangles in $\mathrm{Del}\,S|_\sigma$ sharing an edge is greater than $\pi - 52\omega$. Notice that for any edge in $\mathrm{Del}\,S|_\sigma$ with both endpoints weighted, the edge is incident to only one triangle in $\mathrm{Del}\,S|_\sigma$; and for any edge in $\mathrm{Del}\,S|_\sigma$ with an unweighted endpoint, the edge is incident to two triangles in $\mathrm{Del}\,S|_\sigma$.

Let t_1 and t_2 be two triangles in $\mathrm{Del}\,S|_\sigma$ that share an edge pq. For $i \in \{1,2\}$, let z_i denote $V_{t_i}|_\sigma$ and let $\vec{d_i}$ be a unit vector parallel to V_{t_i} and making an acute angle with $n_\sigma(z_i)$. TriangleNormal ensures that $\angle V_{t_1}, V_{t_2} \leqslant \angle n_\sigma(p), V_{t_1} + \angle n_\sigma(p), V_{t_2} < 52\omega$. Consequently, the dihedral angle between t_1 and t_2 is either less than 52ω or greater than $\pi - 52\omega$. To complete the argument, it suffices to show that both t_1 and t_2 can be oriented in the clockwise order as seen from infinity in $\vec{d_1}$ such that pq is assigned opposite orientations.

Since SurfaceNormal(p,σ,ϕ_s) returns null, $\angle n_\sigma(p), n_\sigma(z_i) < \phi_s < \pi/16$. So

$$\begin{aligned} \angle n_\sigma(z_i), \vec{d_i} &= \angle n_\sigma(z_i), V_{t_i} \\ &\leqslant \angle n_\sigma(p), V_{t_i} + \angle n_\sigma(p), n_\sigma(z_i) \\ &< 26\omega + \phi_s \\ &< \pi/2. \end{aligned} \tag{1}$$

Denote by O_i the clockwise ordering of the three restricted Voronoi cells locally around z_i if we view from infinity in $\vec{d_i}$. As $n_\sigma(z_i), \vec{d_i} < \pi/2$, O_i is identical to the clockwise ordering of the restricted Voronoi cells locally around z_i as seen from infinity in $n_\sigma(z_i)$.

If we view from infinity in $\vec{d_i}$, the clockwise ordering of t_i is consistent with the clockwise ordering of the three Voronoi cells incident to V_{t_i}. As $n_\sigma(z_i), \vec{d_i} < \pi/2$, this ordering is consistent with the clockwise ordering of the three restricted Voronoi cells locally around z_i as seen from infinity in $n_\sigma(z_i)$. We conclude that O_i is consistent with the clockwise ordering of t_i as seen from infinity in $\vec{d_i}$.

By Theorem 5.1, $V_{pq}|_\sigma$ is an open curve with endpoints z_1 and z_2. Therefore, the ordering of the two restricted Voronoi cells incident to $V_{pq}|_\sigma$ in O_1 is the reverse of that in O_2. This implies that pq is assigned opposite orientations when we orient t_i in the clockwise order as viewed from $\vec{d_i}$ for $i \in \{1, 2\}$. Since $\mathsf{SurfaceNormal}(p, \sigma, \phi_s)$ returns null,

$$\begin{aligned}
\angle \vec{d_1}, \vec{d_2} &\leqslant \angle n_\sigma(z_1), \vec{d_1} + \angle n_\sigma(z_2), \vec{d_2} + \angle n_\sigma(p), n_\sigma(z_1) + \angle n_\sigma(p), n_\sigma(z_2) \\
&< \angle n_\sigma(z_1), \vec{d_1} + \angle n_\sigma(z_2), \vec{d_2} + 2\phi_s \\
&\overset{(1)}{<} 52\omega + 4\phi_s \\
&< \pi/2.
\end{aligned}$$

So the clockwise ordering of t_2 remains the same if we change the viewpoint to infinity in $\vec{d_1}$. In all, if we orient t_1 and t_2 in the clockwise order as seen from $\vec{d_1}$, pq is assigned two opposite orientations.

We summarize with the following theorem.

Theorem 6.1 *Let $\omega \in (0, 0.01]$.* $\mathsf{QualMeshPSC}$ *terminates with the following guarantees.*

- *The output mesh is homeomorphic to $|\mathcal{D}|$.*
- *Each mesh triangle with unweighted vertices has radius-edge ratio less than 1.*
- *Let S be the output vertex set. For each element $\sigma \in \mathcal{D}$, the vertices of $\mathrm{Del}\, S|_\sigma$ belong to σ and $\mathrm{Del}\, S|_\sigma$ is homeomorphic to σ.*
- *For each mesh triangle, the angle between its normal and the surface normal at any of its vertex is less than 26ω.*
- *For each mesh edge with an unweighted endpoint, the two mesh triangles incident to it make a dihedral angle greater than $\pi - 52\omega$.*

7 Conclusions

We have presented a Delaunay refinement algorithm to mesh PSC which preserves topology as well as non-smooth features. We did not handle explicitly three-dimensional elements in PSC. Our algorithm can be extended to handle these inputs. This will require further analysis along the line of [28, 31]. Our bound of 0.01 for λ and ω is very small. So is the lower bound of $0.03\, \mathrm{lfs}^*_\varepsilon$ on the distances among the points inserted during Delaunay refinement, where $\varepsilon = \lambda\omega/2500$. We believe that these pessimistic bounds are not the tightest possible and that the algorithm performs better in practice. A main question is how to make this algorithm practical. The main bottlenecks in practice are the surface normal and curve normal variation tests. Can these be eliminated or replaced with some other easier computations? Some answer to this question has been provided in [12] recently. We propose to handle parametric surfaces by first implicitizing them [24]. It is an interesting question whether our algorithm can be adapted to work directly with parametric surfaces. Our analysis guarantees homeomorphism between the input and the output. A stronger and more desirable condition would be to ensure an isotopy between the two. Does it require a different approach? Or, is it true that topological ball property indeed guarantees an isotopy at least for surfaces in three dimensions? We leave these questions as open.

Acknowledgment

We thank the anonymous referees for comments and feedbacks that helped improving the paper. We also thank `Skype` for making communication and collaboration so much easier.

References

[1] P. Alliez, D. Cohen-Steiner, M. Yvinec, and M. Desbrun. Variational tetrahedral meshing. *ACM Transactions on Graphics*, 24 (2005), 617–625.

[2] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete and Computational Geometry*, 22 (1999), 481–504.

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

[4] N. Amenta and T. K. Dey. Normal variation with adaptive feature size. http://www.cse.ohio-state.edu/ tamaldey/paper/norvar/norvar.pdf.

[5] F. Aurenhammer. Voronoi diagrams – a survey of a fundamental geometric data structure. *ACM Computing Surveys*, 23 (1991), 345-405.

[6] J.-D. Boissonnat and F. Cazal. Natural neighbor coordinates of points on a surface. *Computational Geometry: Theory and Applications*, 19 (2001), 155–173.

[7] J.-D. Boissonnat and S. Oudot. Provably good sampling and meshing of Lipschitz surfaces. *Proceedings of the 22nd Annual Symposium on Computational Geometry*, 2006, 337–346.

[8] J.-D. Boissonnat and S. Oudot. Provably good surface sampling and meshing of surfaces. *Graphical Models*, 67 (2005), 405–451.

[9] Y.-S. Chang and H. Qin. A unified subdivision approach for multi-dimensional non-manifold modeling. *Computer-Aided Design*, 38 (2006), 770–785.

[10] F. Chazal, D. Cohen-Steiner, and A. Lieutier. A sampling theory for compact sets in Euclidean space. *Proc. 22nd Ann. Sympos. Comput. Geom.* (2006), 319–326.

[11] S.-W. Cheng and S.-H. Poon. Three-dimensional Delaunay mesh generation. *Discrete and Computational Geometry*, 36 (2006), 419–456.

[12] S.-W. Cheng, T. K. Dey, and J. A. Levine. A practical Delaunay meshing algorithm for a large class of domains. *Proceedings of the 16th International Meshing Roundtable*, 2007, 477–494.

[13] S.-W. Cheng, T. K. Dey, E. A. Ramos and T. Ray. Quality meshing for polyhedra with small angles. *International Journal of Computational Geometry and Applications*, 15 (2005), 421–461.

[14] S.-W. Cheng, T. K. Dey, E. A. Ramos and T. Ray. Sampling and meshing a surface with guaranteed topology and geometry. To appear in *SIAM Journal on Computing*. A preliminary version appeared in*Proceedings of the 20th Annual Symposium on Computational Geometry*, 2004, 280–289.

[15] L. P. Chew. Guaranteed-quality triangular meshes. Report TR-98-983, Computer Science Department, Cornell University, Ithaca, New York, 1989.

[16] L. P. Chew. Guaranteed-quality mesh generation for curved surfaces. *Proceedings of the 9th Annual Symposium Computational Geometry*, 1993, 274–280.

[17] D. Cohen-Steiner, E. C. de Verdière and M. Yvinec. Conforming Delaunay triangulations in 3D. *Proceedings of the 18th Annual Symposium on Computational Geometry*, 2002, 199–208.

[18] T. K. Dey, G. Li, and T. Ray. Polygonal surface remeshing with Delaunay refinement. *Proceedings of the 14th International Meshing Roundtable*, 2005, 343–361.

[19] T. K. Dey. Curve and surface reconstruction : Algorithms with mathematical analysis. Cambridge University Press, New York, 2007.

[20] H. Edelsbrunner and N. Shah. Triangulating topological spaces. *International Journal of Computational Geometry and Applications*, 7 (1997), 365–378.

[21] L. De Floriani, P. Magillo, E. Puppo, and D. Sobrero. A multi-resolution topological representation for non-manifold meshes. *Computer-Aided Design*, 36 (2004), 141–159.

[22] S. Funke and E.A. Ramos. Smooth-surface reconstruction in near-linear time. *Proc. 13th Annu. ACM-SIAM Sympos. Discrete. Alg.*, 2002, 781–790.

[23] J. Giesen and U. Wagner. Shape Dimension and Intrinsic Metric from Samples of Manifolds. *Discrete and Computational Geometry*, 32 (2004), 245–267.

[24] D. Manocha and J.F. Canny. Implicit representation of rational parametric surfaces. *Journal of Symbolic Computation*, 13 (1992), 485–510.

[25] S.-H. Lee and K. Lee. Partial entity structure: a compact non-manifold boundary representation based on partial topological entities. *Proceedings of the 6th ACM Symposium on Solid Modeling and Application*, 2001, 159–170.

[26] J. Milnor. *Morse Theory.* Princeton University Press, 1963.

[27] M. Murphy, D. M. Mount and C. W. Gable. A point placement strategy for conforming Delaunay tetrahedralization. *International Journal of Computational Geometry and Applications*, 11 (2001), 669–682.

[28] S. Oudot, L. Rineau, and M. Yvinec. Meshing volumes bounded by smooth surfaces. *Proceedings of the 14th International Meshing Roundtable*, 2005, 203–219.

[29] S. Pav and N. Walkington. Robust three dimensional Delaunay refinement. *Proceedings of the 13th International Meshing Roundtable* 2004.

[30] A.A.G. Requicha and J.R. Rossignac. Solid modeling and beyond. *Computer Graphics and Applications*, 12 (1992), 31–44.

[31] L. Rineau and M. Yvinec. Meshing 3D domains bounded by piecewise smooth surfaces. *Proceedings of the 16th International Meshing Roundtable*, 2007, 443–460.

[32] J. Ruppert. A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *Journal of Algorithms*, 18 (1995), 548–585.

[33] J. R. Shewchuk. Tetrahedral mesh generation by Delaunay refinement. *Proceedings of the 14th Annual Symposium on Computational Geometry*, 1998, 86–95.

[34] J. R. Shewchuk. Mesh generation for domains with small angles. *Proceedings of the 16th Annual Symposium on Computational Geometry*, 2000, 1–10.

[35] S. Sternberg. *Lectures on Differential Geometry*, Chelsea Publishing Company, 1983.

[36] D. Talmor, *Well-spaced points for numerical methods*, Report CMU-CS-97-164, Dept. Comput. Sci., Carnegie-Mellon Univ., Pittsburgh, Penn., 1997.

