Localized Delaunay Refinement for Piecewise-Smooth Complexes

Tamal K. Dey* Andrew Slatton†

Abstract

Delaunay refinement, a versatile method of mesh generation, is plagued by memory thrashing when required to generate large output meshes. To address this space issue, a localized version of Delaunay refinement was proposed for generating meshes for smooth surfaces and volumes bounded by them. The method embodies a divide-and-conquer paradigm in that it maintains the growing set of sample points with an octree and produces a local mesh within each individual node, and stitches these local meshes seamlessly. The proofs of termination and global consistency for localized methods exploit recently developed sampling theory for smooth surfaces. Unfortunately, these proofs break down for a larger class called piecewise smooth complexes (PSCs) that allow smooth surface patches that are joined along ridges and corners. In this work, we adapt a recently developed sampling and meshing algorithm for PSCs into the localization framework. This requires revisiting the original algorithm, and more importantly re-establishing the correctness proofs to accommodate the localization framework. Our implementation of the algorithm exhibits that it can indeed generate large meshes with significantly less time and memory than the original algorithm without localization. In fact, it beats a state-of-the-art meshing tool of CGAL for generating large meshes.

*Department of Computer Science and Engineering, The Ohio State University, Columbus, OH 43210, USA. Email: tamaldey@cse.ohio-state.edu

†Department of Computer Science and Engineering, The Ohio State University, Columbus, OH 43210, USA. Email: slatton@cse.ohio-state.edu

1 Introduction

The technique of Delaunay refinement pioneered by Chew [10], Ruppert [14], and Shewchuk [15] has been recognized as an effective tool for sampling and meshing smooth surfaces [3, 9] and volumes bounded by them [13]. The approach utilizes the fact that the sampling of a smooth surface at locally furthest points eventually becomes sufficiently dense to capture its topology and geometry [1, 2]. However, as the sample grows, the Delaunay triangulation starts stressing the available memory, and may eventually trigger memory thrashing if the output mesh is large. This limits the applicability of a theoretically sound technique to generate large meshes. To address this scalability issue, a technique termed localized Delaunay refinement was proposed in [11, 12].

Localized Delaunay refinement works on the divide-and-conquer principle. It maintains the growing set of sampled points with an octree. Each node containing a subset of points is processed individually for Delaunay refinement. By putting a limit on the number of points a node can hold, the size of the Delaunay triangulation is tamed for each node. Of course, now the challenge becomes how to merge the individual meshes seamlessly and also how to guarantee a lower bound on the inter-point distance globally when inserting points for a local meshing. These two questions were addressed in [11, 12] using recent results from the sampling theory [2, 3, 9] for smooth surfaces. These theories do not hold for non-smooth surfaces [4] and for a larger class called piecewise smooth complexes (PSCs) [7] that allow several smooth surface patches to be attached along ridges and corners. This class appears abundantly in meshing applications and the need for computing large meshes for PSCs arises naturally. The localization technique for smooth surfaces obviously does not work for this class and needs to be revisited.

In this paper, we address the problem of localized Delaunay refinement for PSCs. In particular, we localize a ball-protection technique of [7] that was modified further for adaptive refinement in [6] and later in the Delaunay mesh generation book [8]. On the theoretical side, we show how one can adapt the results in [7, 8] to the localization framework of [11, 12]. This requires a careful design of the local refinement steps, some additional analysis of the ball-protection step, and revisiting the sampling theory developed for weighted Delaunay refinement [8]. The Delaunay refinement algorithms for PSC meshing employ weighted points, or *protecting balls*, to preserve sharp features of the input, and these balls may be refined (i.e. - replaced by sets of smaller balls) at any time during the refinement in order to satisfy some refinement criteria. The localization brings up the question of how to refine the balls in the context of the local mesh and still be globally consistent. We show that ball refinement can indeed be performed locally. This requires selecting the point sets for local refinement and reprocessing some of the nodes with some care.

On the practical side, we exhibit that the localization can produce large meshes in the range of millions of simplices with much less time and memory than the original refinement method without localization. In some cases, the localization succeeds where the original refinement fails due to memory shortage. We compare our results with a state of the art CGAL meshing tool, and find that our implementation yields significantly better running times and memory footprints for large meshes.

2 Background

Piecewise-Smooth Complexes. Our input domain $\mathcal{D}$ is a *piecewise smooth complex* (PSC) where each element is a compact subset of a smooth (C^2) k-manifold, $0 \leq k \leq 2$. Each element is closed and hence contains its boundaries. For simplicity we assume that each element has a non-empty boundary. We use $\mathcal{D}_k$ to denote the subset of all k-dimensional elements, the kth stratum. $\mathcal{D}_0$ is a set of *vertices*; $\mathcal{D}_1$ is a set of ridges called 1-*faces*; $\mathcal{D}_2$ is a set of surface patches called 2-*faces*. For $1 \leq k \leq 2$, we use $\mathcal{D}_{\leq k}$ to denote $\mathcal{D}_0 \cup \ldots \cup \mathcal{D}_k$. We use $\operatorname{int} \mathbb{X}$ and $\operatorname{bd} \mathbb{X}$ to denote the interior and boundary of a topological space $\mathbb{X}$, respectively.

The domain $\mathcal{D}$ satisfies the usual proper requirements for being a complex: (i) interiors of the elements are pairwise disjoint and for any $\sigma \in \mathcal{D}$, $\operatorname{bd} \sigma \subset \mathcal{D}$; (ii) for any $\sigma, \sigma' \in \mathcal{D}$, either $\sigma \cap \sigma' = \emptyset$ or $\sigma \cap \sigma'$ is a union of elements in $\mathcal{D}$. We use $|\mathcal{D}|$ to denote the underlying space of $\mathcal{D}$. For $0 \leq k \leq 2$, we also use $|\mathcal{D}_k|$ to denote the underlying space of $\mathcal{D}_k$.

Delaunay Triangulations and Voronoi Diagrams. For a set of points $P \subset \mathbb{R}^3$, we denote its Voronoi diagram and Delaunay triangulation as $\operatorname{Vor} P$ and Del P respectively. Each k-dimensional Delaunay simplex σ (vertex, edge, triangle, and tetrahedron) is dual to a $(3-k)$-dimensional Voronoi face V_σ (cell, facet, edge, vertex respectively).

In this paper, P will be sampled from a PSC. The set of Delaunay simplices whose dual Voronoi faces intersect $\mathcal{D}$ will be of special interest to us. Throughout the refinement, we use a special sub-complex of Del P called the *restricted Delaunay complex* with respect to $\mathcal{D}$. It is defined as:

$$\mathrm{Del}|_{\mathcal{D}}\, P = \{f \in \text{Del}\, P :\ V_f \cap \sigma \neq \varnothing, \sigma \in \mathcal{D}\}.$$

We employ Delaunay triangulations of weighted point sets for PSC refinement. The weighted distance d_w between two points p and q with weights ω_p and ω_q respectively is given as:

$$d_w(p,q)^2 = d(p,q)^2 - \omega_p^2 - \omega_q^2$$

where $d(\cdot,\cdot)$ is the unweighted Euclidean distance. The Voronoi diagram and (restricted) Delaunay triangulation of the weighted point set then follow their standard definitions using the distance d_w. In the algorithm, the points of the triangulation having strictly positive weights are used to “protect” the sharp features of the input geometry, as in [7], and will often be referred to as protecting balls. Specifically, given the weighted point q and its weight ω_q, the protecting ball is a ball $b_q = B(q, \omega_q)$ with center q and radius ω_q.

For a weighted Delaunay triangle f, let size$(f) = \max_{x \in e_f \cap \mathcal{D}}\{d_w(p,x)\}$, where e_f is the edge in the weighted Voronoi diagram dual to f and p is a vertex of f. Observe that size(f) is the radius of an empty ball called a *surface Delaunay ball* of f.

Overall algorithm Our algorithm begins by running PROTECT (from [8]) on the input surface to generate an initial set of protecting balls to cover sharp features; this operation is performed globally without localization, but usually with a small set of points. The algorithm then divides the sample set P using an octree, and processes only one leaf node of the octree at a time. During the processing of a node ν, some conditions are checked. When one of these conditions is violated, the algorithm refines the node's local triangulation accordingly. Note that when a protecting ball is refined, it is refined with regards only to the local point set, and may require that we *remove* some points from our sample set. Showing that this local ball refinement can be accomplished without sacrificing guaranteed termination or mesh integrity is nontrivial, and is a key leading to the correctness of the algorithm. When there are no more violations in ν, it begins processing another node of the octree. When none of the refinement criteria are violated in any of the nodes, all local meshes are brought together to form the final output. Throughout the algorithm, we utilize a point insertion strategy that is necessary for our proof of termination, and we may select some nodes for *reprocessing* in order to maintain a global consistency across meshes.

2.1 Protecting Sharp Features

In this section, we reproduce procedures from [8] that are used for generating and maintaining a set of protecting balls that cover $\mathcal{D}_{\leq 1}$. COVER is used to cover a segment of a ridge in $\mathcal{D}_1$ with balls; REFINEBALL replaces one set of balls with a new set of smaller balls; and PROTECT generates an initial set of protecting balls to cover $\mathcal{D}_{\leq 1}$. We describe the main steps of the algorithms here for completeness and push other details to Appendix A. We use the following notations.

- $B(c, r)$ denotes a geometric ball of radius r centered at c.
- $\text{seg}_{\gamma}(b_p)$ denotes a segment of $b_p \cap \gamma$, where $\gamma \in \mathcal{D}_1$ is a ridge and b_p is a geometric ball centered at p. Specifically, it denotes the connected component of $b_p \cap \gamma$ containing p.
- $\gamma(p, q)$ denotes the segment of $\gamma \in \mathcal{D}_1$ bounded by p and q.
- $\text{lfs}_{min}(\gamma)$ is the minimum feature size w.r.t. γ over all points in γ, where the feature size at a point x in γ is the distance from x to the medial axis of γ.

In many cases, we discuss indexed sequences of protecting balls $b_j, b_{j+1}, ..., b_{k-1}, b_k$. On these occasions, we denote by ω_i and c_i the radius and center respectively of a ball b_i, $j \leq i \leq k$.

Definition 2.1 (Ball Adjacency) *Two protecting balls b_p and b_q are adjacent with respect to sample set P iff p and q both lie on a $\gamma \in \mathcal{D}_1$, and $\text{int}\,(\gamma(p, q))$ is empty of sample points, that is, $\text{int}\,(\gamma(p, q)) \cap P = \varnothing$.*

Definition 2.2 (Contiguous Set of Balls) *A set of protecting balls $\mathcal{B} = \{b_0, ..., b_k\}$ with centers located at $\mathcal{C} = \{c_0, ..., c_k\}$ is contiguous with respect to sample set P iff all balls in $\mathcal{B}$ lie on the same ridge $\gamma \in \mathcal{D}_1$ and, for every pair of balls $b_i, b_j \in \mathcal{B}$, $\text{int}\,(\gamma(c_i, c_j)) \cap (P \setminus \mathcal{C}) = \varnothing$.*

COVER. Given a ridge γ and two disjoint adjacent balls b_0 and b_k centered on γ, we would like to cover the subridge between x and z with protecting balls, where x and z are the endpoints of $\text{seg}_{\gamma}(b_0)$ and $\text{seg}_{\gamma}(b_k)$ (respectively) delimiting the shortest subridge that can be made from the four combinations of endpoints. The method COVER from [8] accomplishes this and can be found in Appendix A, page 13.

REFINEBALL **and** PROTECT. During our refinement, we may find that a protecting ball b is too large to allow our meshing criteria to be satisfied. When this occurs, we call REFINEBALL to replace b (or possibly a contiguous set of balls containing b) with smaller balls. Note that *replacement* of a set of balls entails *removal* of them from the sample; the new set of replacement balls is generated by COVER. It is also possible that we remove some zero-weighted points from our sample during this process. In our algorithm, REFINEBALL is always performed with regards only to a local point set. These methods are taken from [8], and their pseudo-code is provided in Appendix A, page 15.

Given the input domain $\mathcal{D}$, we use PROTECT to initialize our sample set with balls that protect the vertices and ridges of $\mathcal{D}_{\leq 1}$ using an user-input parameter λ_2. See the psuedocode in Appendix A, page 15, for a more detailed description. We use the terms $\mathcal{D}_0$-ball and $\mathcal{D}_1$-ball to denote a protecting ball with its center in $\mathcal{D}_0$ and $\text{int}\,(\gamma \in \mathcal{D}_1)$ respectively.

3 Refinement

3.1 Node Processing

The nodes of the octree to be processed are maintained in a queue Q, and each node is processed when it reaches the head of Q. A node may be processed by one of two actions: split or refine. Each node ν of the octree maintains a set of points $P_{\nu} = P \cap \nu$. When the number of points in ν exceeds a user-defined parameter κ, that is, $|P_{\nu}| > \kappa$, we invoke a split; if $|P_{\nu}| \leq \kappa$ when ν reaches the head of Q, we invoke a refine.

In a split, ν is divided into eight children of equal size, each of which is a scaled down version of ν. The points of P_{ν} are then divided among these children, with each child taking the points that lie within its volume, and then these children are enqueued in Q.

When a node ν is refined, we begin by computing its local triangulation Del R_ν, where R_ν is a *superset* of P_ν. Specifically, we initialize $R_\nu := N_\nu \bigcup P_\nu$, where $N_\nu \subseteq P$ contains the points of P that lie within an unweighted distance $2\omega_{max} + 2\sqrt{(\lambda^2 + \omega_{max}^2)}$ of the boundary of ν, where ω_{max} is the weight of the largest protecting ball generated by PROTECT and λ is a user-defined parameter. While violations of our refinement criteria persist, we refine the local triangulation of ν. If $|P_\nu| > \kappa$ at any point during the refinement of the local triangulation, we invoke a split of ν.

When a node ν is not being processed we clear its local Delaunay triangulation (along with all the data structures associated with it) in order to save memory, maintaining only its $P_\nu \subseteq P$ and a list of restricted triangles inside ν.

3.2 Localized Refinement

$$
\begin{aligned}
F(p) &= \{f \,:\, f \in \text{Del}\, R_\nu|_{\mathcal{D}} \text{ is a triangle incident to } p \\
&\qquad \text{or a sub-simplex of such a triangle}\} \\
F_\sigma(p) &= \{f \,:\, f \in \text{Del}\, R_\nu|_{\sigma} \text{ is a triangle incident to } p \\
&\qquad \text{or a sub-simplex of such a triangle}\} \\
s_{max}(p) &= \max\{\text{size}(f) \,:\, f \in F(p) \text{ is a triangle}\} \\
r_{min}(p) &= \min\{\omega_q \,:\, q \in F(p) \text{ is a vertex with } \omega_q > 0\} \\
d_{min}(p) &= \min\{d(q,s), : \, q, s \in F(p) \text{ are vertices with} \\
&\qquad \omega_q, \omega_s = 0\}.
\end{aligned}
$$

For each point p in a node ν, we want the local triangulation around p to be nice, that is, surface triangles around p form a topological disk. Our ultimate goal is to fit all these individual local triangulations seamlessly into a global one. Toward that goal, we define the *surface star* $F(p)$ ($F_\sigma(p)$) of a point $p \in P_\nu$ as the set of triangles incident to p that are restricted to $\mathcal{D}$ (respectively, $\sigma \in \mathcal{D}_2$) in the local triangulation, and the sub-simplices of all such triangles. We also define the s_{max}, r_{min}, and d_{min}, which are used in handling violations of the refinement criteria, C1-C7 giving priority to CI over CJ if I<J:

1. (C1) No $\mathcal{D}_0$-ball centered on $p \in P_\nu$ contains the center of a $\mathcal{D}_1$-ball;
2. (C2) $\forall p \in P_\nu$ and $l \in R_\nu$, if $b_p \in \gamma_i$ and $b_l \in \gamma_j$ given $\gamma_i, \gamma_j \in \mathcal{D}_1$ and $i \neq j$, then $b_p \cap b_l = \varnothing$;
3. (C3) $\forall p \in P_\nu$, if ps is an edge in $F(p)$ and $\omega_p, \omega_s > 0$, then b_p and b_s are adjacent in some $\gamma \in \mathcal{D}_1$;
4. (C4) $\forall p \in P_\nu$, $\sigma \in \mathcal{D}_2$ contains all vertices of $F_\sigma(p)$;
5. (C5) $\forall p \in P_\nu$, $F_\sigma(p)$ is a topological disk $\forall \sigma \in \mathcal{D}_2$ containing p;
6. (C6) $\forall p \in P_\nu$, $\forall \sigma \in \mathcal{D}_2$ containing p, $p \in \text{int}\,\sigma$ iff $p \in \text{int}\, F_\sigma(p)$;
7. (C7) $\forall f \in F(p)$, where f is a triangle and $p \in P_\nu$, size$(f) < \lambda$.

In the event of a violation, either a point is to be inserted, or a ball is to be refined. In both cases, a pair (p, q) is returned. In the first case q is a zero-weighted candidate for insertion and p is the closest point (by distance d_w) to q in R_ν. In the second case q is the center of the protecting ball to be refined and p is a *null* pointer. Let x be the farthest point from p where the Voronoi edge dual to argmax$\{\text{size}(f) \,|\, f \in F(p) \text{ is a triangle}\}$ intersects $|\mathcal{D}|$. If C1 is violated, we call REFINEBALL on the violating $\mathcal{D}_0$-ball b_p. If C2 is violated, we call REFINEBALL on the larger of b_p and b_l. In the event C3 is violated, if $s_{max}(p) \geq 0.03\text{min}\{\omega_p, \omega_s\}$ then q is set to x; if $s_{max}(p) < 0.03\text{min}\{r_{min}(p), r_{min}(s)\}$ then REFINEBALL is called on the larger of the two balls violating this criterion. If C4 is violated, then let t be the errant vertex of $F_\sigma(p)$, that is $t \notin \sigma$. If

$s_{max}(p) < 0.03\omega_t$ then we call REFINEBALL on b_t; otherwise q is set to x. In the event C5 or C6 is violated, if $s_{max}(p) \geq \min\{0.03r_{min}(p), d_{min}(p)\}$ then q is set to x; if $s_{max}(p) < \min\{0.03r_{min}(p), d_{min}(p)\}$ then REFINEBALL is called on the largest of the balls in $F(p)$. If C7 is violated, q is set to x.

3.3 Point Insertion and reprocessing

Algorithm 1 LOCPSC($\partial\mathsf{O}, \mathcal{D}, \kappa, \lambda, \lambda_2$)

```
P ←PROTECT(D, λ2); ω_max ←argmax_{p∈P} ω_p
Compute a bounding box to initialize Q;
while Q is not empty do
   ν :=DEQUEUE(Q);
   while (p, q) :=VIOLATION(∂O, ν, λ) is not null do
      S :=INSERTREFINE(ν, p, q, λ)
      for all s ∈ S do
         NODEENQUEUE(ν, s, λ, ω_max)
      end for
      if |P_ν| ≥ κ then
         Split ν and enqueue its eight children to Q
      end if
   end while
end while
Return P and ∪_p F_p.
```

Let (p, q) be a pair returned by some violation during refinement. Then q may be a candidate for insertion, but it may lie arbitrarily close to some point in P despite being locally far in R_ν. In order to disallow arbitrarily close insertions, we find the closest point in P to q. If this existing point is not p (recall that p is the nearest point to q in R_ν) and lies sufficiently close to q, then we insert the existing point instead of q. Specifically, if q lies within weighted distance λ of a protecting ball $b \in P \setminus R_\nu$, then we add b to R_ν; if no such b exists and $s \in P \setminus R_\nu$ lies within distance λ of q, then we throw away q and insert s into R_ν; if there is no such b and we do not insert s, we insert q into R_ν and add it to P. Note that R_ν is augmented in all cases, and P is augmented in only the last case.

Furthermore, when a new point is added to or deleted from P it is possible that some nodes are enqueued for reprocessing.

The local mesh of a node ν in the octree is comprised of triangles having at least one vertex in ν. As we show in Lemma 4.4, all vertices of such triangles lie either inside ν or within unweighted distance $2\sqrt{(\lambda^2 + \omega_{max}^2)}$ of its boundary when the algorithm LOCPSC terminates. Call this set of points K_ν. It is possible that when a point $q \notin P$ is inserted during the refinement of node ν, it modifies the content of some $K_{\nu'}$, $\nu' \neq \nu$, through insertion or deletion, and so may affect the part of the final output generated by ν'. Then in order to maintain consistency between the meshes of ν and ν' we must reprocess ν'. Therefore, whenever some $q \notin P$ is inserted by node ν or some $q \in P$ is removed from P by node ν, we enqueue all nodes $\nu' \neq \nu$ within distance $2\sqrt{(\lambda^2 + \omega_{max}^2)}$ of q.

3.4 Algorithm LocPSC

Our algorithm LocPSC first encloses the input surface $\mathcal{D}$ in a bounding box which becomes the root of the octree subdivision. PROTECT is then used to generate a set of balls (positive-weighted points) that protect the $\mathcal{D}_1$ ridges and $\mathcal{D}_0$ vertices, and these balls are used to initialize P.

When a node ν is not being processed, we maintain its sample set $P_\nu = P \bigcap \nu \subseteq P$ and a list of triangles $\bigcup_{p \in P_\nu} F(p)$. This list is maintained to avoid recomputing the mesh in a node while finalizing the output mesh. When ν is extracted from Q for processing, we compute Del R_ν which is also updated with each insertion and deletion of point(s). To check violations C1-C7, we compute and maintain the restricted surface triangulation $\mathrm{Del}|_{\mathcal{D}}\, R_\nu$. Note that REFINEBALL is performed locally, removing only those zero- and positive-weighted points in R_ν. At termination we output $\bigcup_{p \in P} F(p)$.

4 Guarantees

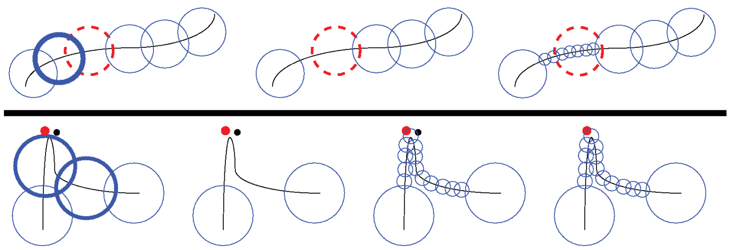

Figure 1: Depiction of complications due to local ball refinement. Top: complication (1) – we refine the emboldened blue ball, but don't see the dashed red ball because it's in $P \setminus R_\nu$, and so place new balls arbitrarily close to and inside of this ball. Bottom: complication (2) – we remove both emboldened balls, and one of the new balls placed contains the red point in $P \setminus R_\nu$, and so this point is not removed from P.

The guarantee that the output is a manifold Delaunay mesh respecting the input stratification hinges on the fact that the algorithm terminate because the refinement conditions C3-C6 would otherwise make the algorithm continue to run. We prove that our algorithm inserts and deletes only finitely many points and thus terminates. We argue that the protecting balls and triangles cannot get refined below a threshold and thus cannot afford insertions/deletions incessantly. On this high level it is similar to the argument in [8]. However, the proofs differ significantly at the detailed level since we have to deal with ball and triangle refinements in a local setting.

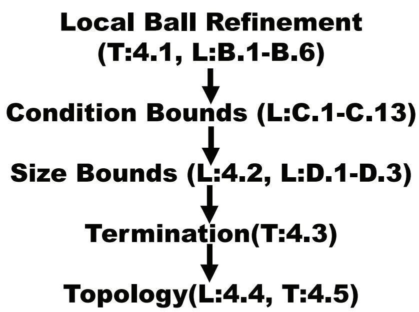

Figure 2: Dependencies of results including the ones in the Appendix [L: Lemma; T: Theorem].