Appendix

A Background results

Proof of Lemma 1.1. If σ is a 1-face, $B(x,r) \cap \mathrm{mani}(\sigma)$ may be split by the endpoints of σ into at most three subcurves, one of which is $B(x,r)\cap\sigma$. Suppose that σ is a 2-face. If $x \in \mathrm{int}\,\sigma$, since $B(x,r)$ avoids 1-faces in this case by assumption, $B(x,r) \cap \sigma = B(x,r) \cap \mathrm{mani}(\sigma)$ and we are done. Suppose that $x \in \mathrm{bd}\,\sigma$. We claim that $B(x,r) \cap \mathrm{bd}\,\sigma$ is an open curve.

By assumption, $B(x,r)$ avoids any 1-face that does not contain x. If x is incident to only one 1-face σ_1 (i.e., $x \in \mathrm{int}\,\sigma_1$), $B(x,r)\cap\mathrm{bd}\,\sigma = B(x,r)\cap\sigma_1$ and we have just shown that $B(x,r)\cap\sigma_1$ is an open curve. If x is a common endpoint of two 1-faces σ_1 and σ_2 in $\mathrm{bd}\,\sigma$, $B(x,r)$ does not contain any vertex of σ_i other than x as $B(x,r)$ avoids other 1-faces by assumption. We know that $B(x,r) \cap \sigma_1$ and $B(x,r) \cap \sigma_2$ are open curves, and so they concatenate to form one open curve. This proves our claim.

Since $B(x,r) \cap \mathrm{bd}\,\sigma$ is an open curve, it splits $B(x,r) \cap \mathrm{mani}(\sigma)$ into two topological disks, one of which is $B(x,r) \cap \sigma$.

Proof of Lemma 1.2. We use $\mathcal{C}$ to denote $H \cap B \cap \mathrm{mani}(\sigma)$. No curve in $\mathcal{C}$ is closed. Otherwise, the tangent to some point $z \in \mathcal{C}$ would be parallel to the projection $\vec{d}$ of $n_\sigma(x)$ on H, implying that $\angle n_\sigma(x), n_\sigma(z) \geqslant \angle n_\sigma(z), \vec{d} - \angle n_\sigma(x), \vec{d} \geqslant \pi/2 - \theta$. But this is a contradiction because $\angle n_\sigma(x), n_\sigma(z) < \theta$. Notice that the endpoints of each curve in $\mathcal{C}$ lie on the boundary of B.

Notice that y is the center of $H \cap B$ and some curve in $\mathcal{C}$ passes through y. Assume to the contrary that $\mathcal{C}$ consists of at least two open curves. Then, there exists a concentric disk D inside the disk $H \cap B$ such that D is tangent to some curve in $\mathcal{C}$ at a point z. Let ℓ be the support line of yz. Since ℓ intersects $B \cap \mathrm{mani}(\sigma)$ at y and z, we can find a point $z' \in \ell \cap B \cap \mathrm{mani}(\sigma)$ such that z and z' are consecutive intersections points in $\ell \cap B \cap \mathrm{mani}(\sigma)$. As B is convex, z and z' are also consecutive intersection points in $\ell \cap \mathrm{mani}(\sigma)$. We have $\angle n_\sigma(z), \ell = \angle n_\sigma(z), H \leqslant \angle n_\sigma(x), H + \angle n_\sigma(x), n_\sigma(z) < 2\theta$ by assumption. Also, $\angle n_\sigma(z'), \ell \leqslant \angle n_\sigma(z), \ell + \angle n_\sigma(z), n_\sigma(z') \leqslant \angle n_\sigma(z), \ell + \angle n_\sigma(x), n_\sigma(z) + \angle n_\sigma(x), n_\sigma(z') < 4\theta$. If we walk along ℓ, we enter $\mathrm{mani}(\sigma)$ at z and exit $\mathrm{mani}(\sigma)$ at z' or vice versa. We conclude that $\angle n_\sigma(z), n_\sigma(z') > \pi - 6\theta$ because $n_\sigma(z)$ and $n_\sigma(z')$ are outward surface normals. But this is a contradiction because $\angle n_\sigma(z), n_\sigma(z') \leqslant \angle n_\sigma(x), n_\sigma(z) + \angle n_\sigma(x), n_\sigma(z') < 2\theta \leqslant \pi - 6\theta$ as $\theta \leqslant \pi/8$.

Proof of Lemma 1.3. Consider (i). We claim that z lies strictly inside the double cone with apex y, angular aperture 2θ, and axis through y parallel to the tangent to $\mathrm{mani}(\sigma)$ at x. This claim immediately implies that $\angle n_\sigma(x), yz > \pi/2 - \theta$. Assume to the contrary that the claim is false. Since $\angle n_\sigma(x), n_\sigma(y) < \theta$ by assumption, the tangent to $\mathrm{mani}(\sigma)$ at y lies strictly inside the double cone. This implies that if we walk along the subcurve of $B(x,r) \cap \mathrm{mani}(\sigma)$ from y to z, we stay inside the double cone initially. We must reach the double cone boundary for the first time at some point z_1 after leaving y. Let H be the plane tangent to the double cone at z_1. We translate H to a plane that is tangent to the subcurve between y and z_1 at a point z_2. So $\angle n_\sigma(x), n_\sigma(z_2) \geqslant \theta$. Because z_2 lies on the subcurve between y and z_1, $z_2 \in B(x,r) \cap \mathrm{mani}(\sigma)$. So by assumption, $\angle n_\sigma(x), n_\sigma(z_2) < \theta$, a contradiction.

Consider (ii). We are done if $\angle n_\sigma(x), yz > \pi/2 - \theta$ for any point $z \in B(x, \frac{1}{3}r) \cap \mathrm{mani}(\sigma)$. Otherwise, because $B(x, \frac{1}{3}r) \cap \mathrm{mani}(\sigma)$ is connected, by continuity, there exists a point $z_1 \in B(x, \frac{1}{3}r) \cap \mathrm{mani}(\sigma)$ such that $\angle n_\sigma(x), yz_1 = \pi/2 - \theta$. Let H be the plane containing yz_1

and parallel to $n_\sigma(x)$. Notice that $B(y, \|y - z_1\|) \subset B(y, \frac{2}{3}r) \subset B(x, r)$. By Lemma 1.2, $H \cap B(y, \frac{2}{3}r) \cap \text{mani}(\sigma)$ is an open curve, and let ξ be its subcurve between y and z_1. The tangent to ξ at some point z_2 is parallel to yz_1 and thus it makes an angle $\pi/2 - \theta$ with the support line of $n_\sigma(x)$. So $\angle n_\sigma(x), n_\sigma(z_2) \geqslant \theta$. But this is a contradiction because $\angle n_\sigma(x), n_\sigma(z_2) < \theta$ by assumption.

B The computation of $d_\omega(x)$, $g(x)$, and $b(x)$

Let x be a point on a 1-face. We discuss the details of computing $d_\omega(x)$, $g(x)$, and $b(x)$.

Computing $d_\omega(x)$ for a non-critical ω. For any 1- or 2-face σ containing x, we compute the ω-deviation radius of x with respect to σ as follows.

Suppose that σ is a 1-face. Let $E_1(z) = 0$ and $E_2(z) = 0$ be the equations of the two given surfaces whose intersection contains σ. For any point $z \in \text{mani}(\sigma)$, $\angle n_\sigma(x), n_\sigma(z) = \omega$ if and only if the angle between the tangents at x and z is ω. Define $G(z) = \nabla E_1(z) \times \nabla E_2(z)$. We solve the following system for z: $E_1(z) = 0$, $E_2(z) = 0$, $\langle G(x), G(z)\rangle = \|G(x)\| \cdot \|G(z)\| \cdot \cos\omega$. Then, we return the distance between x and the closest solution.

Suppose that σ is a 2-face. Let $E_\sigma(z) = 0$ be the equation of mani(σ). Define $G_\sigma(z) = \langle \nabla E_\sigma(x), \nabla E_\sigma(z)\rangle / (\|\nabla E_\sigma(x)\| \cdot \|\nabla E_\sigma(z)\|)$. Then, $\sigma_{x,\omega} = \{z \in \text{mani}(\sigma) : \cos(\angle n_\sigma(x), n_\sigma(z)) = \cos\omega\}$ is described by the system: $E_\sigma(z) = 0$, $G_\sigma(z) = \cos\omega$. If $\nabla G_\sigma(z)$ is not parallel to $\nabla E_\sigma(z)$ for any point $z \in \sigma_{x,\omega}$, $\sigma_{x,\omega}$ is a collection of disjoint smooth closed curves by the implicit function theorem. If $\nabla G_\sigma(z)$ is parallel to $\nabla E_\sigma(z)$ for some point $z \in \sigma_{x,\omega}$, $\cos\omega$ is a *critical value* for the function $G_\sigma(z)$ restricted to mani(σ). We can decide whether $\cos\omega$ and hence ω is critical by testing the solvability of the following system of equations: $E_\sigma(z) = 0, G_\sigma(z) = \cos\omega, \nabla G_\sigma(z) \times \nabla E_\sigma(z) = 0$. If the system is solvable, ω is critical; otherwise, ω is not critical.

Suppose that ω is not critical. For each smooth closed curve in $\sigma_{x,\omega}$, its closest point to x is a tangential contact point with $B(x, r)$ for some $r > 0$. Each contact point z is characterized by the fact that xz is orthogonal to the tangent to $\sigma_{x,\omega}$ at z. Thus, it suffices to find all the tangential contact points and pick the one closest to x. The tangent at a point $z \in \sigma_{x,\omega}$ is parallel to $\nabla G_\sigma(z) \times \nabla E_\sigma(z)$. Thus, we solve the system: $E_\sigma(z) = 0$, $G_\sigma(z) = \cos\omega$, $\langle \nabla G_\sigma(z) \times \nabla E_\sigma(z), (x - z)\rangle = 0$, and return the distance between x and the closest solution.

The minimum distance returned over all faces containing x is $d_\omega(x)$.

Perturbation. The previous computation of $d_\omega(x)$ cannot be performed if ω is found critical with respect to some 2-face σ containing x. We can get around this problem by exploiting the strong version of the Sard's theorem [35].

Since mani(σ) is assumed to be C^3-smooth, the function $G_\sigma(z) = \langle \nabla E_\sigma(x), \nabla E_\sigma(z)\rangle / (\|\nabla E_\sigma(x)\| \cdot \|\nabla E_\sigma(z)\|)$ is C^2-smooth. Under these conditions, the strong version of Sard's theorem states that the critical values of the restriction of $G_\sigma(z)$ to mani(σ) are isolated [35]. This allows us to perturb ω to a non-critical value slightly less than ω. The details are as follows.

We set α to be a positive constant as small as we wish. We check as before whether $\omega - \alpha$ is critical for any 2-face containing x. That is, for any 2-face σ containing x, we check whether the system: $E_\sigma(z) = 0, G_\sigma(z) = \cos(\omega - \alpha), \nabla G_\sigma(z) \times \nabla E_\sigma(z) = 0$ is solvable. If $\omega - \alpha$ is not critical for any 2-face containing x, we set $\bar\omega = \omega - \alpha$ and compute $d_{\bar\omega}(x)$. If $\omega - \alpha$ is critical for some 2-face containing x, we halve α and test the criticality of $\omega - \alpha$ again. By the strong version of the Sard's theorem, we must eventually find a small enough α such that $\omega - \alpha$ is not critical for any 2-face containing x.

Computing $g(x)$. The minimum distance between x and other faces not containing x is $g(x)$. The distances from x to other vertices can easily be computed in linear time. For any 1- or 2-face σ not containing x, we compute the distance between x and $\mathrm{int}\,\sigma$ as follows.

Suppose that σ is a 1-face. Let $E_1(z) = 0$ and $E_2(z) = 0$ be the equations of the two given surfaces Σ_1 and Σ_2 whose intersection contains σ. Any ball centered at x is tangent to a curve in $\Sigma_1 \cap \Sigma_2$ at a point z if and only if xz is orthogonal to the tangent to $\Sigma_1 \cap \Sigma_2$ at z. We find all such contact points z by solving the system: $E_1(z) = 0$, $E_2(z) = 0$, $\langle \nabla E_1(z) \times \nabla E_2(z), (x - z) \rangle = 0$. Then, we return the distance between x and the closest contact point on σ.