A major hurdle to be overcome is to ensure that a ball refinement that considers only the local R_ν is still good globally. We achieve this by selecting the additional points carefully when constructing R_ν from P_ν. Due to space shortage, we only include the main line of argument in the proof and defer other details to Appendix B, page 19. The lower bounds on distances among adjacent balls help to prove the lower bounds for ball sizes and triangle sizes, see Figure 2. These proofs are also deferred to Appendix C.

Consider the complications that could arise with regards to ball refinement when using only a local point set, see also Figure 1: (1) By nature of a locally performed REFINEBALL, the set of removed balls is contiguous with respect to R_ν, but not necessarily contiguous with respect to P. If this set is not contiguous with respect to P, then the new set of balls placed by COVER may be placed on top of and without regards to some balls in $P \setminus R_\nu$, thus making it impossible to place lower bounds on the geodesic distance along the ridge between adjacent pairs of balls, and such lower bounds are requisite for our proof of termination (specifically, the C3 Bound Lemma C.3 and C5 and C6 Bound Lemmas C.13, see Figure 2) (2) There may exist zero-weighted points in $P \setminus R_\nu$ that lie inside newly placed protecting balls, and it is necessary that this not be the case for our proof of termination (specifically, C5 and C6 Bound Lemma). In the proof of Theorem 4.1, we mention three supporting lemmas that appear in Appendix B, two of which (B.1 and B.3) handle (1) and the last of which (B.5) handles (2).

Theorem 4.1 (Local Ball Refinement) *When* REFINEBALL *is called on a point p during the refinement of some local sample set R_ν, there is a set of points added to and a set removed from R_ν. The constituent points of these two sets would remain the same if R_ν were augmented with any subset of $P \setminus R_\nu$.*

Proof: Exposed Ridge Segment Lemma (B.1) shows that the center of the ball b_p to be refined must lie very close to ν because it must be in the star of some vertex in ν and must lie in sufficiently close proximity to this vertex in ν in order to incur a violation that will call REFINEBALL. This lemma goes on to show that all points on the ridge segment that is exposed by removing balls at the beginning of REFINEBALL must lie close to p – specifically, the entirety of this segment must lie within $3.52\omega_{max}$ of ν. In the Endball Distance Lemma (B.3), we show that the centers of b_0 and b_k (the endballs for COVER) must lie within $4\omega_{max}$ of ν, and so these along with all the balls centered on the exposed segment are in $P_\nu \cup N_\nu$; thus, adding more points to R_ν from $P \setminus R_\nu$ would not yield a different b_0, b_k, or exposed ridge segment, implying that the set of positive-weighted points deleted from P and R_ν cannot be altered by augmenting R_ν with any subset of $P \setminus R_\nu$. Note that this also implies that the set of inserted points cannot be altered by augmenting R_ν with a subset of $P \setminus R_\nu$, because for a given b_0, b_k, and set of removed balls, COVER by nature of being deterministic must generate a fixed set of new balls. We then show in the Point Removal Lemma (B.5) that all zero-weighted points in the union of newly placed balls lie within $4\omega_{max}$ of ν, implying that all zero-weighted sample points that are removed must lie close enough to ν that they are in $P_\nu \cup N_\nu$, so there are no points in $P \setminus R_\nu$ that lie close enough to be removed. ■

The following result is an integral part of our proof for termination.

Lemma 4.2 *$\mathcal{D}_0$-balls with radii less than some surface dependent constant will not be refined.*

Proof Sketch:

1. C1. For $\mathcal{D}_0$-ball b_v to satisfy C1, we require that it not contain the center of any $\mathcal{D}_1$-ball. A violation of this implies that either b_v intersects a ridge disjoint from v or b_v intersects a ridge containing v in multiple disjoint segments. The former is impossible when ω_v is less than the distance to a ridge disjoint from v; the latter is impossible when ω_v is small with respect to the local feature size of the ridges containing v as an endpoint.

2. C2. A violation of C2 involving $\mathcal{D}_0$-ball b_v implies that b_v intersects some ball b_q centered on a ridge γ disjoint from v. As the distance from v to γ must be strictly positive, C2 must be satisfied when b_v and all balls centered on γ have radii less than half the distance from v to γ.

3. C3. A violation of C3 involving $\mathcal{D}_0$-ball b_v implies that b_v and some non-adjacent ball b_q are both vertices of the same restricted Delaunay triangle t. Let b_q lie on a ridge γ disjoint from v. When ω_v and ω_q are less than $d(v, \gamma)/3$, b_v will not be refined because size$(t) \geq 0.03\omega_v$. So let γ contain v as one of its endpoints. In Lemma A.4 on page 16, we show that there is a minimum geodesic distance along γ (as a function of ω_q and ω_v) between q and v, and so when ω_q and ω_v are small enough with respect to local feature size, the Euclidean distance between them implies that any triangle t having both q and v as vertices must have size$(t) > 0.03 \min\{\omega_q, \omega_v\}$, and so a triangle must be refined.

4. C4. A violation of C4 involving $\mathcal{D}_0$-ball b_v implies that v is a vertex of $F_\sigma(p)$ for some σ disjoint from v. Satisfaction of C3 precludes a positive-weighted p. Then when ω_v is less than $d(\sigma, v)/2$, size$(t) > \omega_v/2$ for any t containing both p and v as vertices, so a triangle must be refined.

5. C5 and C6. Let b_v be a $\mathcal{D}_0$-ball and σ be a patch in $\mathcal{D}_2$ such that $v \in \mathrm{bd}\,\sigma$, and let b_v be adjacent to b_p on $\gamma_1 \in \mathcal{D}_1$ and b_q on $\gamma_2 \in \mathcal{D}_2$ along $\mathrm{bd}\,(\sigma)$. We proceed by showing that $\gamma_1(p, v)$ intersects the Voronoi facet $F_{pv} = V_p \cap V_v$ exactly once when ω_v and ω_q are less than $0.06\mathrm{lfs}_{min}(\gamma_1)$, and that $\gamma_1(p, v)$ can intersect no other Voronoi facets when ω_v and ω_q are this small. These follow from the observations that $\gamma_1(p, v) \subset b_p \cup b_q$ and that multiple intersections of the segment with the Voronoi facet imply the existence of a medial axis point within distance less than $0.5\mathrm{lfs}_{min}(\gamma)$ (a contradiction of the definition of lfs_{min}). We also show that the circumradius of a restricted triangle is bounded with respect to its size and the weights of its vertices. Bounding the circumradius allows us to employ some of the lemmas of [9], which, with the knowledge of how ridge segments must intersect Voronoi

facets, enable us to show that $V_v \cap \sigma$ is a half disk with v on its boundary. This implies satisfaction of C5 and C6 almost immediately.

■

Specific bounds for Lemma 4.2 are in the proof of Lemma D.1, which employs results of the Condition Bound Lemmas (C.1, C.2, C.3, C.7, C.13) on page 20. The reader may note that a proof sketch showing that $\mathcal{D}_1$-balls with radii less than some surface dependent constant will not be refined (Lemma D.2) is very similar to the previous lemma for $\mathcal{D}_0$-balls. We provide detailed proofs of these as well as proofs of minimum triangle size (Lemma D.3) in Appendix D.

Theorem 4.3 LocPSC *terminates.*

Proof: The lemmas D.1 (alternatively, 4.2), D.2, and D.3 on page 31, showing that the ball and triangle sizes necessary to satisfy all criteria depend solely on surface geometry imply that the processing of a given node terminates when these size conditions are satisfied. Furthermore, our point-insertion strategy prevents any new point s from being inserted within weighted distance λ of its nearest-neighbor point p in the global point set unless the local triangulation already includes p. This implies that there is a minimum inter-point distance for the global point set: it is the minimum of λ and the inter-point distance necessary to satisfy all criteria. A minimum ball size implies a finite number of calls to RefineBall and therefore a finite number of point removals. Since our domain is bounded, this implies that P consists of a finite number of points, the number of which is ultimately dependent on surface geometry and λ, and once $|P|$ reaches this number there can be no more violations in any node, so the algorithm terminates. ■

Lemma 4.4 (Consistency) *At termination, triangle t_{pqs} with vertices p, q, s is in $F_\sigma(p)$ iff it is in $F_\sigma(q)$.*

Proof: Assume that t_{pqs} is in $F_\sigma(p)$ but is not in $F_\sigma(q)$. This implies that t_{pqs} is either not Delaunay or is not restricted w.r.t. σ when we consider $F_\sigma(q)$. This further implies that $F_\sigma(p)$ is defined w.r.t. a different sample set R_ν than that by which $F_\sigma(q)$ is defined, which requires p and q to lie in different nodes – call them ν and ν' respectively. By C7, we know that size$(t_{pqs}) \leq \lambda$, so $d(p,q), d(q,s), d(p,q) \leq 2\lambda + 2\omega_{max} < 2\omega_{max} + 2\sqrt(\omega_{max}^2 + \lambda^2)$, and therefore p and s must be in R'_ν by the gathering of N'_ν. So there must be some sample point u such that $u \in R'_\nu$, $u \notin R_\nu$, and $d_w(u,x) < d_w(p,x)$, where x is the point at which the Voronoi edge e_x dual to t_{pqs} pierces σ. We know $d_w(p,x) \leq \lambda$ by satisfaction of C7, and $d(p,x) = \sqrt(d_w(p,x)^2 + \omega_p^2)$. This implies $d(u,x) < \sqrt(d_w(p,x)^2 + \omega_p^2)$, and therefore $d(p,u) < 2\sqrt(d_w(p,x)^2 + \omega_p^2) \leq 2\sqrt(\lambda^2 + \omega_{max}^2)$. Then if u had been added to P before the last time we processed ν, it would have been in R_ν by the gathering of N_ν, and so t_{pqs} would not be in $F_\sigma(p)$, a contradiction. But if u had been added to P after the last time we processed ν it would have been in K_ν and ν would have been reprocessed, a contradiction. The only other alternative is for u to have been added to P during the last time we processed ν, in which case it would have been added to R_ν as well, and so t_{pqs} would not be in $F_\sigma(p)$, a contradiction. ■

Theorem 4.5 (Topology) *The output mesh $T = \cup_p F(p)$ of* LocPSC *satisfies the following:*

1. *T is a subcomplex of the restricted Delaunay complex $\mathrm{Del}|_{\mathcal{D}}\, P$;*

2. *Each point in T is at most distance λ from $\mathcal{D}$;*

3. *The underlying space of the set of edges E_σ in T between positive-weighted points in σ (given $\sigma \in \mathcal{D}_2$) is homeomorphic to $\mathrm{bd}\,(\sigma)$, with every vertex of E_σ lying on $\mathrm{bd}\,(\sigma)$, and $\cup_{p \in \sigma \cap P} F_\sigma(p)$ is a 2-manifold with boundary; Furthermore, $\exists \lambda^* > 0$ such that if $\lambda < \lambda^*$ then $\mathrm{bd}\,(\cup_{p \in \sigma \cap P} F_\sigma(p)) = E_\sigma$ and $\exists h$, a homeomorphism from $|\mathcal{D}|$ to $|T|$ such that $h(p) = p\ \forall p \in P$, $h(|\mathrm{bd}\,(\sigma)|) = |E_\sigma|$, and $h(\sigma) = |\cup_{p \in \sigma \cap P} F_\sigma(p)|$.*

Proof: (1) follows from Lemma 4.4, and (2) follows immediately from satisfaction of refinement condition C7. Given consistency, (3) follows from chapter 15 of [8]. ■

5 Experimental Results

In experiments, we ran several examples with different values of κ, see Figure 3. This helped us to find a value of κ which performs well in terms of CPU time and memory footprint. We then compared our implementation to a state of the art CGAL meshing tool using this value of κ which is 1000. The experiments were conducted on a PC with 2GB 667MHz RAM and a 2.8GHz processor running Ubuntu 11.10.

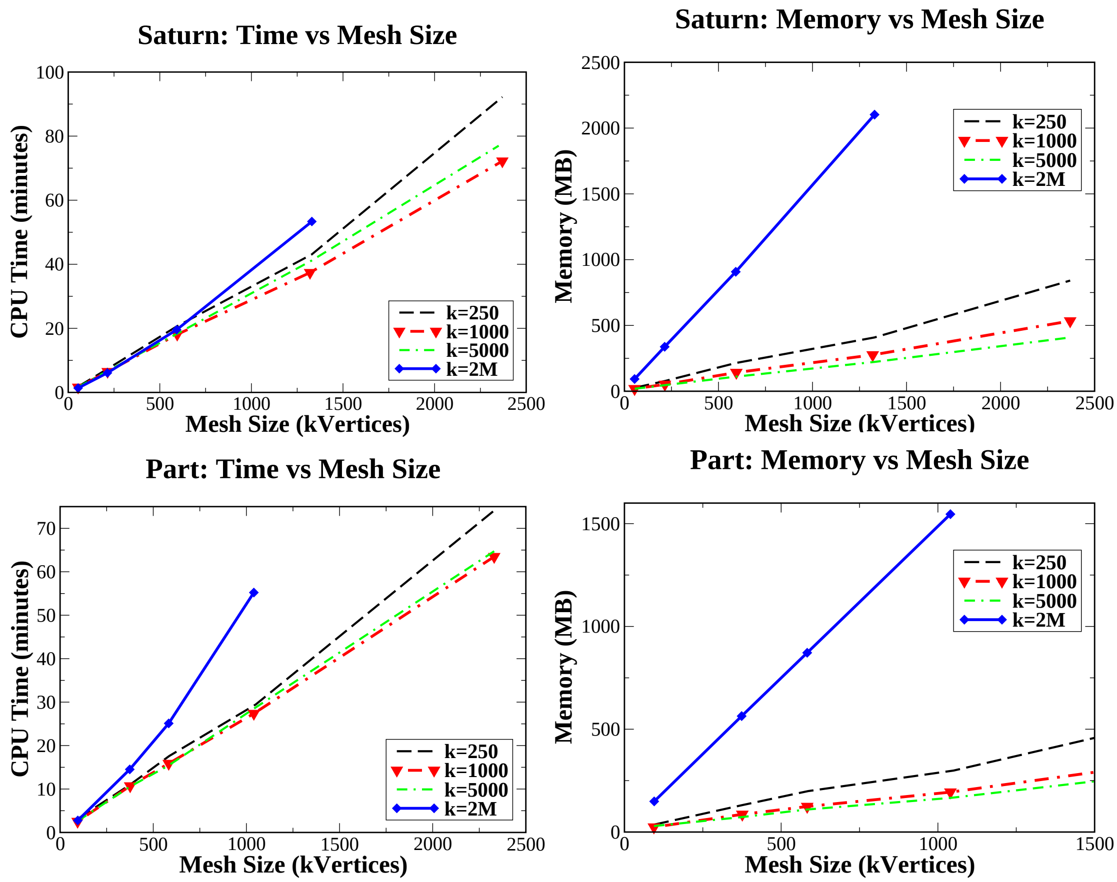

Figure 3: CPU times and memory footprints for varying mesh size and κ, performed on *Saturn* and *Part*. These results reveal $\kappa = 1000$ to be a good choice for CPU time for meshes larger than 350,000 vertices, and for memory footprint. Experiments with $\kappa = 2,000,000$ have only a single node, and therefore are equivalent to running Delaunay refinement without taking advantage of localization.

CGAL Comparison. In Table 1, we show results of experiments on several models that allow comparison of our implementation to a state of the art CGAL meshing tool. The CGAL meshing tool was chosen for comparison because it also employs Delaunay refinement in its meshing algorithm. Therefore, the resultant meshes are similar in both size and quality for a given set of input parameters. Observe that our localized implementation significantly outperforms CGAL in terms of both CPU time and memory footprint for large meshes. Also observe that our localized method is capable of comfortably producing meshes of approximately 8 million vertices on our machine, while CGAL cannot generate meshes much larger than 1.4 million vertices due to memory constraints (and exhibits memory thrashing when generating meshes larger than 1 million vertices). The values of λ here are expressed as a factor of the smallest dimension of the bounding box of the input surface. More results appear in Appendix E on page 33. Times listed were acquired using the CGAL::Timer class and vary from those listed in the CGAL manual due to differences in processor and in experimental parameters.

model	λ	Version	#Vertices (thousand)	#Simplices (thousand)	mem (MB)	CPU Time (sec.)
3Holes	0.0011	LocPSC: $\kappa = 1k$	8272	49632	1530	20302
	0.0011	CGAL	NA	NA	NA	NA
Fandisk	0.0035	LocPSC: $\kappa = 1k$	469	2816	96	802
	0.0035	CGAL	456	2736	695	2550
Rocker	0.0045	LocPSC: $\kappa = 1k$	469	2814	98	1028
	0.0045	CGAL	459	2756	745	2038

Table 1: Time and memory usage for different models for CGAL (CGAL 3.8, release mode, -O3 optimization) and LocPSC results for $\kappa = 1000$. Number of vertices and simplices expressed in thousand unit($\times 1000$); NA indicates that an experiment could not be completed due to memory constraints.

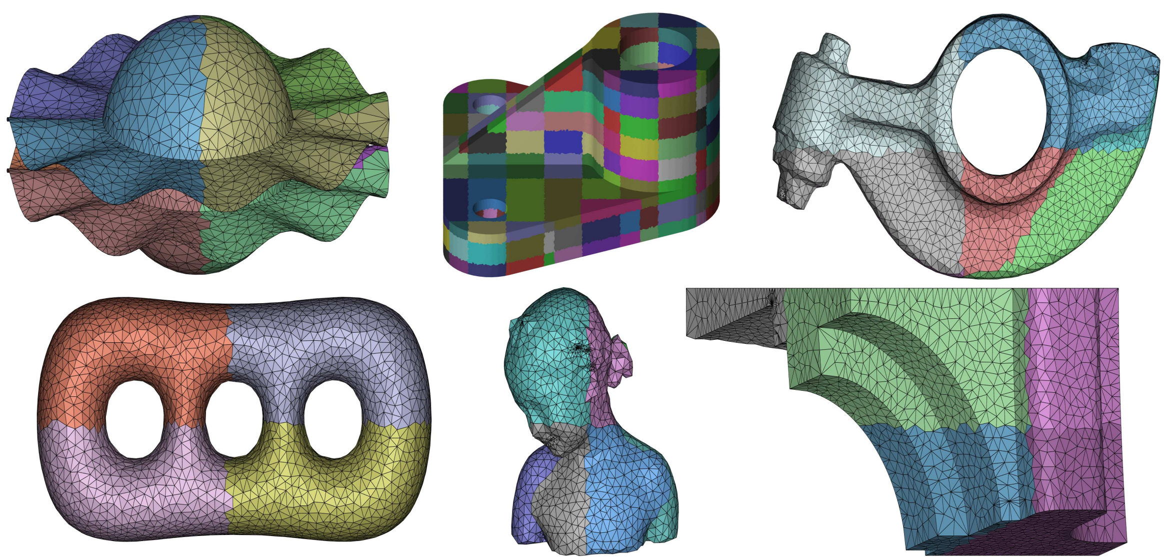

Figure 4: Output meshes of LocPSC. Local meshes from different nodes have different colors.

6 Conclusion

In summary, we have given a localized Delaunay refinement algorithm for PSCs with guarantees for topology, geometry, and termination, the implementation of which outperforms a state of the art CGAL meshing tool in terms of both CPU time and memory footprint for large meshes.

It is possible that localized Delaunay refinement may be the basis for an efficient parallel or distributed Delaunay refinement algorithm for myriad domains with guarantees for topology and geometry, as well as the guarantee that the output is a subcomplex of the restricted Delaunay triangulation of the global point set. Such a result would significantly reduce generation time for Delaunay meshes.

Acknowledgment

We acknowledge the support of the NSF grants CCF-1116258 and CCF-1064416 and thank all the referees for their valuable comments.

References

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

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

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

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

[5] H.-L. Cheng, T. K. Dey, H. Edelsbrunner, and J. Sullivan. Dynamic skin triangulation. *Discrete & Computational Geometry*, 25:525–568, 2001.

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

[7] S.-W. Cheng, T. K. Dey, and E. A. Ramos. Delaunay refinement for piecewise smooth complexes. *Discrete & Computational Geometry*, 2008.

[8] S.-W. Cheng, T. K. Dey, and J. R. Shewchuk. *Delaunay Mesh Generation*. CRC Press, Boca Raton, Florida, 2012.

[9] S.-W. Cheng, T.K. Dey, E.A. Ramos, and T. Ray. Sampling and meshing a surface with guaranteed topology and geometry. *SIAM Journal on Computing*, 37:1199–1227, 2007.

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

[11] T. K. Dey, J. A. Levine, and A. G. Slatton. Localized Delaunay refinement for sampling and meshing. *Computer Graphics Forum*, 29:1723–1732, 2010.

[12] T. K. Dey and A. G. Slatton. Localized Delaunay refinement for volumes. *Computer Graphics Forum*, 30:1417–1426, 2011.

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

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

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

A Methods

In this section, we provide pseudocode for methods from [8], along with related lemmas that are used to support our main results.

We begin by describing the method COVER. Given a ridge γ and two disjoint adjacent balls b_0 and b_k centered on γ, we would like to cover the subridge between x and z with protecting balls, where x and z are the endpoints of $seg_\gamma(b_0)$ and $seg_\gamma(b_k)$ (respectively) delimiting the shortest subridge that can be made from the four combinations of endpoints. Let the protecting balls generated by this procedure be $b_1, ..., b_{k-1}$. Assume the ball b_i has just been generated, and let $y_{i,j+1}$ be the endpoint of $seg_\gamma(B(y_{i,j}, \alpha/12))$ closest to z along γ, with $y_{i,0}$ being the endpoint of $seg_\gamma(b_i)$ closest to z along γ, and α being a constant. Then COVER proceeds as follows:

- If $z \notin \gamma(y_{i,0}, y_{i,5})$ and $z \notin seg_\gamma(B(y_{i,4}, \alpha))$, then $b_{i+1} = B(y_{i,4}, \alpha)$;
- Otherwise, construct the last ball as follows, and return $b_1, ..., b_{k-1}$:
  - If $z \notin \gamma(y_{i,0}, y_{i,5})$ and $z \in seg_\gamma(B(y_{i,4}, \alpha))$, then $b_{i+1} = B(y_{i,4}, 5\alpha/4)$;
  - If $z \in \gamma(y_{i,2}, y_{i,5})$, then $b_{i+1} = B(y_{i,1}, \alpha)$;
  - If $z \in \gamma(y_{i,0}, y_{i,2})$, then replace b_i with $B(c_i, 5\alpha/4)$.

These cases are depicted in Figure 5. COVER(α, x, z) then returns the set of balls generated by this process.

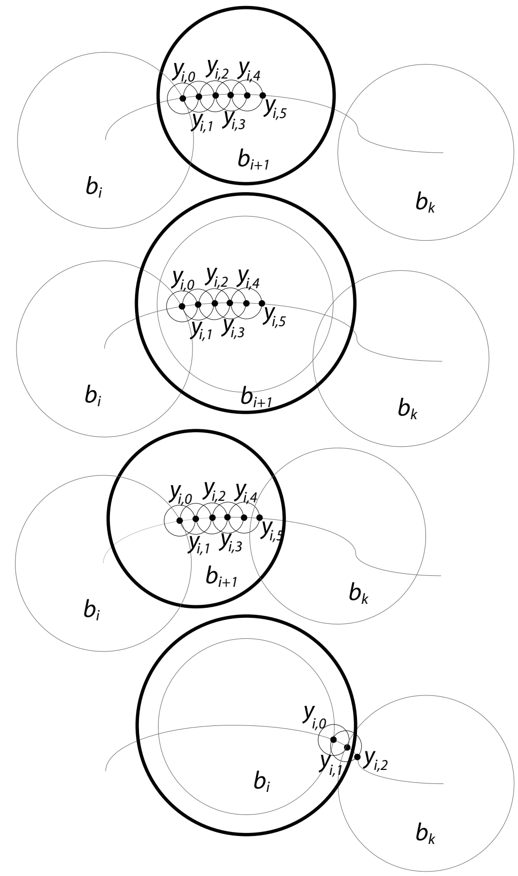

Figure 5: Various cases for ball generation in COVER, depicted in the same order as they are listed in the text. The placed ball is emboldened

In the lemmas and proofs that follow, we use the following notations:

- $d_\gamma(p, q)$ denotes the geodesic distance along γ between p and q.
- $\text{lfs}_\sigma(x)$ is local feature size at x with respect to $\sigma \in \mathcal{D}_1 \cup \mathcal{D}_2$.