Suppose that σ is a 2-face. Let $E(z) = 0$ be the equation of $\mathrm{mani}(\sigma)$. Any ball centered at x is tangent to $\mathrm{mani}(\sigma)$ at a point z if and only if xz is normal to $\mathrm{mani}(\sigma)$. We find all such tangential contact points z by solving the system: $E(z) = 0$, $\nabla E(z) \times (x - z) = 0$. Then, we return the distance between x and the closest contact point on σ.

Computing $b(x)$. For any 1- or 2-face σ containing x, we determine the largest value $b_\sigma(x) > 0$ such that for any $r < b_\sigma(x)$, $B(x, r) \cap \mathrm{mani}(\sigma)$ is a closed ball of dimension $\dim(\sigma)$. This computation is similar to the computation of $g(x)$. The only difference is that we now return the distance between x and the closest contact point instead of the closest contact point on σ. Afterwards, $b(x) = \min_{x \in \sigma} b_\sigma(x)$.

C Proof of Lemma 3.1

In Section C.1 we first prove some properties of the protecting balls constructed. Then, we use these properties to show in Section C.2 that the protecting ball construction procedure terminates correctly. Lemma 3.1(i) is obvious by construction. Lemma 3.1(ii)–(v) are proved in Section C.3–C.6.

Throughout this section we use σ to denote the 1-face that we are processing. We use $u = x_1$ and v to denote the endpoints of σ, and $(x_2, \ldots, x_{m-1}, y_m, x_0)$ to denote the protecting ball centers in $\mathrm{int}\,\sigma$. The quantities r_k and r_{0k} are defined as in Section 3. Although the value m is used in Section C.1, the results in Section C.1 hold independent of whether the protecting ball construction procedure terminates or not. If the procedure were not to terminate, m would be equal to ∞.

We define the *host* of x_k as follows. For $1 \leqslant k \leqslant m$, if $r_k = \lambda \mathrm{f}_{\omega_j}(x_j) + \lambda \|x_j - x_k\|$ for some $0 \leqslant j \leqslant k$, define $\mathrm{host}(x_k) = x_k$; otherwise, define $\mathrm{host}(x_k) = x_b$, where $b = \max\{i : i < k \text{ and } \mathrm{host}(x_i) = x_i\}$. Notice that $\mathrm{host}(x_1) = x_1$, $\mathrm{host}(x_2) = x_2$, and if $\mathrm{host}(x_k) \neq x_k$, $r_k = \|x_{k-1} - x_k\|/2$.

C.1 Technical lemmas

We know that if $\mathrm{host}(x_k) \neq x_k$, $r_k = \|x_{k-1} - x_k\|/2$. The case of $\mathrm{host}(x_k) = x_k$ is discussed below.

Lemma C.1 *$\forall\ 2 \leqslant k \leqslant m$, if* $\mathrm{host}(x_k) = x_k$, $r_k = \min_{0 \leqslant j \leqslant k}\{\lambda \mathrm{f}_{\omega_j}(x_j) + \lambda \|x_j - x_k\|\}$.

Proof. The lemma is true by construction for $k \geqslant 3$. It suffices to show that $\|x_i - x_2\| \geqslant \mathrm{f}_{\omega_2}(x_2)$ for $i \in \{0, 1\}$ so that $\lambda \mathrm{f}_{\omega_2}(x_2) < \min_{0 \leqslant i \leqslant 1}\{\lambda \mathrm{f}_{\omega_i}(x_i) + \lambda \|x_i - x_2\|\}$. It follows from definition that $\|x_1 - x_2\| = \|u - x_2\| \geqslant \mathrm{f}_{\omega_2}(x_2)$. Observe that $\lambda \|u - v\| \geqslant \max\{\lambda g(u), \lambda g(v)\} \geqslant$

$\max\{\lambda f_{\omega_u}(u), \lambda f_{\omega_v}(v)\} = \max\{\|u-x_2\|, \|v-x_0\|\}$. So $\|x_0-x_2\| \geqslant \|u-v\|-\|u-x_2\|-\|v-x_0\| \geqslant (1-2\lambda)\|u-v\| \geqslant (1-2\lambda)f_{\omega_u}(u)$. Since $f_{\omega_2}(x_2) \leqslant \|u-x_2\| = \lambda f_{\omega_u}(u)$, $\|x_0-x_2\| \geqslant (1-2\lambda)f_{\omega_2}(x_2)/\lambda > f_{\omega_2}(x_2)$ as $\lambda \leqslant 0.01$. ◻

We relate radius(B_{x_k}) to radius($B_{\mathrm{host}(x_k)}$) and prove a bound on $\|x_k - \mathrm{host}(x_k)\|$.

Lemma C.2 *$\forall\ 1 \leqslant k \leqslant m$, if* $\mathrm{host}(x_k) = x_b$, $r_k = \left(\frac{3}{5}\right)^{k-b} r_b$ *and* $\|x_b - x_k\| \leqslant 3r_b$.

Proof. The lemma is trivial if $x_b = x_k$. Assume that $x_b \neq x_k$ and so $k \geqslant 3$. For $b < i \leqslant k$, $r_i = \frac{1}{2}\|x_{i-1} - x_i\|$ as $\mathrm{host}(x_i) \neq x_i$. Because $r_{b+1} = \frac{1}{2}\|x_b - x_{b+1}\| = \frac{3}{5}r_b$, a simple induction shows that $r_i = \left(\frac{3}{5}\right)^{i-b} r_b$ for $b < i \leqslant k$. We have $\|x_b - x_k\| \leqslant \sum_{i=b+1}^{k} \|x_{i-1} - x_i\| \leqslant 2r_b \sum_{j=1}^{\infty} \left(\frac{3}{5}\right)^j = 3r_b$. ◻

We show that the radii of protecting balls vary gradually along $\mathrm{int}\,\sigma$.

Lemma C.3

(i) $\forall\ 2 \leqslant k \leqslant m$, $\forall\ 0 \leqslant j \leqslant k$, $r_k \leqslant r_j + \lambda\|x_j - x_k\|$.

(ii) $\forall\ 2 \leqslant k < m$, $r_k = \frac{5}{6}\|x_k - x_{k+1}\|$ *and* $r_{k+1} \leqslant \frac{5+6\lambda}{6}\|x_k - x_{k+1}\| = \frac{5+6\lambda}{5} r_k$.

(iii) $\forall\ 2 \leqslant k \leqslant m$, $r_{0k} \geqslant r_k - \lambda\|x_k - x_0\|$.

Proof. Consider (i). For $0 \leqslant j \leqslant k$, by definition, $r_j \geqslant \lambda f_{\omega_{i_j}}(x_{i_j}) + \lambda\|x_j - x_{i_j}\|$ for some $i_j \leqslant j$. If $\mathrm{host}(x_k) = x_k$, by Lemma C.1, $r_k \leqslant \lambda f_{\omega_{i_j}}(x_{i_j}) + \lambda\|x_k - x_{i_j}\| \leqslant \lambda f_{\omega_{i_j}}(x_{i_j}) + \lambda\|x_j - x_{i_j}\| + \lambda\|x_j - x_k\| \leqslant r_j + \lambda\|x_j - x_k\|$. Suppose that $\mathrm{host}(x_k) = x_b$ for some $b < k$. For $b \leqslant j < k$, since $\mathrm{host}(x_j) = x_b$, Lemma C.2 implies that $r_k = \left(\frac{3}{5}\right)^{k-b} r_b < \left(\frac{3}{5}\right)^{j-b} r_b = r_j$. For $0 \leqslant j < b$, $r_j + \lambda\|x_j - x_k\| \geqslant r_j + \lambda\|x_j - x_b\| - \lambda\|x_b - x_k\|$. We have shown that $r_b \leqslant r_j + \lambda\|x_j - x_b\|$ as $\mathrm{host}(x_b) = x_b$. Also, $\|x_b - x_k\| \leqslant 3r_b$ by Lemma C.2. Thus, $r_j + \lambda\|x_j - x_k\| \geqslant (1-3\lambda)r_b = (1-3\lambda)5^{k-b}r_k/3^{k-b} \geqslant r_k$ as $\lambda \leqslant 0.01$. This proves (i). For $2 \leqslant k < m$, $r_k = \frac{5}{6}\|x_k - x_{k+1}\|$ by construction, and by (i), $r_{k+1} \leqslant r_k + \lambda\|x_k - x_{k+1}\| = \frac{5+6\lambda}{6}\|x_k - x_{k+1}\| = \frac{5+6\lambda}{5} r_k$. This proves (ii). By definition, there exists $0 \leqslant j \leqslant k$ such that $r_{0k} = r_j + \lambda\|x_j - x_0\|$, so by (i), $r_{0k} \geqslant r_k - \lambda\|x_j - x_k\| + \lambda\|x_j - x_0\| \geqslant r_k - \lambda\|x_k - x_0\|$. This proves (iii). ◻

C.2 Proof of termination

The next lemma shows the correctness of our strategy of defining x_k for $3 \leqslant k \leqslant m$.

Lemma C.4 *For $3 \leqslant k \leqslant m$, $B(x_{k-1}, \frac{6}{5}r_{k-1}) \cap \sigma$ is an open curve and x_k is its only endpoint that satisfies $\angle x_{k-2}x_{k-1}x_k > \pi/2$. Also, x_k lies between x_{k-1} and v along σ.*

Proof. Let $x_b = \mathrm{host}(x_{k-2})$. Choose any number r from the range $(7\lambda f_{\omega_b}(x_b), f_{\omega_b}(x_b))$. By Lemma 2.5(ii)(b), $B(x_b, r) \cap \sigma$ is an open curve. Moreover, since $k \geqslant 3$, we have $x_b \in \mathrm{int}\,\sigma$. Then, as $f_{\omega_b}(x_b) \leqslant g(x_b)$ by definition, the endpoints of $B(x_b, r) \cap \sigma$ lie on the boundary of $B(x_b, r)$.

We first show that $B(x_{k-1}, \frac{6}{5}r_{k-1}) \subset B(x_b, r)$. If $k = 3$, $x_b = x_1$ and $\|x_1 - x_2\| + \frac{6}{5}r_2 \leqslant r_1 + \frac{6}{5}r_1 < 3r_1 = 3\lambda f_{\omega_1}(x_1) < r$. If $k > 3$, Lemma C.2 and Lemma C.3(ii) imply that

$\|x_b - x_{k-2}\| + \frac{6}{5}r_{k-2} + \frac{6}{5}r_{k-1} < 7r_b \leqslant 7\lambda \mathrm{f}_{\omega_b}(x_b) < r$. We conclude that $\|x_b - x_{k-1}\| + \frac{6}{5}r_{k-1} < r$ which implies that $B(x_{k-1}, \frac{6}{5}r_{k-1}) \subset B(x_b, r)$.

Next, we show that $B(x_{k-1}, \frac{6}{5}r_{k-1}) \cap \sigma$ is an open curve. Suppose not. Then, we can walk from x_{k-1} along $B(x_b, r) \cap \sigma$ so that we leave $B(x_{k-1}, \frac{6}{5}r_{k-1})$ at a point y and reenter $B(x_{k-1}, \frac{6}{5}r_{k-1})$ later. Thus, after leaving $B(x_{k-1}, \frac{6}{5}r_{k-1})$ at y, we must return to another point y' on the plane H tangent to $B(x_{k-1}, \frac{6}{5}r_{k-1})$ at y. We can translate H to a plane that is tangent to the subcurve between y and y' at a point z. Notice that $x_{k-1}y \perp H$ which implies that $x_{k-1}y$ is parallel to $n_\sigma(z)$. But this is a contradiction because, by applying Lemma 2.5(ii)(a) and (iii) to $B(x_b, r) \cap \sigma$, we get $\angle n_\sigma(z), x_{k-1}y \geqslant \angle n_\sigma(x_b), x_{k-1}y - \angle n_\sigma(x_b), n_\sigma(z) > \pi/2 - 2\omega$.