Lemma A.1 (Ball Segment) *For sufficiently small protecting balls,* $b_p \cap \gamma =$*seg*$_\gamma(b_p)$.

Proof: Assume otherwise: that $\omega_p < 0.06\text{lfs}_{min}(\gamma)$ (where $\text{lfs}_{min}(\gamma) = \min_{x\in\gamma}\{\text{lfs}_\gamma(x)\}$) and $b_p \cap \gamma$ consists of multiple disjoint segments. Then by Lemma 1.1 of Dey06, b_p must contain a medial axis point of γ, and so this medial axis point lies within distance $0.06\text{lfs}_{min}(\gamma)$ of p, but this is contradictory, as all medial axis points of γ must lie at least distance $\text{lfs}_{min}(\gamma)$ from p. ■

Lemma A.2 (Subridge Covering) $\gamma(x, z)$ *is a subset of the union of balls placed by* COVER(α, x, z).

Proof: Let p and q be a pair of consecutive points along γ placed by a given call to COVER such that q was placed after p. Then by construction all points in $\gamma(p,q) - \text{seg}_\gamma(b_p)$ lie within the union of a set of four balls with radius $\alpha/12$, with the center of the farthest of these balls being no farther than distance $\alpha/3$ from q. This implies that all points in $\gamma(p,q) - \text{seg}_\gamma(b_p)$ lie within $5\alpha/12$ of q, and therefore lie within b_q since b_q is constructed with radius at least α. Since $\text{seg}_\gamma(b_p) \subset b_p$, this means $\gamma(p,q) \subset b_p \cup b_q$, so for a set of balls $b_1, ..., b_{k-1}$ placed by a call to COVER $\gamma(c_1, c_{k-1}) \subset b_1 \cup ... \cup b_{k-1}$. $\gamma(x, c_1)$ follows this same trend by construction of b_1, so $\gamma(x, c_{k-1}) \subset b_1 \cup ... \cup b_{k-1}$. In the placement of the final ball, either $\gamma(c_{k-1}, z) \in \text{seg}_\gamma(B(c_{k-1}, \alpha))$ and b_{k-1} has radius $5\alpha/4$; or $\gamma(c_{k-1}, z)$ lies in the union of a set of three balls of radius $\alpha/12$, with the farthest center being at most distance $\alpha/4$ from c_{k-1}, and radius b_{k-1} is α; or $\gamma(c_{k-1}, z)$ lies in the union of three balls, one with radius α centered on c_{k-1} and the other two of radius $\alpha/12$ with centers no farther than $13\alpha/12$ from c_{k-1}, and radius $b_{k-1} = 5\alpha/4$. In each of these three cases, we find that $\gamma(c_{k-1}, z)$ lies in b_{k-1}, and so $\gamma(x, z) \subset b_1 \cup ... \cup b_{k-1}$. ∎

Lemma A.3 (Euclidean-Geodesic Distance) *For a pair of intersecting balls b_p, b_q centered on $\gamma \in \mathcal{D}_1$, $d(p,q)/d_\gamma(p,q) \geq 0.99939$ when $\omega_p, \omega_q \leq 0.06\textit{lfs}_{min}(\gamma)$.*

Proof: For any two points x and y on curve γ, their Euclidean distance can be related to their geodesic distance along γ by

$$d(x,y) \geq \begin{cases} 2\eta \sin(d_\gamma(x,y)/(2\eta)) & \text{if } d_\gamma(x,y) \leq \pi\eta \\ 2\eta & \text{if } d_\gamma(x,y) \geq \pi\eta \end{cases},$$

where $\eta = \text{lfs}_{min}(\gamma)$. Since the two balls intersect, we see that $d(p,q) \leq \omega_p + \omega_q \leq 0.12\eta$. Because $d(p,q) \leq 2\eta$, $d(p,q) \leq d_\gamma(p,q) \leq 2\eta \arcsin(d(p,q)/(2\eta)) \leq 2\eta \arcsin(0.06)$. So $1 = d(p,q)/d(p,q) \geq d(p,q)/d_\gamma(p,q) \geq 0.12\eta/(2\eta \arcsin(0.06)) > 0.99939$. ∎

Algorithm 4 INSERTREFINE($P, \nu, p, q, \lambda, \mathcal{D}$)

if $p = null$ **then**
  S :=REFINEBALL($P, \mathcal{D}, q$)
  update R_ν; return S
end if
$s \leftarrow \text{argmin}_{u \in P}\ d_w(q,u)$
if $d_w(q,s) \leq \lambda$ and $s \neq p$ **then**
  $R_\nu := R_\nu \cup \{s\}$; return $\{\}$
end if
$P := P \cup \{q\}$; update R_ν; return $\{q\}$

Algorithm 5 NODEENQUEUE($\nu, s, \lambda, \omega_{max}$)

Compute $W := \{\nu' \neq \nu | d(s, \nu') \leq 2\sqrt{(}\lambda^2 + \omega_{max}^2)\}$
for each $\nu' \in W$ **do**
  enqueue(ν', Q)
end for

Lemma A.4 (Distance-Weight Relation) *A call to* COVER *made from* REFINEBALL *or* PROTECT *terminates, with each pair of adjacent balls b_p and b_q satisfying one of the following distance-weight relations:*

i) $\omega_p = \omega_q$ *and* $d_\gamma(p,q) \geq 13\omega_p/12$

ii) $\omega_p = 5\omega_q/4$ *and* $d_\gamma(p,q) \geq 16\omega_p/15$

Algorithm 2 REFINEBALL($P, \mathcal{D}, q$)

1: $S := \varnothing$,
2: **if** $q \in P \cap \mathcal{D}_0$ **then**
3: Cut radius of b_q in half.
4: $S = S \cup \{q\}$
5: **for all** b_p adjacent to b_q **do**
6: S :=REFINEBALL($P, \mathcal{D}, q$)$\cup S$
7: **end for**
8: **else if** $q \in P \cap \mathcal{D}_1$ **then**
9: $\mathcal{B} = \varnothing$
10: Let $\gamma \ni q$, $\gamma \in \mathcal{D}_1$
11: Let b', b'' be the balls adjacent to b_q on γ
12: α :=min{radius(b_q),radius(b'),radius(b'')}$/4$
13: $S := S \cup \{q\}$; $P := P \backslash \{q\}$
14: **if** $d(b', b'') < 2\alpha$ **then**
15: Let b' be a non-$\mathcal{D}_0$-ball.
16: $b_q := b'$; Goto line 9
17: **end if**
18: Let x and z be the endpoints of seg$_\gamma(b')$ and seg$_\gamma(b'')$ delimiting the shortest subridge of the possible combinations of endpoints. $\mathcal{B}$:=COVER(α, x, z).
19: Let $\mathcal{P}$ be the weighted points of $\mathcal{B}$.
20: $S := S \cup \mathcal{P}$; $P := P \cup \mathcal{P}$
21: **for all** p with $\omega_p = 0$ lying in $\cup_{b \in \mathcal{B}} b$ **do**
22: $S = S \cup \{p\}$
23: $P = P \setminus \{p\}$
24: **end for**
25: **end if**
26: return S

Algorithm 3 PROTECT($\lambda_2, \mathcal{D}$)

$\mathcal{B} := \varnothing$
for all $v \in \mathcal{D}_0$ **do**
r_v :=min$_{u \in \mathcal{D}_0, u \neq v}\{d(u, v)\}/3$
$b_v := B(v$,min$\{r_v, \lambda_2\})$
$\mathcal{B} := \mathcal{B} \cup b_v$
end for
for all $\gamma \in \mathcal{D}_1$ **do**
Let u, v be endpoints of γ
α :=min{λ_2,radius(b_u),radius(b_v)}$/4$
Let x and z be the endpoints of seg$_\gamma(b')$ and seg$_\gamma(b'')$ delimiting the shortest subridge of the possible combinations of endpoints.
$\mathcal{B} := \mathcal{B} \cup$COVER$(\alpha, x, z)$
end for
return $\mathcal{B}$

iii) $5\omega_p/4 = \omega_q$ *and* $d_\gamma(p, q) \geq 4\omega_p/3$

iv) $3\omega_p \leq \omega_q$ *and* $d_\gamma(p,q) \geq \omega_p/12 + \omega_q > 13\omega_p/12$

v) $\omega_p \geq 3\omega_q$ *and* $d_\gamma(p,q) \geq \omega_p + \omega_q/12$*,*

where at least one of b_p *and* b_q *was placed by this call to* COVER*.*

Proof:

Claim A.5 *Distance-weight relations are satisfied.*

Proof: Let $b_1, ..., b_{k-1}$ be the sequence of balls placed by a call to COVER(α, x, z) such that b_i is placed immediately after b_{i-1} for $1 < i < k-1$, and let b_0 and b_k be the balls immediately before b_1 and immediately following b_{k-1} along γ respectively.

- Consider the first ball b_1 placed by COVER. It will satisfy $\omega_1 = \alpha$ and $d_\gamma(\text{seg}_\gamma(b_0), c_1) \geq \alpha/3 = \omega_1/3$; otherwise it is the last ball placed, but this would imply that b_0 and b_k satisfy $d(b_0, b_k) \leq 4\alpha/3$, and this cannot be the case because we ensure that $d(b_0, b_k) \geq 2\alpha$ before calling COVER from REFINEBALL, and that $d(b_0, b_k) \geq 4\alpha$ before calling COVER from PROTECT. Also, because $\omega_1 = \alpha$ and $\omega_0 \geq 4\alpha$ by the selection of α in REFINEBALL (or PROTECT), we see $\omega_0 \geq 3\omega_1$. Furthermore, $d_\gamma(c_0, c_1) \geq \omega_0 + d_\gamma(\text{seg}_\gamma(b_0), c_1) \geq \omega_0 + \omega_1/3$, thereby satisfying relation (v). Swapping the roles of p and q, it also satisfies (iv).

- Now consider each pair of adjacent balls b_i and b_{i-1} after that for $i < k-1$. Each of these will have radius α, and will satisfy $d_\gamma(\text{seg}_\gamma(b_{i-1}), c_i) \geq \alpha/3 = \omega_{i-1}/3$; otherwise, either b_{i-1} or b_i must be the last ball placed, but this cannot be because b_{k-1} is the last ball placed and $i < k-1$. So $d_\gamma(c_{i-1}, c_i) \geq \omega_{i-1} + d_\gamma(\text{seg}_\gamma(b_{i-1}), c_i) \geq 4\omega_{i-1}/3 = 4\omega_i/3$, which satisfies relation (i) for both $p = c_i$ and $p = c_{i-1}$.

- Now consider the placement and size of b_{k-1}. This is the last ball placed, and so must be constructed as such.
  - If $z \notin \gamma(y_{i,0}, y_{i,5})$ and $z \in seg_\gamma(B(y_{i,4}, \alpha))$, then $b_{i+1} = B(y_{i,4}, 5\alpha/4)$. This implies $5\omega_{k-2}/4 = \omega_{k-1}$ because $\omega_{k-2} = \alpha$, and $d_\gamma(c_{k-2}, c_{k-1}) \geq 4\alpha/3 = 4\omega_{k-2}/3 = 16\omega_{k-1}/15$, so b_{k-2} and b_{k-1} satisfy (ii) for $p = c_{k-1}$, and by swapping p and q they also satisfy (iii). We also have $3\omega_{k-1} = 15\alpha/4 < 4\alpha \leq \omega_k$, yielding $\omega_k \geq 3\omega_{k-1}$. $d_\gamma(c_k, c_{k-1}) = d_\gamma(c_k, z) + d_\gamma(c_{k-1}, z)$, and $d_\gamma(c_k, z) \geq \omega_k$ and $d_\gamma(c_{k-1}, z) = d_\gamma(c_{k-2}, z) - d_\gamma(c_{k-1}, c_{k-2}) \geq \alpha/12$, so $d_\gamma(c_k, c_{k-1}) \geq \omega_k + \omega_{k-1}/12$. This satisfies (v) for $p = c_k$ and (iv) for $q = c_k$.
  - If $z \in \gamma(y_{i,2}, y_{i,5})$, then $b_{i+1} = B(y_{i,1}, \alpha)$. By the same reasoning as above, $d_\gamma(c_k, c_{k-1}) \geq \omega_k + \omega_{k-1}/12$. Furthermore, $\omega_k \geq 4\omega_{k-1}$, so b_k and b_{k-1} satisfy (v) for $p = c_k$ and (iv) for $q = c_k$. $\omega_{k-2} = \omega_{k-1} = \alpha$, and $d_\gamma(c_{k-1}, c_{k-2}) \geq \omega_{k-2} + \alpha/12 = 13\omega_{k-2}/12 = 13\omega_{k-1}/12$. So b_{k-2} and b_{k-1} satisfy (i).
  - If $z \in \gamma(y_{i,0}, y_{i,2})$, then replace b_i with $B(c_i, 5\alpha/4)$. Since $z \notin B(c_{k-1}, \alpha)$, we have $d_\gamma(c_k, c_{k-1}) \geq 4\omega_{k-1}/5 + \omega_k$, and again $3\omega_{k-1} = 15\alpha/4 < 4\alpha \leq \omega_k$, yielding $\omega_k \geq 3\omega_{k-1}$. This satisfies (v) for $p = c_k$ and (iv) for $q = c_k$. $5\omega_{k-2}/4 = \omega_{k-1}$ because $\omega_{k-2} = \alpha$, and $d_\gamma(c_{k-2}, c_{k-1}) \geq \omega_{k-2} + 4\alpha/12 = 4\omega_{k-2}/3 = 16\omega_{k-1}/15$. So b_{k-1} and b_{k-2} satisfy (ii) for $p = c_{k-1}$, and by swapping p and q they also satisfy (iii).

■

Claim A.6 COVER *terminates.*

Proof: Every pair of consecutive balls along γ satisfies at least one of the above size-distance relations, and both $\omega_p \geq \alpha$ and $\omega_q \geq \alpha$, where α is a strictly positive constant for a given call to COVER. This implies that we move at least some fixed positive distance along γ with each new ball that we place, at least $16\alpha/15$. Since $d_\gamma(x, z)$ must be finite (closure of $\gamma(x, z)$ is compact and embedded in Euclidean space, and so must be totally bounded), this implies that a finite number of balls are placed. ■ ■

Lemma A.7 (Positive Power Distance) *No zero-weighted point lies inside a protecting ball.*

Proof: There are only two ways in which we could have a zero-weighted point p inside a protecting ball b_q: either (1) p is inserted within distance ω_q of q, or (2) q is inserted within distance ω_q of p and p is not removed. In the event that $q \in R_\nu$, the former case cannot occur because it would imply that a Voronoi edge e intersects $\mathcal{D}$ inside b_q, which further implies that the power distance of this intersection point to each vertex of t =dual(e) is negative, and this can only be the case when all three vertices of t are positive-weighted. This implies a violation of C3, so we cannot insert p into b_q when C3 is satisfied. We insert no zero-weighted points in handling violations of C1 and C2, so if p is ever inserted into b_q then it is inserted in response to a violation of C3, but this cannot be the case, because the weights of the vertices of t are all strictly positive and size(t) must be negative, so $s_{max}(q) < 0.03 r_{min}(q)$, and a ball is refined rather than t. In the event that $q \in P \setminus R_\nu$, then we find q in INSERTREFINE, and $d_w(q, p) < 0 < \lambda$, so we add q to R_ν instead of adding p to P and R_ν. So p is not inserted within distance ω_q of q. In the event that $p \in R_\nu$, the latter case, in which q is inserted within distance ω_q of p and p is not removed, is precluded by the final step of REFINEBALL, which removes all zero-weighted points within distance ω_q of q immediately after q is inserted. Note that $p \notin P \setminus R_\nu$, as we incorporate all points within distance $2\omega_{max} + 2\sqrt{(}\lambda^2 + \omega_{max}^2)$ of ν into R_ν, and the greatest distance from ν at which REFINEBALL removes points is $4\omega_{max}$ (Lemma B.5). ■

Lemma A.8 (Ridge-Facet Intersection) *Let $p, q \in \gamma \cap R_\nu$, $\gamma \in \mathcal{D}_1$, $p \in \nu$, and let F be a Voronoi facet. For protecting balls b_p and b_q with radii less than 0.06lfs(p) and 0.06lfs(q) respectively, we have the following:*

i) *When C1, C2, and C3 are satisfied, $\gamma(p, q) \cap F \neq \varnothing$ iff $F = V_p \cap V_q$ and $\operatorname{int}(\gamma(p, q)) \cap R_\nu = \varnothing$;*

ii) *If $\gamma(p, q) \cap F \neq \varnothing$ then $\gamma(p, q) \cap F$ is a single point.*

Proof: Assume $\exists s \in P$ such that $V_s \cap \gamma(p, q) \neq \varnothing$. In the proof of the Subridge Covering Lemma, we show that for any two protecting balls b_i and b_j adjacent along some ridge γ, $\gamma(c_i, c_j) \subset b_i \cup b_j$, where at least one of b_i, b_j was placed by COVER(α, x, z), and the other was either placed by the same call to COVER or satisfies $z \in \operatorname{bd}(\text{seg}_\gamma(b_j))$. We proceed by first showing $\gamma(p, q) \subseteq b_p \cup b_q$. Assume the contrary, that $\gamma(p, q) \setminus b_p \cup b_q \neq \varnothing$. Then one of the following must be true: (1) p and q are not adjacent on γ; (2) in some call to COVER(α, x, z), either $x \notin$seg(b_0) or $z \notin$seg(b_k); (3) $\exists t \in P \setminus R_\nu$ such $t \in \gamma(p, q)$. (1) is false by assumption. (2) is false because we only call COVER from REFINEBALL and PROTECT, and both of these satisfy $x \in$seg(b_0) and $z \in$seg(b_k) when making a call to COVER(α, x, z). (3) also cannot be true because a point $t \in \gamma(p, q)$ satisfying $t \in P \setminus R_\nu$ implies $d(t, p) > 4\omega_{max}$, but as pointed out in the proof of the Near Ball Refinement Lemma (LemmaB.3) if there is such a t in $P \setminus R_\nu$ then $\exists t' \in P \cap \gamma(p, q)$ such that $d(t', p) \leq 5\omega_{max}/4 < 4\omega_{max}$, and so this t' would have been found in the initialization of R_ν, thus precluding the condition $\operatorname{int}(\gamma(p, q)) \cap R_\nu = \varnothing$, and so $\gamma(p, q) \subset b_p \cup b_q$ when $p \in \nu$. So every point in $\gamma(p, q)$ has a strictly negative weighted distance to at least one of p, q. This implies $\min_{x \in \gamma(p,q)}\{d_w(x, p), d_w(x, q)\} < d_w(x, s)\ \forall x \in \gamma(p, q)$, where $s \in P$ is a zero-weighted point, so $x \notin V_s$ if $\omega_s = 0$. This further implies that if $V_s \cap \gamma(p, q) \neq \varnothing$ then $\omega_s > 0$. If s does not lie on γ then b_s does not intersect any ball on γ by satisfaction of C2, and therefore does not intersect γ, again yielding $\min\{d_w(x, p), d_w(x, q)\} < d_w(x, s)$. Then $s \in \gamma$. By the Ball Segment Lemma, if protecting balls are sufficiently small then $b_s \cap \gamma$ =seg$_\gamma(b_s)$, and this implies that either $s \in \gamma(p, q)$ or $p \in$seg$_\gamma(b_s)$. The former is precluded by the conditions of this Lemma ($\text{int}\,(\gamma(p,q)) \cap R_\nu$), and the latter is precluded by the Size-Distance Relation Lemma (these imply $d_\gamma(s,p) > \omega_s$ and $d_\gamma(s,p) > \omega_p$). Then no such s can exist, and $\gamma(p,q) \subset V_p \cup V_q$. Since $\gamma(p,q) \subset V_p \cup V_q$ and $\gamma(p,q) \cap V_s = \varnothing$ for all other $s \in P$, $\gamma(p,q)$ can traverse only the Voronoi facet $F = V_p \cap V_q$, and it must traverse this facet because it is connected and contains points in both V_p and V_q. This concludes the proof of (i).

Let $F_{pq} = V_p \cap V_q$. Consider (ii) and assume the contrary: that $F_{pq} \cap \gamma(p,q)$ contains multiple points, $x_1, ..., x_{n-1}$. Let $x_1, ..., x_{n-1}$ be ordered such that in moving from p to q along γ one encounters them in order (x_1 first, x_2 second, etc.), and let $x_0 = p$ and $x_n = q$. Then either $\gamma(x_i, x_{i+1}) \subset V_p$ (specifically, when i is even) or $\gamma(x_i, x_{i+1}) \subset V_q$ (when i is odd) for all $0 \leq i < n$. Since $\gamma(p,q)$ begins in V_p and terminates in V_q, and F_{pq} is the only Voronoi facet traversed by it, then there must be an odd number of points in $F_{pq} \cap \gamma(p,q)$. This implies $n > 3$, so let us consider the smallest ball $B_{2,3}$ containing x_2 and x_3 on its boundary, and grow it into $B'_{2,3}$ by moving its center into V_p along the normal to F_{pq} while maintaining $\{x_2, x_3\} \subset \text{bd}\,(B'_{2,3})$ until $B'_{2,3} \cap \gamma(x_0, x_1) \neq \varnothing$. Notice that no matter how much we grow $B'_{2,3}$, $F_{pq} \cap B'_{2,3}$ is always a circle of diameter $d(x_2, x_3)$. So if x_1 does not lie in this circle then $\gamma \cap B'_{2,3}$ contains at least two disjoint connected components and therefore contains a medial axis point of γ. This is depicted in Figure 1. Let $\omega_p, \omega_q \leq 0.06\text{lfs}_{min}(\gamma)$. From here, we see that $R^2 = d_1^2 + (R - d_2)^2$, and $d_1^2 = r^2 + d_2^2$, where

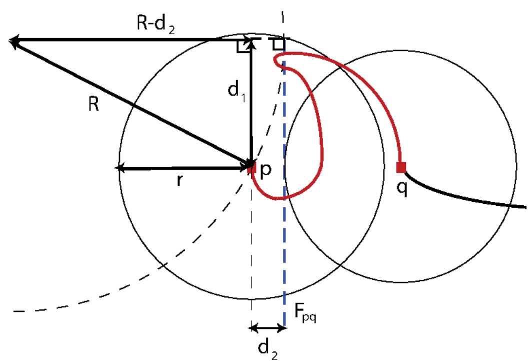

Figure 6: Ridge-Facet Figure.

$r = \omega_p$ and $d_2 = d(p, F_{pq})$. Since b_p and b_q intersect, F_{pq} must contain at least one point in b_q, so from our Distance-Weight Relations and Euclidean-Geodesic Distance we obtain $d_2/\omega_p \geq (d(p,q) - \omega_q)/\omega_p > (0.99939 * 13/12 - 1)$ in the worst case. Then the radius R of $B'_{2,3}$ when it is fully grown (meaning, when that it is large enough to contain p) can be no greater than $R \leq \omega_p^2/(2d_2) \leq 6.06 * 0.06\text{lfs}_{min}(\gamma) < 0.4\text{lfs}_{min}(\gamma)$, and therefore cannot contain a medial axis point of γ, a contradiction. Then x_1 must lie in $F_{pq} \cap B'_{2,3}$. But then we can construct $B_{1,2}$ as the smallest ball containing x_1 and x_2 on its boundary, and grow it into $B'_{1,2}$ by moving its center into V_q along the normal to F_{pq} while maintaining $\{x_1, x_2\} \subset \text{bd}\,(B'_{1,2})$ until $B'_{1,2} \cap \gamma(x_{n-1}, x_n) \neq \varnothing$. Notice that no matter how much we grow $B'_{1,2}$, $F_{pq} \cap B'_{1,2}$ is always a circle of diameter $d(x_1, x_2) < d(x_2, x_3)$. So x_3 cannot lie in this circle, and $\gamma \cap B'_{1,2}$ contains at least two disjoint connected components and therefore contains a medial axis point of γ. Again, the radius of $B'_{1,2}$ when it is fully grown can be no greater than $6.06 * 0.06\text{lfs}_{min}(\gamma)$, and therefore cannot contain a medial axis point of γ, a contradiction. So $F_{pq} \cap \gamma(p,q)$ cannot contain multiple points, $x_1, ..., x_{n-1}$, and (ii) must be true for $\omega_p, \omega_q \leq 0.06\text{lfs}_{min}(\gamma)$. ■

Lemma A.9 (Domain Sample) $P \subseteq \mathcal{D}$.