Since $B(x_{k-1}, \frac{6}{5}r_{k-1}) \cap \sigma$ is an open curve, it has one endpoint x_k between x_{k-1} and v and one endpoint z' between u and x_{k-1}. We have $\angle x_{k-2}x_{k-1}x_k \geqslant \angle n_\sigma(x_{k-1}), x_{k-1}x_{k-2} + \angle n_\sigma(x_{k-1}), x_{k-1}x_k \geqslant \angle n_\sigma(x_b), x_{k-1}x_{k-2} + \angle n_\sigma(x_b), x_{k-1}x_k - 2 \cdot \angle n_\sigma(x_b), n_\sigma(x_{k-1})$. Thus, by Lemma 2.5(ii)(a) and (iii), we conclude that $\angle x_{k-1}x_{k-1}x_k > \pi - 4\omega$ which is greater than $\pi/2$. On the other hand, $\angle x_{k-2}x_{k-1}z' \leqslant \pi - \angle n_\sigma(x_{k-1}), x_{k-1}x_{k-2} - \angle n_\sigma(x_{k-1}), x_{k-1}z' \leqslant \pi - \angle n_\sigma(x_b), x_{k-1}x_{k-2} - \angle n_\sigma(x_b), x_{k-1}z' + 2 \cdot \angle n_\sigma(x_b), n_\sigma(x_{k-1}) < 4\omega$, which is less than $\pi/2$. □

We prove in Lemma C.5 below that the protecting ball construction procedure terminates correctly. Properties (i)-(iii) will be needed in some subsequent proofs.

Lemma C.5

(i) $\frac{2}{1+\lambda}r_{m-1} < \|x_{m-1} - x_0\| < 5r_{m-1}$.

(ii) $\max\{\|x_{m-1} - y_m\|, \|y_m - x_0\|\} < \min\{4r_{0\,m-1}, 5r_{m-1}\}$.

(iii) $0.09r_{m-1} \leqslant 0.1r_{0\,m-1} \leqslant \mathrm{radius}(B_{y_m}) < \min\{5\lambda g(y_m), 4r_{0\,m-1}, 5r_{m-1}\}$.

(iv) *The protecting ball construction procedure terminates correctly.*

Proof. We claim that $B_{x_2} \cap B(x_0, r_0) = \emptyset$. By Lemma 2.5(i), $r_0 = \lambda \mathrm{f}_{\omega_0}(x_0) \leqslant \lambda g(x_0) \leqslant \lambda\|v - x_0\|$. Similarly, $r_2 \leqslant \lambda\|u - x_2\|$ and $\lambda\|u - v\| \geqslant \max\{\lambda g(u), \lambda g(v)\} \geqslant \max\{\lambda \mathrm{f}_{\omega_u}(u), \lambda \mathrm{f}_{\omega_v}(v)\} = \max\{\|u - x_2\|, \|v - x_0\|\}$. Thus, $\|u - x_2\| + r_2 + r_0 + \|v - x_0\| \leqslant 2\lambda(1 + \lambda)\|u - v\| < \|u - v\|$. This implies our claim.

Since $r_{02} \leqslant r_0$, by our claim, the protecting ball construction starts with the right condition that $B_{x_2} \cap B(x_0, r_{02}) = \emptyset$. By Lemma C.4, our procedure correctly defines x_k for $3 \leqslant k \leqslant m$. For $2 \leqslant k < m$, since $B_{x_k} \cap B(x_0, r_{0k}) = \emptyset$, $\|x_k - x_0\| > r_k + r_{0k} \geqslant 2r_k - \lambda\|x_k - x_0\|$ by Lemma C.3(iii). Thus, $\|x_k - x_0\| > \frac{2}{1+\lambda}r_k > \frac{6}{5}r_k$. It follows that the protecting ball construction procedure never advances beyond x_0. This also proves the first inequality in (i) by substituting $k = m - 1$. For $k \geqslant 3$, both the distance advanced along σ by B_{x_k} and r_{0k} are bounded away from zero. Thus, we must eventually place the ball $B(x_m, r_m)$ that overlaps with $B(x_0, r_{0m})$.

We claim that

$$\|x_{m-1} - x_0\| < 4r_{0\,m-1} \leqslant 4r_0 \tag{2}$$

By Lemma C.3(iii), $r_{m-1} \leqslant r_{0\,m-1} + \lambda\|x_{m-1} - x_0\|$. Then, by Lemma C.3(i) and (ii), $r_m \leqslant r_{m-1} + \lambda\|x_{m-1} - x_m\| = \frac{5+6\lambda}{5}r_{m-1} \leqslant \frac{5+6\lambda}{5}r_{0\,m-1} + \frac{5\lambda+6\lambda^2}{5}\|x_{m-1} - x_0\|$. Substituting these inequalities into $\|x_{m-1} - x_0\| \leqslant \frac{6}{5}r_{m-1} + r_m + r_{0m}$ (recall that $B(x_m, r_m)$ overlaps with $B(x_0, r_{0m})$), we get $\|x_{m-1} - x_0\| \leqslant \frac{11+6\lambda}{5}r_{0\,m-1} + r_{0m} + \frac{11\lambda+6\lambda^2}{5}\|x_{m-1} - x_0\|$. Since $r_{0m} \leqslant r_{0\,m-1}$, we get $\|x_{m-1} - x_0\| \leqslant \frac{16+6\lambda}{5}r_{0\,m-1} + \frac{11\lambda+6\lambda^2}{5}\|x_{m-1} - x_0\| \Rightarrow \|x_{m-1} - x_0\| \leqslant \frac{16+6\lambda}{5-11\lambda-6\lambda^2}r_{0\,m-1} < 4r_{0\,m-1} \leqslant 4r_0$. This proves equation (2).

By equation (2) and the definition of $r_{0\,m-1}$, $\|x_{m-1} - x_0\| < 4r_{0\,m-1} \leqslant 4r_{m-1} + 4\lambda\|x_{m-1} - x_0\| \Rightarrow \|x_{m-1} - x_0\| < \frac{4}{1-4\lambda} r_{m-1} < 5r_{m-1}$. This proves the second inequality in (i).

Since $4r_0 = 4\lambda \mathrm{f}_{\omega_0}(x_0) < \mathrm{f}_{\omega_0}(x_0)$, by Lemma 2.5(ii), $B(x_0, 4r_0) \cap \sigma$ is an open curve and for any point $y \in B(x_0, 4r_0) \cap \sigma$, $\angle n_\sigma(x_0), n_\sigma(y) < \omega$. By Lemma 2.5(iii), $\angle n_\sigma(x_0), x_0x_{m-1} > \pi/2 - \omega$, which implies that $n_\sigma(x_0)$ makes an angle less than ω with the bisector plane of B_{x_0} and $B_{x_{m-1}}$. It follows that $B(x_0, 4r_0) \cap \sigma$ contains the subcurve between x_{m-1} and x_0, which intersects the bisector plane of $B_{x_{m-1}}$ and B_{x_0} exactly once. This intersection point is y_m and $\|y_m - x_0\| \leqslant 4r_0$. This proves that the procedure terminates correctly, i.e., (iv) is true.

Since $B(x_0, 4r_0) \cap \sigma$ contains the subcurve between x_{m-1} and x_0, by Lemma 2.5(ii)(a) and (iii), $\angle x_{m-1}y_mx_0 \geqslant \angle n_\sigma(y_m), y_mx_{m-1} + \angle n_\sigma(y_m), y_mx_0 > \angle n_\sigma(x_0), y_mx_{m-1} + \angle n_\sigma(x_0), y_mx_0 - 2 \cdot \angle n_\sigma(x_0), n_\sigma(y_m) > \pi - 4\omega > \pi/2$. Thus, $\max\{\|x_{m-1} - y_m\|, \|y_m - x_0\|\} < \|x_{m-1} - x_0\|$. Then (ii) follows from (i) and equation (2).

Since $B_{x_{m-1}} \cap B(x_0, r_{0\,m-1}) = \emptyset$ and $\mathrm{radius}(B_{x_0}) = \frac{4}{5} r_{0\,m-1}$, the distance between $B_{x_{m-1}}$ and B_{x_0} is at least $\frac{1}{5} r_{0\,m-1}$. Because B_{y_m} covers the gap between $B_{x_{m-1}}$ and B_{x_0}, we get $\mathrm{radius}(B_{y_m}) \geqslant 0.1 r_{0\,m-1}$. By (i) and Lemma C.3(iii), $r_{0\,m-1} \geqslant r_{m-1} - \lambda\|x_{m-1} - x_0\| \geqslant r_{m-1} - 5\lambda r_{m-1} \geqslant 0.9 r_{m-1}$. Thus, $\mathrm{radius}(B_{y_m}) \geqslant 0.1 r_{0\,m-1} \geqslant 0.09 r_{m-1}$. This proves the first two inequalities in (iii).

Since x_0 lies outside B_{y_m}, by (ii), $\mathrm{radius}(B_{y_m}) < \|y_m - x_0\| < \min\{4r_{0\,m-1}, 5r_{m-1}\}$. This proves part of the third inequality in (iii). Starting with $\|y_m - x_0\| < 4r_{0\,m-1} \leqslant 4r_0$, we get $\|y_m - x_0\| < 4r_0 = 4\lambda \mathrm{f}_{\omega_0}(x_0) \leqslant 4\lambda g(x_0)$. Since g is 1-Lipschitz over $\mathrm{int}\,\sigma$ by Lemma 2.2, $\|y_m - x_0\| < 4\lambda g(y_m) + 4\lambda\|y_m - x_0\| \Rightarrow \mathrm{radius}(B_{y_m}) < \|y_m - x_0\| < \frac{4\lambda}{1-4\lambda} g(y_m) < 5\lambda g(y_m)$. ∎

C.3 Lemma 3.1(ii)

The next result implies that $u = x_1$ does not influence the setting of r_k for $3 \leqslant k < m$ and $r_{0\,m-1}$.

Lemma C.6 *$\forall\, 3 \leqslant k < m$, $\lambda \mathrm{f}_{\omega_1}(x_1) + \lambda\|x_1 - x_k\| > \lambda \mathrm{f}_{\omega_2}(x_2) + \lambda\|x_2 - x_k\|$ and $r_1 + \lambda\|x_1 - x_0\| > r_2 + \lambda\|x_2 - x_0\|$.*

Proof. Because $\lambda \leqslant 0.01$ and $\mathrm{f}_{\omega_2}(x_2) \leqslant \|x_1 - x_2\|$, $\mathrm{f}_{\omega_1}(x_1) = \frac{1}{\lambda}\|x_1 - x_2\| > \mathrm{f}_{\omega_2}(x_2) + \|x_1 - x_2\|$. It follows immediately that $\lambda \mathrm{f}_{\omega_1}(x_1) + \lambda\|x_1 - x_k\| > \lambda \mathrm{f}_{\omega_2}(x_2) + \lambda\|x_2 - x_k\|$ and $r_1 + \lambda\|x_1 - x_0\| = \lambda \mathrm{f}_{\omega_1}(x_1) + \lambda\|x_1 - x_0\| > \lambda \mathrm{f}_{\omega_2}(x_2) + \lambda\|x_1 - x_2\| + \lambda\|x_1 - x_0\| \geqslant r_2 + \lambda\|x_2 - x_0\|$. ∎

We are ready to prove Lemma 3.1(ii).

Proof of Lemma 3.1(ii). For any $p \in \{u, x_2, v\}$, since $\mathrm{radius}(B_p) = \lambda \mathrm{f}_{\bar\omega}(p)$, where $\bar\omega = \omega_u$, ω_2 or ω_v depending on the identity of p, by Lemma 2.5(i), $\mathrm{radius}(\mathcal{B}_p) \in [\,\frac{1}{2}\lambda\bar\omega f(p)\,,\, \lambda g(p)\,] \subseteq [\,\frac{99}{200}\lambda\omega f(p)\,,\, \lambda g(p)\,]$ because $\bar\omega \in [0.99\omega, \omega]$.