Proof: Consider all the ways in which we insert new points into P. Positive-weighted points are inserted via COVER or PROTECT, which insert points only in $\mathcal{D}_0 \cup \mathcal{D}_1 \subset \mathcal{D}$. All zero-weighted points that are added to P lie in $E(\mathrm{Vor}\,(P)) \cap \mathcal{D} \subset \mathcal{D}$ at the moment before their insertion, where $E(\mathrm{Vor}\,(P))$ is the set of edges in the Voronoi diagram of P. There are no negative-weighted points, so all points in P must lie in $\mathcal{D}$. ■

B RefineBall Operates Locally

Lemma B.1 (Exposed Ridge Segment) *When* RefineBall *is called from a node ν, all points in the ridge segment exposed by the removal of existing protecting balls lie either inside ν or within distance $3.52\omega_{max}$ of its boundary.*

Proof: For this proof, we consider a violation of each criterion.

1. C1. Here, we refine a $\mathcal{D}_0$-ball b_v with center inside ν. By definition of ω_{max}, its radius is initially at most ω_{max}, and this ball's radius is decreased by a factor of two in the refinement – call the shrunken ball b'_v. In the worst case, this leads us to refine a $\mathcal{D}_1$-ball b_q of radius at most $5\omega_{max}/16$ that is adjacent (with respect to R_ν) to b_v. By nature of Cover, the center of every ball adjacent to b_v with respect to the global point set must lie within distance $21\omega_{max}/16$ of v, implying that they are in N_ν and therefore also in R_ν. This implies that $d(v, q) \leq 21\omega_{max}/16$. By nature of RefineBall, the farthest ball removed b_p must either be b_q or must satisfy $d(b_p, b'_v) < \omega_{max}/8$ (if b_q is not the only ball removed from this ridge, then at least one ball of radius $\omega_{max}/4$ or smaller must be removed, and any ball removed b_p must then satisfy $d(b_p, b'_v) < 2\alpha \leq \omega_{max}/8$). So for $p \neq q$, $d(p, v) \leq \omega_v/2 + 2\alpha + \omega_p \leq (1/2 + 1/8 + 5/16)\omega_{max} = 15/16\omega_{max}$. This implies that $b_p = b_q$, and the farthest point from ν exposed by ball removal is at most distance $\omega_v + 2\omega_q \leq 13\omega_{max}/8$ from ν.

2. C2. Here, we could refine a $\mathcal{D}_0$-ball or a $\mathcal{D}_1$-ball. First we consider the case of refining a $\mathcal{D}_0$-ball b_v. To violate C2, b_v must intersect some ball b_s having its center in ν, and b_s must be a $\mathcal{D}_1$-ball because Protect precludes non-empty intersections of all pairs of $\mathcal{D}_0$-balls. So $d(s, v) \leq (1 + 5/16)\omega_{max}$, and we may refine a $\mathcal{D}_1$-ball b_q adjacent to b_v that lies farther from ν than b_v does. By the same argument used for C1, we see that the center of the farthest ball b_p from b_v removed by this refinement must satisfy $d(v, p) \leq \omega_v + \omega_q \leq 21\omega_{max}/16$, so the farthest point exposed by ball removal lies at most distance $d(s, v) + d(v, p) + \omega_p \leq (21/16 + 21/16 + 5/16)\omega_{max} = 47\omega_{max}/16$ from ν, concluding our analysis of $\mathcal{D}_0$-ball refinement for C2. If we instead refine a $\mathcal{D}_1$-ball b_q, $q \notin \nu$, in response to a violation of C2, then it intersects some smaller ball b_s, $s \in \nu$. If b_q is the only ball removed during RefineBall, then the bound on the distance to the farthest exposed point is small and trivially computed, so assume that the farthest ball removed is $b_p \neq b_q$. In the worst case, b_q was adjacent to a $\mathcal{D}_0$-ball b_v of radius ω_{max}, and the ridge γ containing q wraps around b_v such that the farthest ball b_p from ν satisfying $d(p, v) \leq \omega_v + 2\alpha + \omega_p$ lies on the opposite side of b_v from ν. Then the distance to the farthest exposed point from ν is at most $d(s, q) + d(q, v) + \omega_v + 2\alpha + 2\omega_p \leq (9/16 + 21/16 + 1 + 1/8 + 1/2)\omega_{max} = 3.5\omega_{max}$.

3. C3. The proof for this is identical to that for C2 with the sole exception of the distance between balls necessary to violate the criterion. For C2, this distance had to be less than zero; for C3 it must be less than $0.06 * 5\omega_{max}/16$, so the worst case bound here is the worst case bound for C2 plus this distance, yielding a bound of $(3.5 + 0.06 * 5/16)\omega_{max} < 3.52\omega_{max}$.

4. C4,C5,C6. The refined ball here must either be adjacent to a ball with center in ν, in which case the maximal bound is better than that for C2, or it must share an edge with a zero-weighted point in ν, in which case the maximal bound is better than that for C3.

■

Corollary B.2 *Every ball in P centered on the exposed ridge segment is added to R_ν either when it is initialized or during the current refinement of ν.*

Lemma B.3 (Endball Distance) *When* COVER *is called from* REFINEBALL *while refining node ν, then the centers of the protecting balls b_0 and b_k bounding the exposed ridge segment lie inside ν or within distance $4\omega_{max}$ of its boundary.*

Proof: Let c_k =argmax$_{p\in\{c_0,c_k\}}\{d(\nu,p)\}$. We first consider the case in which b_k is a $\mathcal{D}_1$-ball. If b_k is a $\mathcal{D}_1$-ball then $\omega_k \leq 5\omega_{max}/16$. As the farthest point on the exposed ridge segment from ν is at most $3.52\omega_{max}$ from ν and seg(b_k) bounds this exposed segment, c_k must lie within distance $(3.52 + 5/16)\omega_{max} < 4\omega_{max}$ of ν.

Now we consider the case in which b_k is a $\mathcal{D}_0$-ball. We begin by assuming $d(\nu, c_k) > 4\omega_{max}$, and show this cannot be the case. $\omega_k \leq \omega_{max}$, which implies a point on the exposed ridge at least distance $3\omega_{max}$ from ν. Such a distance between ν and the farthest exposed point implies that a $\mathcal{D}_1$-ball adjacent to a $\mathcal{D}_0$-ball has been refined in response to the violation. The distance from c_k to the nearest $\mathcal{D}_0$-ball center u is at least $d(u, c_k) \geq 3\max\{\omega_k, \omega_u\}$, so $d(b_k, b_u) \geq \max\{\omega_k, \omega_u\}$, but then the farthest point x from ν on the exposed ridge segment must satisfy $d(x, b_u) \leq 2\alpha+2\omega_p$, where b_p is the farthest ball from ν removed during REFINEBALL. Since p lies on a ridge interior, it cannot be a $\mathcal{D}_0$-ball, so $d(x, b_u) \leq (1/8+5/8)\omega_k$, implying $d(b_k, b_u) \leq (1/8+5/8)\max\{\omega_u, \omega_k\}$, but this contradicts $d(b_k, b_u) \geq \max\{\omega_k, \omega_u\}$. So $d(\nu, c_k) \leq 4\omega_{max}$. ■

Corollary B.4 *Endballs b_0 and b_k used to bound the segment to be covered by a call to* COVER *issued by* REFINEBALL *are added to R_ν either when it is initialized or during the current refinement of ν. Furthermore, there is no ball in P centered on the ridge segment to be covered that is not removed by* REFINEBALL*, as the contiguous ball set with respect to P beginning with b_0 and ending with b_k must be a subset of R_ν by the previous two lemmas.*

Lemma B.5 (Point Removal) *Any zero-weighted point removed when* REFINEBALL *is called from a node ν lies either inside ν or within $4\omega_{max}$ of its boundary.*

Proof: Consider new $\mathcal{D}_0$-balls that are placed: they space occupied by them is a subset of the space occupied by the $\mathcal{D}_0$-balls they have replaced because each new $\mathcal{D}_0$-ball has the same center as its predecessor and a smaller radius. Consider new $\mathcal{D}_1$-balls that are placed: their centers must lie within $3.52\omega_{max}$ of ν, and their radii can be no greater than $25\omega_{max}/256$, so they cannot occupy any space farther than $4\omega_{max}$ from ν. Since we only remove zero-weighted points that lie in the space occupied by newly placed balls, all of these points must lie within distance $4\omega_{max}$ of ν. ■

Corollary B.6 *All zero-weighted points removed by* REFINEBALL *are added to R_ν either when it is initialized or during the current refinement of ν. Furthermore there is no zero-weighted point in $P \setminus R_\nu$ that could be removed by this process, as no such point lies within $4\omega_{max}$ of ν.*

We refer the interested reader back to Theorem 4.1 to see how these lemmas are assimilated into the proof that REFINEBALL can operate locally.

C Condition Bound Lemmas

Here we argue termination via a packing argument, showing that each criterion implies lower-bounds on inter-point distances and protecting ball sizes. Throughout our arguments, we use the following definitions:

- $\gamma^* = |\gamma| - (b_u \cup b_v)$, where b_u and b_v are the protecting balls centered on the endpoints of ridge γ;

- $\sigma^* = |\sigma| - \mathcal{B}_\sigma$, where $\mathcal{B}_\sigma$ is the set of protecting balls centered on the boundary of patch σ;
- $D_{min,0,1} = \min_{v\in\mathcal{D}_0}\{\min_{\gamma\in\mathcal{D}_1}\{d(v,\gamma)|v \notin \text{cls}\,(\gamma)\}\}$ (the smallest distance between any $\mathcal{D}_0$ vertex and a non-adjacent $\mathcal{D}_1$ ridge);
- $D_{min,0,2} = \min_{v\in\mathcal{D}_0}\{\min_{\sigma\in\mathcal{D}_2}\{d(v,\sigma)|v \notin \text{cls}\,(\sigma)\}\}$ (the smallest distance between any $\mathcal{D}_0$ vertex and a non-adjacent $\mathcal{D}_2$ patch);
- $D_{min,1,2} = \min_{\sigma\in\mathcal{D}_2}\{\min_{\gamma\in\mathcal{D}_1}\{d(\sigma,\gamma)|\gamma \not\subset \text{cls}\,(\sigma)\}\}$ (the smallest distance between any $\mathcal{D}_1$ ridge and a non-adjacent $\mathcal{D}_2$ patch);
- $D_{min,1,1} = \min_{\gamma_1\in\mathcal{D}_1}\{\min_{\gamma_2\in\mathcal{D}_1}\{d(\gamma_1^*,\gamma_2^*)\}\}$ (the smallest distance between any pair of $\mathcal{D}_1$ ridges excluding portions of those ridges contained in the protecting balls centered on their endpoints);
- $D_{min,2,2} = \min_{\sigma_1\in\mathcal{D}_2}\{\min_{\sigma_2\in\mathcal{D}_2}\{d(\sigma_1^*,\sigma_2^*)\}\}$ (the smallest distance between any pair of $\mathcal{D}_2$ patches excluding portions of those patches contained in the protecting balls centered on their boundaries).

The reader should take notice that $D_{min,0,1}$, $D_{min,0,2}$, and $D_{min,1,2}$ are all surface dependent constants, while $D_{min,1,1}$ is dependent on $\mathcal{D}_0$-ball size and $D_{min,2,2}$ is dependent on $\mathcal{D}_0$ and $\mathcal{D}_1$-ball sizes.