Consider B_{x_k} for any $3 \leqslant k < m$. By definition, there exists $0 \leqslant j \leqslant k$ such that $r_k \geqslant \lambda \mathrm{f}_{\omega_j}(x_j) + \lambda\|x_j - x_k\|$ which is at least $\frac{1}{2}\lambda\omega_j f(x_j) + \lambda\|x_j - x_k\| \geqslant \frac{99}{200}\lambda\omega f(x_j) + \lambda\|x_j - x_k\|$ by Lemma 2.5(i) and the fact that $\omega_j \geqslant 0.99\omega$. We can choose $j > 1$ by Lemma C.6 and so $x_j, x_k \in \mathrm{int}\,\sigma$. Because f is 1-Lipschitz over $\mathrm{int}\,\sigma$ by Lemma 2.2, $r_k \geqslant \frac{99}{200}\lambda\omega f(x_j) + \lambda\|x_j - x_k\| > \frac{99}{200}\lambda\omega f(x_k)$. This proves the lower bound. Let $x_b = \mathrm{host}(x_k)$. Since $k \geqslant 3$, $b \geqslant 2$ and so $x_b, x_k \in \mathrm{int}\,\sigma$. By Lemma C.2 and Lemma 2.5(i), $r_k = \left(\frac{3}{5}\right)^{k-b} r_b \leqslant \left(\frac{3}{5}\right)^{k-b} \lambda \mathrm{f}_{\omega_b}(x_b) \leqslant \left(\frac{3}{5}\right)^{k-b} \lambda g(x_b)$. Since g is 1-Lipschitz over $\mathrm{int}\,\sigma$ by Lemma 2.2, $r_k \leqslant \left(\frac{3}{5}\right)^{k-b} \lambda g(x_k) + \left(\frac{3}{5}\right)^{k-b} \lambda\|x_b - x_k\|$. By

Lemma C.2, $\left(\frac{3}{5}\right)^{k-b}\lambda\|x_b - x_k\| \leqslant \left(\frac{3}{5}\right)^{k-b} 3\lambda r_b = 3\lambda r_k$. This yields $r_k \leqslant \left(\frac{3}{5}\right)^{k-b}\frac{\lambda}{1-3\lambda}g(x_k) < 2\lambda g(x_k)$ as $\lambda \leqslant 0.01$.

Consider B_{x_0}. We have $\mathrm{radius}(B_{x_0}) = 0.8r_{0\,m-1} \leqslant 0.8r_0 = 0.8\lambda \mathrm{f}_{\omega_0}(x_0) \leqslant 0.8\lambda g(x_0)$. By definition and Lemma C.6, $r_{0\,m-1} = r_j + \lambda\|x_j - x_0\|$ for some $j \leqslant m-1$ and $x_j \in \mathrm{int}\,\sigma$. We already know that $r_j \geqslant \frac{99}{200}\lambda\omega f(x_j)$. So $0.8r_{0\,m-1} \geqslant 0.8 \cdot (\frac{99}{200}\lambda\omega f(x_j) + \lambda\|x_j - x_0\|) > 0.8 \cdot \frac{99}{200}\lambda\omega f(x_0) = \frac{99}{250}\lambda\omega f(x_0)$ because f is 1-Lipschitz over $\mathrm{int}\,\sigma$.

Consider B_{y_m}. By Lemma C.5(iii), $\mathrm{radius}(B_{y_m}) < 5\lambda g(y_m)$. By Lemma C.5(ii), $f(y_m) \leqslant f(x_0) + \|y_m - x_0\| < f(x_0) + 4r_{0\,m-1}$. We have shown in the previous paragraph that $0.8r_{0\,m-1} > 0.8 \cdot \frac{99}{200}\lambda\omega f(x_0)$. So $f(y_m) < \frac{200}{99\lambda\omega}r_{0\,m-1} + 4r_{0\,m-1} = \frac{200+396\lambda\omega}{99\lambda\omega}r_{0\,m-1}$. Then, Lemma C.5(iii) implies that $\mathrm{radius}(B_{y_m}) \geqslant 0.1r_{0\,m-1} > \frac{99\lambda\omega}{10(200+396\lambda\omega)}f(y_m) > \frac{\lambda\omega}{21}f(y_m)$ as $\lambda \leqslant 0.01$ and $\omega \leqslant 0.01$.

C.4 Lemma 3.1(iii)

The next lemma shows that a ball $B(p,r)$ satisfies Lemma 3.1(iii) assuming that certain conditions are satisfied.

Lemma C.7 *Let p and p' be two points point on the 1-face σ. Let $r' < \mathrm{f}_\omega(p')$. Let E be the set consisting of σ and the 2-faces incident to σ. For any $r < g(p)$ such that $B(p,r) \subset B(p', \frac{1}{3}r')$, the following results hold.*

(i) *For any $\tau \in E$ and for any point $z \in B(p,r) \cap \mathrm{mani}(\tau)$, $\angle n_\tau(p), n_\tau(z) < 2\omega$.*

(ii) *For any $\tau \in E$, $B(p,r) \cap \mathrm{mani}(\tau)$ and $B(p,r) \cap \tau$ are closed balls of dimension* $\dim(\tau)$.

Proof. Since $r' < \mathrm{f}_\omega(p')$ and $B(p,r) \subset B(p', \frac{1}{3}r')$, for any point $z \in B(p,r) \cap \mathrm{mani}(\tau)$, $\angle n_\tau(p), n_\tau(z) \leqslant \angle n_\tau(p'), n_\tau(p) + \angle n_\tau(p'), n_\tau(z) < 2\omega$ by Lemma 2.5(ii)(a). This proves (i).

Take any $\tau \in E$. Assume to the contrary that $B(p,r) \cap \mathrm{mani}(\tau)$ is not a closed ball of dimension $\dim(\tau)$. Then, there exists a concentric ball $B \subset B(p,r)$ such that B is tangent to $B(p,r) \cap \mathrm{mani}(\tau)$ at a point q. We have $\angle n_\tau(p), pq \leqslant \angle n_\tau(p), n_\tau(q) < 2\omega$ by (i). (Note that $\angle n_\tau(p), pq = \angle n_\tau(p), n_\tau(q)$ if τ is a 2-face and $\angle n_\sigma(p), pq \leqslant \angle n_\sigma(p), n_\sigma(q)$.) However, as $p, q \in B(p', \frac{1}{3}r') \cap \mathrm{mani}(\tau)$ and $r' < \mathrm{f}_\omega(p')$, by Lemma 2.5(ii)(a), (iii) and (iv), $\angle n_\tau(p), pq \geqslant \angle n_\tau(p'), pq - \angle n_\tau(p'), n_\tau(p) > \pi/2 - 2\omega$, a contradiction.

Because $B(p,r) \cap \mathrm{mani}(\tau)$ is a closed ball and $r < g(p)$ by assumption, we can invoke Lemma 1.1 to conclude that $B(p,r) \cap \tau$ is a closed ball of dimension $\dim(\tau)$. ◻

The next two lemmas show that the conditions of Lemma C.7 can be satisfied for y_m and x_k for $3 \leqslant k < m$.

Lemma C.8 *Let $r \leqslant 20 \cdot \mathrm{radius}(B_{y_m})$. Let $r' = 0.9\mathrm{f}_{\omega_v}(v)$. Then $r < g(y_m)$ and $B(y_m, r) \subset B(v, \frac{1}{3}r')$. Hence, $B(y_m, r)$ and $B(v, r')$ satisfy the conditions of Lemma C.7.*

Proof. It follows from Lemma C.5(iii) that $r < g(y_m)$: $r \leqslant 20\,\mathrm{radius}(B_{y_m}) < 100\lambda g(y_m) \leqslant g(y_m)$ as $\lambda \leqslant 0.01$. We have

$$r_{0\,m-1} \leqslant r_0 = \lambda \mathrm{f}_{\omega_0}(x_0) \leqslant \lambda\|x_0 - v\| = \lambda^2 \mathrm{f}_{\omega_v}(v). \tag{3}$$

By Lemma C.5(iii) and (3), $r \leqslant 20\,\mathrm{radius}(B_{y_m}) < 80r_{0\,m-1} \leqslant 80\lambda^2\mathrm{f}_{\omega_v}(v)$. We have $\|y_m - v\| \leqslant \|x_0 - v\| + \|y_m - x_0\| = \lambda\mathrm{f}_{\omega_v}(v) + \|y_m - x_0\|$ which is less than $(\lambda + 4\lambda^2)\mathrm{f}_{\omega_v}(v)$ by Lemma C.5(ii) and (3). Therefore, $\|y_m - v\| + r < (\lambda + 84\lambda^2)\mathrm{f}_{\omega_v}(v) < 0.3\mathrm{f}_{\omega_v}(v) = r'/3$ as $\lambda \leqslant 0.01$. So $B(y_m, r) \subset B(v, \frac{1}{3}r')$. ◻

Lemma C.9 *Let $3 \leqslant k < m$. Let $x_b = \mathrm{host}(x_k)$. Let $r \leqslant 20 \cdot \mathrm{radius}(B_{x_k})$. Let $r' = 0.9\mathrm{f}_{\omega_b}(x_b)$. Then $r < g(x_k)$ and $B(x_k, r) \subset B(x_b, \frac{1}{3}r')$. Hence, $B(x_k, r)$ and $B(x_b, r')$ satisfy the conditions of Lemma C.7.*

Proof. Recall that $r_k = \mathrm{radius}(B_{x_k})$ for $3 \leqslant k < m$. By Lemma 3.1(ii), $r \leqslant 20r_k < 100\lambda g(x_k) \leqslant g(x_k)$ as $\lambda \leqslant 0.01$. We have $\|x_b - x_k\| + r \leqslant \|x_b - x_k\| + 20r_k$ which is at most $3r_b + 20r_b = 23r_b$ by Lemma C.2. Since $23r_b \leqslant 23\lambda\mathrm{f}_{\omega_b}(x_b) < 0.3\mathrm{f}_{\omega_b}(x_b)$ as $\lambda \leqslant 0.01$, we conclude that $B(x_k, r) \subset B(x_b, 0.3\mathrm{f}_{\omega_b}(x_b)) = B(x_b, \frac{1}{3}r')$. ◻

We are ready to prove Lemma 3.1(iii).

Proof of Lemma 3.1(iii). Take any $p \in \{u, x_0, x_2, v\}$. For any $c \leqslant 20$, since $c \cdot \mathrm{radius}(B_p) \leqslant c\lambda\mathrm{f}_\omega(p) < \mathrm{f}_\omega(p)$, Lemma 2.5(ii) implies that Lemma 3.1(iii) holds for p. By Lemma C.8 and Lemma C.7, Lemma 3.1(iii) holds for y_m. Similarly, by Lemma C.9 and Lemma C.7, Lemma 3.1(iii) holds for x_k for $3 \leqslant k < m$.

C.5 Lemma 3.1(iv)

C.5.1 Proof of Lemma 3.1(iv)(a)

By construction, x_0 and x_2 lie on the boundaries of B_v and B_u, respectively. Since $\mathrm{radius}(B_{x_0}) \leqslant r_0 = \lambda\mathrm{f}_{\omega_0}(x_0) < g(x_0) \leqslant \|x_0 - v\|$, v lies outside B_{x_0}. Similarly, u lies outside B_{x_2}.

For $3 \leqslant k < m$, x_k lies on the boundary of $B(x_{k-1}, \frac{6}{5}r_{k-1})$ and hence outside $B_{x_{k-1}}$. By Lemma C.3(ii), for $3 \leqslant k < m$, $\|x_{k-1} - x_k\| \geqslant \frac{6}{5+6\lambda}r_k > r_k$ as $\lambda \leqslant 0.01$. So x_{k-1} lies outside B_{x_k}.

Since $B_{x_{m-1}} \cap B_{x_0} = \emptyset$, they do not intersect their bisector plane. It follows that y_m lies outside $B_{x_{m-1}}$ and B_{x_0}. Since B_{y_m} is orthogonal to $B_{x_{m-1}}$ and B_{x_0}, x_{m-1} and x_0 lie outside B_{y_m}.

C.5.2 Proof of Lemma 3.1(iv)(b) and Lemma 3.1(iv)(c)

We first show two lemmas.

Lemma C.10 *Let $B(a, r)$ and $B(a', r')$ be two protecting balls such that $a, a' \in \sigma$ and $B(a', r') \subset B(a, 4r)$. Let $\sigma(a, a')$ denote the subcurve between a and a'. For any point $z \in \sigma(a, a')$, $\angle zaa' < 4\omega$ and $\angle za'a < 8\omega$.*

Proof. Note that $B(a, 4r) \cap \sigma$ is an open curve by Lemma 3.1(iii)(b). As $a' \in B(a, 4r)$ by assumption, it follows that $\sigma(a, a') \subset B(a, 4r) \cap \sigma$. We apply Lemma 3.1(iii) and then Lemma 1.3(i) to $B(a, 4r) \cap \sigma$. This allows us to conclude that:

$$\angle n_\sigma(a), aa' > \pi/2 - 2\omega, \tag{4}$$

$$\forall\, z \in \sigma(a, a'),\ \angle n_\sigma(a), az > \pi/2 - 2\omega, \tag{5}$$

$$\angle n_\sigma(a'), aa' \geqslant \angle n_\sigma(a), aa' - \angle n_\sigma(a), n_\sigma(a') > \pi/2 - 4\omega, \tag{6}$$

$$\forall\, z \in \sigma(a, a'),\ \angle n_\sigma(a'), a'z \geqslant \angle n_\sigma(a), a'z - \angle n_\sigma(a), n_\sigma(a') > \pi/2 - 4\omega. \tag{7}$$

For any point $z \in \sigma(a, a')$, (4) and (5) imply that $\angle zaa' < 4\omega$; (6) and (7) imply that $\angle za'a < 8\omega$. ◻

Lemma C.11 *Let $B(a,r)$ and $B(a',r')$ be two consecutive protecting balls such that $a, a' \in \operatorname{int}\sigma$ and $r \geqslant r'$. For any point b on the circle $\operatorname{bd} B(a,r) \cap \operatorname{bd} B(a',r')$, $\angle baa' \geqslant \arctan(0.09) > 8\omega$.*

Proof. Consider the pairs of adjacent weighted points $\{(x_k, x_{k+1}) : 2 \leqslant k \leqslant m-2\}$. We have $\|a-a'\| = \|x_k - x_{k+1}\| = \frac{6}{5}r_k$. By construction, $r_{k+1} \geqslant \frac{1}{2}\|x_k - x_{k+1}\| = \frac{3}{5}r_k$. By Lemma C.3(ii), $r_{k+1} \leqslant \frac{5+6\lambda}{5}r_k < \frac{6}{5}r_k = \|a-a'\|$. This implies that

$$r' \geqslant 3r/5, \tag{8}$$

$$\|a-a'\| \leqslant 6r/5. \tag{9}$$

Drop a perpendicular from b to the point b' on aa'. If $\angle ba'a \geqslant \pi/4$, by (8), $\|b-b'\| = r'\sin\angle ba'a \geqslant 3r/(5\sqrt{2})$, which implies that $\angle baa' = \arcsin(\|b-b'\|/r) \geqslant \arcsin(3/(5\sqrt{2})) > \arctan(0.09)$. If $\angle ba'a < \pi/4$, $\|a-b'\| = \|a-a'\| - r'\cos\angle ba'a$. So by (8) and (9), $\|a-b'\| \leqslant 6r/5 - 3r/(5\sqrt{2}) < 0.8r$. Thus, $\angle baa' = \arccos(\|a-b'\|/r) > \arccos(0.8) > \arctan(0.09)$.

Consider the pair (x_{m-1}, y_m). By construction, $\operatorname{radius}(B_{x_{m-1}}) = r_{m-1}$. By Lemma C.5(iii), $0.09r_{m-1} \leqslant \operatorname{radius}(B_{y_m}) < 5r_{m-1}$. This implies that $r' \geqslant 0.09r$. Since $B_{x_{m-1}}$ and B_{y_m} are orthogonal, $\angle aba' = \pi/2$. So $\angle baa' = \arctan(r'/r) \geqslant \arctan(0.09)$.

Consider the pair (y_m, x_0). By construction, $\operatorname{radius}(B_{x_0}) = 0.8r_{0\,m-1}$. By Lemma C.5(iii), $0.1r_{0\,m-1} \leqslant \operatorname{radius}(B_{y_m}) < 4r_{0\,m-1}$. This implies that $r' \geqslant 0.125r$. Since B_{y_m} and B_{x_0} are orthogonal, $\angle aba' = \pi/2$. So $\angle baa' = \arctan(r'/r) \geqslant \arctan(0.125)$. ∎

Proof of Lemma 3.1(iv)(b). Let B_p and B_q be two consecutive protecting balls. Assume that $\operatorname{radius}(B_p) \geqslant \operatorname{radius}(B_q)$. Take any point b on the circle $\operatorname{bd} B_p \cap \operatorname{bd} B_q$. We have $\angle bqp \geqslant \angle bpq$ as $\operatorname{radius}(B_p) \geqslant \operatorname{radius}(B_q)$. If $p, q \in \operatorname{int}\sigma$, $\angle bpq \geqslant \arctan(0.09)$ by Lemma C.11 and so $\angle pbq \leqslant \pi - 2\arctan(0.09) < 170°$. The remaining case is that p is an endpoint of σ as $\operatorname{radius}(B_p) \geqslant \operatorname{radius}(B_q)$ by assumption. That is, $p = u$ and $q = x_2$, or $p = v$ and $q = x_0$. In this case since $q \in \operatorname{bd} B_p$ and $\operatorname{radius}(B_q) \leqslant \operatorname{radius}(B_p)$, $\angle pbq < \pi/2$.

Proof of Lemma 3.1(iv)(c). Let B_p and B_q be two consecutive protecting balls. Let $\sigma(p,q)$ denote the subcurve of σ between p and q. We show that $\sigma(p,q) \subset B_p \cup B_q$ by examining the possibilities for p and q.

Consider the pairs $\{(u, x_2), (x_0, v)\}$. Because $x_2 \in \operatorname{bd} B_u$ and $B_u \cap \sigma$ is an open curve by Lemma 3.1(iii)(b), $B_u \cap \sigma$ is exactly the subcurve between u and x_2. Similarly, $B_v \cap \sigma$ is exactly the subcurve between x_0 and v.

The remaining possible pairs are adjacent weighted points in $\operatorname{int}\sigma$. Let $B(a,r)$ denote the larger of B_p and B_q. Let $B(a',r')$ denote the smaller one. Notice that both B_p and B_q are contained in $B(a,3r)$. By Lemma C.10, $\sigma(p,q)$ stays inside the cone with apex a, axis aa' and angular aperture 8ω. Thus, Lemma C.11 implies that $\sigma(p,q)$ enters $B(a',r')$ before leaving $B(a,r)$. After $\sigma(p,q)$ enters $B(a',r')$, $\sigma(p,q)$ cannot leave $B(a',r')$ and reenter $B(a',r')$ later because $B(a',r') \cap \sigma$ is an open curve by Lemma 3.1(iii)(b). Hence, $\sigma(p,q) \subset B_p \cup B_q$.

C.5.3 Proof of Lemma 3.1(iv)(d)

Let z be a point on the subcurve $\sigma(p,q)$. If $z = p$ or q, by Lemma 3.1(ii), the ball centered at z with radius $\frac{\lambda\omega}{2500}f(z)$ lies within B_z and we are done. Assume that $z \in \operatorname{int}\sigma(p,q)$ in the following.

Let $B(a,r)$ denote the larger of B_p and B_q. Let $B(a',r')$ denote the smaller one. Let H be the plane containing the circle $\mathrm{bd}\, B_p \cap \mathrm{bd}\, B_q$. Let C be the circle in H concentric to $\mathrm{bd}\, B_p \cap \mathrm{bd}\, B_q$ such that the the angular aperture of the cone K formed by connecting a to C is 16ω. Since $r \geqslant r'$, the angular aperture of the cone K' formed by connecting a' to C is at least 16ω. Refer to Figure 4.

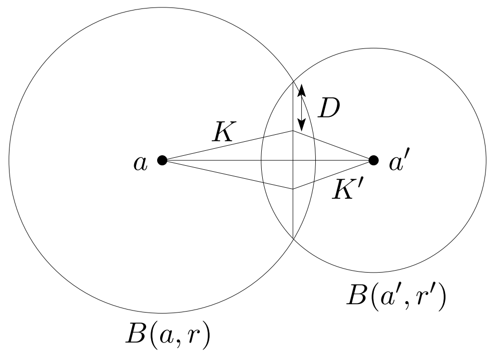

Figure 4: Proof of Lemma 3.1(iv)(d).

Notice that both B_p and B_q are contained in $B(a,3r)$. By Lemma C.10, $\sigma(p,q)$ resides in $K \cup K'$, which means that the minimum distance between $\sigma(p,q)$ and $\mathrm{bd}\,(B_p \cup B_q)$ is at least the minimum distance between $K \cup K'$ and $\mathrm{bd}\,(B_p \cup B_q)$. Denote the latter distance by D. Observe that D is equal to the distance between C and $\mathrm{bd}\, B_p \cap \mathrm{bd}\, B_q$. So by Lemma C.11,

$$D \geqslant r\sin(\arctan(0.09)) - r\sin(8\omega) > 0.009r. \tag{10}$$

By Lemma 3.1(iv)(c), $\sigma(p,q) \subset B_p \cup B_q$ which is contained in $B(a,3r)$. Therefore, for any point $z \in \sigma(p,q)$,

$$\|a - z\| \leqslant 3r, \tag{11}$$

$$\|a' - z\| \leqslant \|a - a'\| + \|a - z\| \leqslant 5r. \tag{12}$$

Suppose that $a \in \mathrm{int}\,\sigma$. As f is 1-Lipschitz over $\mathrm{int}\,\sigma$, $f(z) \leqslant f(a) + \|a - z\| \leqslant f(a) + 3r$ by (11). By Lemma 3.1(ii), $r \geqslant \frac{\lambda\omega}{21} f(a) \Rightarrow f(z) \leqslant \frac{21+3\lambda\omega}{\lambda\omega} r < \frac{22}{\lambda\omega} r$ as $\lambda \leqslant 0.01$ and $\omega \leqslant 0.01$. Substituting this lower bound of r into (10) yields $D > 0.009r > \frac{\lambda\omega}{2500} f(z)$.

If a is an endpoint of σ, then $a' \in \mathrm{int}\,\sigma$ and a' lies on the boundary of $B(a,r)$. So $f(a') \leqslant r$. We have $f(z) \leqslant f(a') + \|a' - z\| \leqslant r + \|a' - z\|$ which is at most $6r$ by (12). Substituting this lower bound of r into (10) yields $D > 0.009r \geqslant 0.0015 f(z)$.

C.6 Lemma 3.1(v)

We first show that the distance between two protecting balls B_p and B_q is at least $0.06 \cdot \min\{\mathrm{radius}(B_p), \mathrm{radius}(B_q)\}$ under certain conditions.

Lemma C.12 *Let (p,s,q) be three weighted points in this order on the 1-face σ (not necessarily consecutive). If $\|p - s\| \geqslant (1+\mu) \cdot \mathrm{radius}(B_p)$ for some $\mu \geqslant 0.02$, the distance between B_p and B_q is at least $0.06 \cdot \min\{\mathrm{radius}(B_p), \mathrm{radius}(B_q)\}$.*

Proof. Let $d(B_p, B_q)$ denote the distance between B_p and B_q. We derive a contradiction below assuming that $d(B_p, B_q) < 0.06 \cdot \min\{\mathrm{radius}(B_p), \mathrm{radius}(B_q)\}$.

By Lemma 3.1(iv)(a), for any protecting ball B_z where $z \in \sigma$, B_z does not strictly contain any weighted point adjacent to z. By Lemma 3.1(iii)(b), $B_p \cap \sigma$ and $B_q \cap \sigma$ are open curves. Therefore, B_p cannot strictly contain s; otherwise, B_p would have to contain a weighted point adjacent to p. Similarly, B_q cannot strictly contain s.

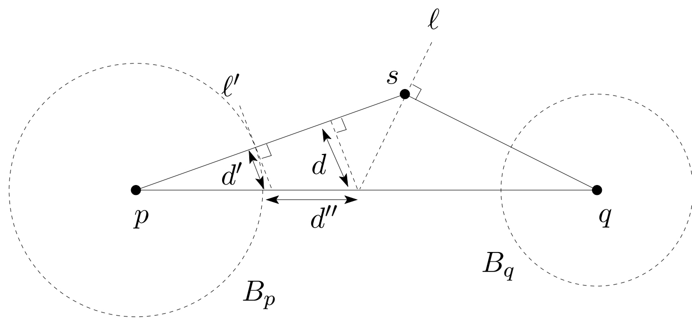

Figure 5: Proof of Lemma C.12.

Refer to Figure 5. Let $\theta_1 = \angle spq$ and let $\theta_2 = \angle sqp$. By the assumption that $d(B_p, B_q) < 0.06 \cdot \min\{\text{radius}(B_p), \text{radius}(B_q)\}$, both B_p and B_q are contained in $B(a, 4r)$, where $B(a, 4r)$ denotes the larger of B_p and B_q. Then, Lemma C.10 implies that $\theta_1 < 8\omega$, $\theta_2 < 8\omega$ and $\angle psq > \pi/2$.

Let ℓ be the line through s perpendicular to qs. Let ℓ' be the line between p and s, perpendicular to ps, and at distance $\mu \cdot \text{radius}(B_p)$ from s. We have $\text{length}(\ell \cap pqs) \geqslant \|p - s\| \sin\theta_1 \geqslant (1+\mu)\sin\theta_1 \cdot \text{radius}(B_p)$. Thus, $d \geqslant (1+\mu)\sin\theta_1\cos(\theta_1+\theta_2)\cdot\text{radius}(B_p)$ and so $d/d' \geqslant (1+\mu)\cos\theta_1\cos(\theta_1+\theta_2)$. In Figure 5, $d(B_p, B_q) \geqslant d''$ and $d'' = \text{radius}(B_p)\cdot(\frac{d}{d'\cos\theta_1} - 1)$. Therefore,

$$d(B_p, B_q) \geqslant \text{radius}(B_p) \cdot ((1+\mu)\cos(\theta_1+\theta_2) - 1). \tag{13}$$

As $\omega \leqslant 0.01$ and $\max\{\theta_1, \theta_2\} < 8\omega$, $\cos(\theta_1+\theta_2) \geqslant \cos(16\omega) \geqslant 0.987$. By this lower bound of $\cos(\theta_1+\theta_2)$ and the assumption that $\mu \geqslant 0.02$, inequality (13) yields $d(B_p, B_q) \geqslant \text{radius}(B_p) \cdot (0.987 + 0.987\mu - 1) > 0.06\,\text{radius}(B_p)$. This contradicts our assumption that $d(B_p, B_q) < 0.06 \cdot \min\{\text{radius}(B_p), \text{radius}(B_q)\}$. ◻

Proof of Lemma 3.1(v). Let B_p and B_q be two non-consecutive protecting balls. We use $d(B_p, B_q)$ to denote the distance between B_p and B_q.

If p and q belong to different 1-faces, then $\|p - q\| \geqslant \max\{g(p), g(q)\}$. For any $x \in \{p, q\}$, Lemma 3.1(ii) implies that $\text{radius}(B_x) < 5\lambda g(x) < g(x)/4$. Therefore, $d(B_p, B_q) \geqslant \|p - q\| - g(p)/4 - g(q)/4 > 0.5 \cdot \max\{g(p), g(q)\}$.

Suppose that p and q belong to the same 1-face σ. We prove Lemma 3.1(v) in this case by verifying that the conditions of Lemma C.12 are satisfied for all pairs of non-adjacent weighted points on σ. We use R_{pq} to denote $0.06 \cdot \min\{\text{radius}(B_p), \text{radius}(B_q)\}$.

Consider the pairs of weighted points $\{(x_j, x_k) : 1 \leqslant j \leqslant k-2 < m-2\}$. By Lemma C.3(ii), $\|x_{k-1} - x_k\| \geqslant \frac{6}{5+6\lambda} r_k = \left(1 + \frac{1-6\lambda}{5+6\lambda}\right) r_k$. Since $\lambda \leqslant 0.01$, $\frac{1-6\lambda}{5+6\lambda} > 0.02$. So we can apply Lemma C.12 to $(p, s, q) = (x_k, x_{k-1}, x_j)$ to obtain $d(B_p, B_q) \geqslant R_{pq}$.

Consider the pairs of weighted points $\{(u, y_m), (u, x_0), (u, v)\}$. We have $x_0 \in B(v, \lambda \mathrm{f}_{\omega_v}(v)) \subset B(v, 0.3\mathrm{f}_{\omega_v}(v))$. By Lemma C.8, $y_m \in B(v, 0.3\mathrm{f}_{\omega_v}(v))$. Since $\|u - v\| \geqslant \max\{g(u), g(v)\} \geqslant$

$\max\{f_{\omega_u}(u), f_{\omega_v}(v)\}$, for any $q \in \{y_m, x_0, v\}$, $d(B_u, B_q) \geqslant (1-\lambda-0.3)\|u-v\| \geqslant \frac{1-\lambda-0.3}{\lambda}\text{radius}(B_u) > 20\,\text{radius}(B_u)$.

Consider the pairs of weighted points $\{(x_k, q) : 2 \leqslant k \leqslant m-2 \text{ and } q \in \{y_m, x_0, v\}\}$. By Lemma C.3(ii), $\|x_k - x_{k+1}\| = 6r_k/5$. So we can apply Lemma C.12 to $(p, s, q) = (x_k, x_{k+1}, q)$ to obtain $d(B_p, B_q) \geqslant R_{pq}$.

The remaining set of weighted points to be considered is $\{(x_{m-1}, x_0), (x_{m-1}, v), (y_m, v)\}$. Since $B_{x_{m-1}} \cap B(x_0, r_{0\,m-1}) = \emptyset$ and $\text{radius}(B_{x_0}) = \frac{4}{5}r_{0\,m-1}$, $d(B_{x_{m-1}}, B_{x_0}) > \frac{1}{5}\text{radius}(B_{x_0})$. By Lemma C.5(i), $\|x_{m-1} - x_0\| > \frac{2}{1+\lambda}r_{m-1} > 1.9r_{m-1}$ as $\lambda \leqslant 0.01$. So we can apply Lemma C.12 to $(p, s, q) = (x_{m-1}, x_0, v)$ to obtain $d(B_p, B_q) \geqslant R_{pq}$. To bound $d(B_{y_m}, B_v)$, we first relate $\text{radius}(B_{y_m})$ and $\|y_m - x_0\|$. Since B_{y_m} and B_{x_0} are orthogonal, $\text{radius}(B_{y_m}) = \sqrt{\|y_m - x_0\|^2 - (\frac{4}{5}r_{0\,m-1})^2}$. By Lemma C.5(ii), $\|y_m - x_0\| \leqslant 4r_{0\,m-1}$. Thus, $\text{radius}(B_{y_m}) \leqslant \sqrt{24/25}\,\|y_m - x_0\| \Rightarrow \|y_m - x_0\| > 1.02\,\text{radius}(B_{y_m})$. So we can apply Lemma C.12 to $(p, s, q) = (y_m, x_0, v)$ to obtain $d(B_p, B_q) \geqslant R_{pq}$.

D Lemma 5.15

We first prove four topological results in the next section. Afterwards, we use them to prove Lemma 5.15.

D.1 Technical lemmas

The next lemma states that for any edge $e \in \text{star}(p, \sigma)$ where p is a point on a 2-face σ, $V_e|_\sigma$ is an open curve under certain conditions. So each facet of V_p contributes at most one curve to $\text{bd}\,(V_p \cap \sigma)$.

Lemma D.1 *Let S be the current admissible point set. Assume that $\|a - b\| \geqslant 0.03\,\text{lfs}^*_\varepsilon$ for any $a, b \in S$. Let p be a point in $S \cap \sigma$ for some 2-face σ. Suppose that* $\mathsf{Infringed}(p, \sigma)$ *and* $\mathsf{CurveNormal}(p, \sigma, \phi_c)$ *return null and* $\text{size}(t, \sigma) < 0.03\,\text{lfs}^*_\varepsilon$ *for any triangle* $t \in \text{star}(p, \sigma)$. *Then, $V_e|_\sigma$ is an open curve for any edge $e \in \text{star}(p, \sigma)$.*

Proof. Take an edge $pq \in \text{star}(p, \sigma)$. As $\mathsf{Infringed}(p, \sigma)$ returns null, $q \in \sigma$. No curve in $V_{pq}|_\sigma$ is closed; otherwise, $\mathsf{CurveNormal}(p, \sigma, \phi_c)$ would return a point. So $V_{pq}|_\sigma$ is a collection of open curve(s). Take any open curve ξ in $V_{pq}|_\sigma$. If V_{pq} avoids $\text{bd}\,\sigma$, both endpoints of ξ belong to $\text{bd}\,V_{pq}$. Otherwise, V_{pq} intersects $\text{bd}\,\sigma$ exactly once by Lemma 5.3 and so at least one endpoint of ξ belongs to $\text{bd}\,V_{pq}$.

Assume that $V_{pq}|_\sigma$ consists of two or more open curves. We derive a contradiction below. There are at least three distinct curve endpoints x_1, x_2 and x_3 in $\text{bd}\,V_{pq}$. By assumption, the weighted distance between p and x_i is less than $0.03\,\text{lfs}^*_\varepsilon$ for $1 \leqslant i \leqslant 3$. Then, $x_i \in B(p, \sqrt{2}\,\text{range}(p))$ for $1 \leqslant i \leqslant 3$ by Lemma 5.8 and $\angle n_\sigma(p), V_{pq} < 4\omega$ by Lemma 5.9. Also, by Lemma 5.7, $n_\sigma(p)$ makes an angle less than 26ω radians with the Voronoi edges containing x_1, x_2 and x_3.

If any two points in $\{x_1, x_2, x_3\}$ lie on the same edge of V_{pq}, say x_1 and x_2, then x_1x_2 is parallel to that edge and so $\angle n_\sigma(p), x_1x_2 < 26\omega$. This is a contradiction because $\angle n_\sigma(p), x_1x_2 > \pi/2 - 2\omega$ by Lemma 5.2(iii).

Suppose that x_1, x_2 and x_3 lie on three distinct Voronoi edges. The support lines of these three edges bound a convex polygon Q containing V_{pq}. (Q is possibly unbounded.) Refer to Figure 6. Let $\vec{d}$ be the projection of $n_\sigma(p)$ onto the plane of V_{pq}. We have $\angle n_\sigma(p), \vec{d} = \angle n_\sigma(p), V_{pq} <$

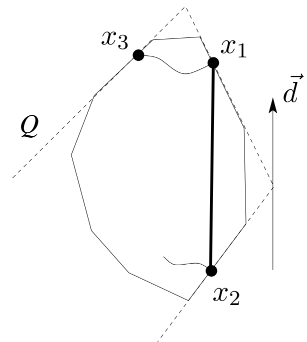

Figure 6: The convex polygon denotes V_{pq}.

4ω. So each side of Q makes an angle less than $4\omega + 26\omega = 30\omega$ radians with $\vec{d}$. Thus, the angle at some vertex of Q is obtuse, say the vertex between x_1 and x_2. It follows that $\angle \vec{d}, x_1x_2 < 30\omega$. However, Lemma 5.2(iii) implies that $\angle \vec{d}, x_1x_2 \geqslant \angle n_\sigma(p), x_1x_2 - \angle n_\sigma(p), \vec{d} > \pi/2 - 6\omega$, a contradiction. ◻

The next result shows that $\mathrm{bd}\,(V_p \cap \sigma)$ is not far from p under certain conditions.

Lemma D.2 *Let S be the current admissible point set. Assume that $\|a - b\| \geqslant 0.03\,\mathrm{lfs}^*_\varepsilon$ for any $a, b \in S$. Let p be a point in $S \cap \sigma$ for some 2-face σ. Suppose that* $\mathsf{Infringed}(p, \sigma)$ *and* $\mathsf{CurveNormal}(p, \sigma, \phi_c)$ *return null and* $\mathrm{size}(t, \sigma) < 0.03\,\mathrm{lfs}^*_\varepsilon$ *for any triangle $t \in \mathrm{star}(p, \sigma)$. Then,* $\mathrm{bd}\,(V_p \cap \sigma) \subset B(p, 6\,\mathrm{range}(p))$.

Proof. If p is unweighted, V_p avoids $\mathrm{bd}\,\sigma$ by Lemma 5.3. So $\mathrm{bd}\,(V_p \cap \sigma)$ contains the curves in $\{V_{pq}|_\sigma : \text{edge } pq \in \mathrm{star}(p, \sigma)\}$. If p is weighted, $\mathrm{bd}\,(V_p \cap \sigma)$ contains the above curves as well as $V_p \cap \mathrm{bd}\,\sigma$.

When p is weighted, by Lemma 5.3, $\mathrm{bd}\,V_p$ intersects $\mathrm{bd}\,\sigma$ exactly twice. So $V_p \cap \mathrm{bd}\,\sigma$ is an open curve. Then, it follows from Lemma 3.1(iv)(c) that $V_p \cap \mathrm{bd}\,\sigma \subset B_p \subset B(p, 6\,\mathrm{range}(p))$.

Take an edge $pq \in \mathrm{star}(p, \sigma)$. By Lemma D.1, $V_{pq}|_\sigma$ is an open curve. Let x_1 and x_2 be its endpoints. If $x_i \in \mathrm{bd}\,V_{pq}$, by assumption, the weighted distance between p and x_i is less than $\mathrm{lfs}^*_\varepsilon$. Then, $\|p - x_i\| < \sqrt{2}\,\mathrm{range}(p)$ by Lemma 5.8. If $x_i \in \mathrm{int}\,V_{pq}$, p must be weighted and $x_i \in V_p \cap \mathrm{bd}\,\sigma \subset B_p$. Hence, $\max_{1 \leqslant i \leqslant 2} \|p - x_i\| < \sqrt{2}\,\mathrm{range}(p)$ and $\|x_1 - x_2\| \leqslant \|p - x_1\| + \|p - x_2\| < 2\sqrt{2}\,\mathrm{range}(p)$.

For any point $z \in \mathrm{int}\,V_{pq}|_\sigma$, since $\mathsf{CurveNormal}(p, \sigma, \phi_c)$ returns null, $\angle x_1 z x_2 > \pi - 4\phi_c > \pi - \pi/4 - 16\omega > \pi/2$. Therefore, $\max_{1 \leqslant i \leqslant 2} \|x_i - z\| \leqslant \|x_1 - x_2\| < 2\sqrt{2}\,\mathrm{range}(p)$. Hence, $\|p - z\| \leqslant \|p - x_1\| + \|x_1 - z\| < 6\,\mathrm{range}(p)$. ◻

Lemma D.2 motivates us to examine $B(p, 6\,\mathrm{range}(p)) \cap \sigma$ which contains $\mathrm{bd}\,(V_p \cap \sigma)$. The next results shows that V_p contains a topological disk in $B(p, 6\,\mathrm{range}(p)) \cap \sigma$ under certain conditions.

Lemma D.3 *Let S be the current admissible point set. Assume that $\|a - b\| \geqslant 0.03\,\mathrm{lfs}^*_\varepsilon$ for any $a, b \in S$. Let p be a point in $S \cap \sigma$ for some 2-face σ. Suppose that* $\mathrm{size}(t, \sigma) < 0.03\,\mathrm{lfs}^*_\varepsilon$ *for any triangle $t \in \mathrm{star}(p, \sigma)$. Let $D \subset B(p, 6\,\mathrm{range}(p)) \cap \sigma$ be a topological disk such that* $\mathrm{bd}\,D \subset \mathrm{bd}\,V_p$, $\mathrm{bd}\,D$ *intersects some edge of V_p, and* $\mathrm{int}\,D \cap \mathrm{bd}\,V_p = \emptyset$. *Then $D \subset V_p$.*

Proof. By assumption, $\mathrm{bd}\,D$ intersects an edge of V_p. So $\mathrm{bd}\,D$ intersects V_t for some triangle $t \in \mathrm{star}(p, \sigma)$. By assumption, $\mathrm{ortho}(t) \leqslant \mathrm{size}(t, \sigma) < 0.03\,\mathrm{lfs}^*_\varepsilon$. So $\angle n_\sigma(p), V_t < 26\omega$ by

Lemma 5.7. Assume to the contrary that $D \not\subset V_p$. As $\mathrm{int}\, D \cap \mathrm{bd}\, V_p = \emptyset$, we conclude that $\mathrm{int}\, D$ is outside V_p. The closed curve $\mathrm{bd}\, D$ splits $\mathrm{bd}\, V_p$ into two topological disks, and let U be one of them. So $D \cup U$ is a topological sphere.

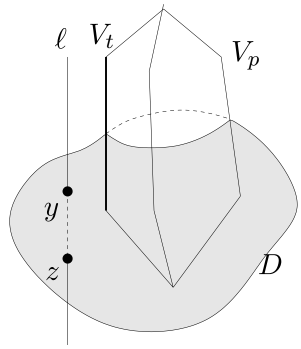

Figure 7: Proof of Lemma D.3.

Refer to Figure 7. Let ℓ be a line outside V_p, parallel to V_t, and arbitrarily close to V_t. Since $\mathrm{int}\, D$ is outside V_p, ℓ must intersect $\mathrm{int}\, D$ at a point y. As $D \cup U$ is a closed surface, ℓ must intersect $D \cup U$ at another point z. Moreover, as ℓ is outside V_p, z must belong to D. Since yz is parallel to ℓ, $\angle n_\sigma(p), yz = \angle n_\sigma(p), V_t < 26\omega$. However, since $y, z \in B(p, 6\,\mathrm{range}(p)) \cap \sigma$, Lemma 5.2(iii) implies that $\angle n_\sigma(p), yz > \pi/2 - 2\omega$, a contradiction. ☐

The possibility that V_p may intersect $\mathrm{bd}\, \sigma$ complicates the analysis because $\mathrm{bd}\,(V_p \cap \sigma)$ needs not be a subset of $\mathrm{bd}\, V_p$. Nonetheless, the next result shows that if $\mathrm{bd}\,(V_p \cap \sigma)$ contains two or more curves, we can always find a "most nested curve" in $\mathrm{bd}\,(V_p \cap \sigma)$ that lies on $\mathrm{bd}\, V_p$.

Lemma D.4 *Let S be the current admissible point set. Let p be a point in $S \cap \sigma$ for some 2-face σ. Let $M = B(p, 6\,\mathrm{range}(p)) \cap \mathrm{mani}(\sigma)$. Suppose that $\mathrm{bd}\,(V_p \cap \sigma) \subset M$ and $\mathrm{bd}\,(V_p \cap \sigma)$ contains two or more closed curves. There exists a closed curve ζ in $\mathrm{bd}\,(V_p \cap \sigma)$ such that $\zeta \subset \mathrm{bd}\, V_p$ and the topological disk bounded by ζ in M does not contain other curves in $\mathrm{bd}\,(V_p \cap \sigma)$.*

Proof. By Lemma 5.2(ii), M is a topological disk. So each closed curve ξ in $\mathrm{bd}\,(V_p \cap \sigma)$ bounds a topological disk in M, which we denote by D_ξ. Pick a most nested curve ξ in $\mathrm{bd}\,(V_p \cap \sigma)$, i.e., D_ξ does not contain other curves in $\mathrm{bd}\,(V_p \cap \sigma)$. If ξ does not contain p, ξ lies on the boundary of $\mathrm{bd}\, V_p$ and we are done.

Suppose that ξ contains p. So p is weighted. By Lemma 5.3, $\mathrm{bd}\, V_p$ intersects $\mathrm{bd}\, \sigma$ exactly twice. Thus, $V_p \cap \mathrm{bd}\, \sigma$ is an open curve contained in ξ and any curve in $\mathrm{bd}\,(V_p \cap \sigma)$ other than ξ must avoid $\mathrm{bd}\, \sigma$. It follows from Lemma 3.1(iii)(b) that $B(p, 6\,\mathrm{range}(p)) \cap \mathrm{bd}\, \sigma$ is an open curve, which cuts across M. We have just shown that for any curve η in $\mathrm{bd}\,(V_p \cap \sigma)$ other than ξ, η cannot cross $B(p, 6\,\mathrm{range}(p)) \cap \mathrm{bd}\, \sigma$. This implies that D_η cannot contain ξ. Hence, either η or the "most nested curve" in $\mathrm{bd}\,(V_p \cap \sigma)$ inside D_η satisfies the lemma. ☐

D.2 Proof of Lemma 5.15

Assume that $\mathrm{bd}\,(V_p \cap \sigma)$ consists of two or more closed curves. We derive a contradiction in the following.

Let $M = B(p, 6\,\text{range}(p)) \cap \text{mani}(\sigma)$. By Lemma D.2, $\text{bd}\,(V_p \cap \sigma) \subset M$. By Lemma 5.2(ii), M is a topological disk. Lemma D.4 implies that there is a closed curve ζ in $\text{bd}\,(V_p \cap \sigma)$ such that $\zeta \subset \text{bd}\,V_p$ and ζ bounds a topological disk D_ζ in M that lies inside V_p.

Consider the support planes of the facets of V_p intersected by ζ. Each such plane bounds a halfspace containing V_p. Let Q denote the polytope at the common intersection of all these halfspaces. So $V_p \subseteq Q$.

Since $M \setminus D_\zeta$ is connected, we can find a path γ in M that connects D_ζ to another curve in $\text{bd}\,(V_p \cap \sigma)$ such that the interior of γ avoids D_ζ. Since $D_\zeta \subset V_p$, γ must leave V_p and Q when it leaves D_ζ. So γ must return to V_p later in order to reach $\text{bd}\,(V_p \cap \sigma)$ again. This implies that γ must return to $\text{bd}\,Q$ at some point y. By construction, the boundary facet of Q containing y must have the same support plane as some facet V_e of V_p that ζ intersects. Let H_e denote the support plane of V_e. Let x be the point in $V_e \cap \zeta$ closest to y. By Lemma D.1, $V_e \cap \zeta = V_e \cap \sigma$ and it is an open curve. We consider the two cases of x lying in the interior of $V_e \cap \sigma$ and x being an endpoint of $V_p \cap \sigma$. We derive a contradiction in each case, thus completing the proof.

Suppose that x lies in the interior of $V_e \cap \sigma$. Any endpoint of $V_e \cap \sigma$ lies on an edge of V_p. So the weighted distances between p and the endpoints of $V_e \cap \sigma$ are less than lfs^*_ε by the assumption of the lemma. Then, Lemma 5.9 implies that $\angle n_\sigma(p), H_e < 4\omega$. Since $x \in \text{int}\,(V_e \cap \sigma)$, xy is normal to $V_e \cap \sigma$. So $\angle n_\sigma(x), H_e = \angle n_\sigma(x), xy$. By Lemma 5.2(i) and (iii), we get $\angle n_\sigma(p), H_e \geqslant \angle n_\sigma(x), xy - \angle n_\sigma(p), n_\sigma(x) \geqslant \angle n_\sigma(p), xy - 2\angle n_\sigma(p), n_\sigma(x) > \pi/2 - 6\omega$, a contradiction.

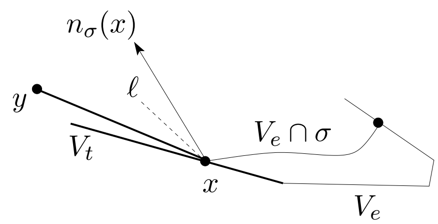

Figure 8: The line ℓ in the plane of V_e is normal to $V_e \cap \sigma$ at x. So $n_\sigma(x)$ projects onto ℓ.

Suppose that x is an endpoint of $V_e \cap \sigma$. Refer to Figure 8. So x lies on some edge of V_p. Let this edge be V_t for some triangle $t \in \text{star}(p, \sigma)$. Since x is the closest point to y, xy must lie inside a sharp wedge bounded by V_t and the line in H_e normal to $V_e \cap \sigma$ at x. Observe that $\angle n_\sigma(x), xy \leqslant \angle n_\sigma(x), V_t$ in this case. By Lemma 5.2(i) and (iii), we have $\angle n_\sigma(x), xy \geqslant \angle n_\sigma(p), xy - \angle n_\sigma(p), n_\sigma(x) > \pi/2 - 4\omega$. On the other hand, by Lemma 5.7, $\angle n_\sigma(x), V_t \leqslant \angle n_\sigma(p), V_t + \angle n_\sigma(p), n_\sigma(x) < 28\omega$, a contradiction.