Lemma C.1 (C1 Bound) *When the radii of all $\mathcal{D}_0$-balls are less than $\min\{0.06\mathit{lfs}_{min}(\gamma), D_{min,0,1}\}$, C1 must be satisfied.*

Proof: Consider b_v and b_q, where b_v is a $\mathcal{D}_0$-ball and b_q is a non-$\mathcal{D}_0$-ball. We claim that b_v and b_q cannot violate C1 when their radii are less than $\min\{0.06\text{lfs}_{min}(\gamma), D_{min,0,1}\}$. Assume otherwise. Let b_v and b_q lie on the same ridge γ. Then by the Ball Segment Lemma $q \in \text{seg}_\gamma(b_v)$. This cannot be the case by the construction of new balls in COVER(x,z,α) when called by PROTECT or REFINEBALL, as the center of each new ball lies in $\gamma(x,z)$, the shortest subridge delimited by the endpoints of $\text{seg}_\gamma(b_0)$ and $\text{seg}_\gamma(b_k)$, and $\gamma(x,z)\cap P = \varnothing$, by the Empty Subridge Lemma. So let b_v and b_q lie on distinct ridges. Then $d(v,q) \geq D_{min,0,1}$ by the definition of $D_{min,0,1}$, but then in order for b_v to contain q it must be that $\omega_v \geq D_{min,0,1}$, contradicting the assumption that the radii of b_v and b_q are less than $\min\{0.06\text{lfs}_{min}(\gamma), D_{min,0,1}\}$, so when all $\mathcal{D}_0$-balls have small radii, C1 must be satisfied. ■

Lemma C.2 (C2 Bound) *When protecting balls are sufficiently small, C2 must be satisfied. Specifically, C2 must be satisfied when $\mathcal{D}_0$-balls have radii less than $D_{min,0,1}/2$ and $\mathcal{D}_1$-balls have radii less than $\min\{D_{min,0,1}, D_{min,1,1}\}/2$.*

Proof: $\mathcal{D}_0$-ball size. For a $\mathcal{D}_0$-ball centered on p and some other weighted point q to incur a violation of C2, p must lie on a $\mathcal{D}_0$ vertex disjoint from the ridge containing q, and so we have $D_{min,0,1} \leq d(p,q)$. Also for violation we require $d(p,q) \leq \omega_p + \omega_q$. So when $\omega_p + \omega_q < D_{min,0,1}$ no violation can occur, implying that when both ω_p and ω_q are less than $D_{min,0,1}/2$ no violation can occur.

$\mathcal{D}_1$-ball size. For points p and q to lie on the interiors of distinct ridges, we require $D_{min,1,1} \leq d(p,q)$. For a violation to occur, we require $d(p,q) \leq \omega_p + \omega_q$. So when $\omega_p + \omega_q < D_{min,1,1}$ no violation can occur, implying that when both ω_p and ω_q are less than $D_{min,1,1}/2$ no violation can occur. ■

Lemma C.3 (C3 Bound) *When $\mathcal{D}_0$-balls are smaller than $\min\{0.06\mathit{lfs}_{min}(\gamma), D_{min,0,1}/4\}$, $\mathcal{D}_1$-balls are smaller than $\min\{0.06\mathit{lfs}_{min}(\gamma), D_{min,0,1}, D_{min,1,1}\}/4$, and triangles have size less than $0.03\min_{p\in\mathcal{D}_{\leq 1}}\{\omega_p\}$, C3 must be satisfied.*

Proof: We will consider two main cases here: (1) the two balls b_p and b_q in violation of C3 lie on different ridges, and (2) b_p and b_q lie on the same ridge.

For the former case, we may have $p \in \mathcal{D}_0$ or we may have both p and q on the interiors of distinct ridges of $\mathcal{D}_1$.

Claim C.4 *If $s_{max}(p), s_{max}(q), \omega_p, \omega_q < D_{min,0,1}/4$ given $p \in \mathcal{D}_0$ and q is another positive-weighted point on a ridge distinct from all those ridges containing p as an endpoint, then p and q cannot violate C3.*

Proof: By assumption regarding the relationship between p and q given in the claim, we have $D_{min,0,1} \leq d(p,q)$. For these vertices to violate C3, it must be that $d(p,q) \leq \sqrt(s_{max}(p)^2+\omega_p^2)+\sqrt(s_{max}(q)^2+\omega_q^2) \leq s_{max}(p) + \omega_p + s_{max}(q) + \omega_q$. By assumption, $s_{max}(p) + \omega_p + s_{max}(q) + \omega_q < D_{min,0,1}$. Combining these equations, we arrive at the contradiction $D_{min,0,1} < D_{min,0,1}$, so p and q cannot violate C3 given the assumptions in our claim. ■

Claim C.5 *If $s_{max}(p), s_{max}(q), \omega_p, \omega_q < D_{min,1,1}/4$ given p and q lie on the interiors of distinct ridges γ_1 and γ_2, then p and q cannot violate C3.*

Proof: By assumption regarding the relationship between p and q given in the claim, we have $D_{min,1,1} \leq d(p,q)$. For these vertices to violate C3, it must be that $d(p,q) \leq \sqrt(s_{max}(p)^2+\omega_p^2)+\sqrt(s_{max}(q)^2+\omega_q^2) \leq s_{max}(p) + \omega_p + s_{max}(q) + \omega_q$. By assumption, $s_{max}(p) + \omega_p + s_{max}(q) + \omega_q < D_{min,1,1}$. Combining these equations, we arrive at the contradiction $D_{min,1,1} < D_{min,1,1}$, so p and q cannot violate C3 given the assumptions in our claim. ■

The previous two claims show that for sufficiently small triangles and protecting balls, vertices p and q on distinct ridges cannot violate C3. It also shows that the $\mathcal{D}_0$-ball size required to ensure no violations depends solely on surface-dependent constant $D_{min,0,1}$. The $\mathcal{D}_1$-ball size and triangle size required to ensure no violations depend on $D_{min,1,1}$, which depends on both surface geometry and $\mathcal{D}_0$-ball size.

Now let us consider the case where p and q lie on the same ridge γ.

Claim C.6 *If $\omega_p, \omega_q < 0.06lfs_{min}(\gamma)$ and $s_{max}(p), s_{max}(q) \leq 0.03 \min\{\omega_p, \omega_q\}$ given p and q lie on the same ridge γ, then p and q cannot violate C3.*

Proof: For violation to occur, we have $d(p,q) \leq \sqrt(s_{max}(p)^2 + \omega_p^2) + \sqrt(s_{max}(q)^2 + \omega_q^2) \leq s_{max}(p) + \omega_p + s_{max}(q) + \omega_q$. Then by assumption $d(p,q) < 0.1236\text{lfs}_{min}(\gamma)$, implying $d_\gamma(p,q)/d(p,q) \leq 1.00064$ and $d(p,q)/d_\gamma(p,q) \geq 0.99936$ by the equation in the Euclidean-Geodesic Distance Lemma. Also, for p and q to violate there must be some other point $s \in \gamma(p,q) \cap P$. For what follows, let $s, v \in \text{int}\,(\gamma(p,q))$ such that s is adjacent to q and v is adjacent to p, and note that s and v may or may not be the same point.

(Case 1) Consider the case in which both pairs (p,s) and (v,q) satisfy one of the first three Distance-Weight Relations. Then $d_\gamma(p,q) \geq d_\gamma(p,s) + d_\gamma(v,q) \geq 16(\omega_q + \omega_p)/15 > 1.00064 * 1.03(\omega_p + \omega_q) \geq 1.00064(\omega_q + \omega_p + 2 * 0.03 \min\{\omega_p, \omega_q\}) \geq 1.00064 d(p,q) \geq d_\gamma(p,q)$, but this equation contradicts itself ($d_\gamma(p,q) > d_\gamma(p,q)$) and so p and q cannot violate C3 when both pairs (p,s) and (v,q) satisfy one of the first three Distance-Weight Relations.

(Case 2) Consider the case in which pair (p,s) satisfies one of the last two Distance-Weight Relations such that $\omega_s > \omega_p$ and pair (v,q) satisfies one of the first three relations. Then $d_\gamma(p,q) \geq 16\omega_q/15 + 37\omega_p/12 \geq 16(\omega_q + \omega_p)/15$. As we know from the previous case, this again results in $d_\gamma(p,q) > d_\gamma(p,q)$, and so p and q cannot violate C3 when pair (p,v) satisfies one of the last two relations such that $\omega_s > \omega_p$ and pair (v,q) satisfies one of the first three relations.

(Case 3) Consider the case in which both pairs (p,s) and (v,q) satisfy one of the last two Distance-Weight Relations such that $\omega_s > \omega_p$ and $\omega_v > \omega_q$. Then $d_\gamma(p,q) \geq 37\omega_q/12 + 37\omega_p/12 \geq 16(\omega_q + \omega_p)/15$. As we know from a previous case, this again results in $d_\gamma(p,q) > d_\gamma(p,q)$, and so p and q cannot violate C3 when both pairs (p,s) and (v,q) satisfy one of the last two Distance-Weight Relations such that $\omega_s > \omega_p$ and $\omega_v > \omega_q$.

(Case 4) Consider the case in which pair (p,s) satisfies one of the last two Distance-Weight Relations such that $\omega_s < \omega_p$ and pair (v,q) satisfies one of the first three relations. Then we know that $d_\gamma(q,v) \geq 16\omega_p/15$ and $d_\gamma(v,b_p) \geq \omega_v/12 \geq \omega_q/15$, implying $d_\gamma(q,b_p) \geq 17\omega_q/15$. Let $x = \gamma(q,p) \cap \text{bd}\,(b_p)$.

The Ball Segment Lemma implies that x is a single point for balls of the sizes we consider here. Let Π_{xp} be the plane tangent to b_p at x, and let $\mathcal{L}_{xq}$ be a line passing through x and q. If we can show that $d(q, \Pi_{xp}) \geq 1.06\omega_q \geq \omega_q + 0.06 \min\{\omega_p, \omega_q\}$ then we know $d(p, q) > \omega_p + \omega_q + 0.06 \min\{\omega_p, \omega_q\}$, which contradicts the requirements for violation. We proceed by lower-bounding the angle $\angle_a \mathcal{L}_{xq}, \Pi_{xp}$ and thereby lower-bounding $d(q, \Pi_{xp})$ in terms of ω_q. To bound $\angle_a \mathcal{L}_{xq}, \Pi_{xp}$, we adapt [9] Lemma 2.3 to curves. The Sandwich Property [5], which is the basis for [9] Lemma 2.3, can be adapted to curves; however, for curves it may be more aptly named the Doughnut Property. For any point $z \in \gamma$ there exists a set of balls $\mathcal{B}_z = \{b : \gamma \text{ is tangent to } b \text{ at } z, \text{and radius of } b \text{ is } \mathrm{lfs}_\gamma(z)\}$ the union of which is empty of points in γ. The centers of the balls in $\mathcal{B}_z$ form a circle that lies in the plane Π_z normal to the tangent of γ at z. This yields $\angle \mathcal{L}_{yz}, t_z \leq \pi/2 - \arccos(d(y, z)/(2\mathrm{lfs}_\gamma(z)))$, where $\mathcal{L}_{yz}$ is a line passing through $y \in \gamma$ and z, and t_z is the tangent vector to γ at z. Since $\mathcal{L}_{xp}$, the line passing through x and p, is perpendicular to Π_{xp}, we find $\angle_a \Pi_{xp}, t_x \geq \pi/2 - \angle \mathcal{L}_{xp}, t_x$, and $\angle_a \Pi_{xp}, \mathcal{L}_{xq} \geq \angle_a \Pi_{xp}, t_x - \angle \mathcal{L}_{xq}, t_x \geq 1.5067$. Now consider the nearest point $u \in \Pi_{xp}$ to q. $\mathcal{L}_{uq}$ must be perpendicular to Π_{xp}, and so $\mathcal{L}_{uq}$, $\mathcal{L}_{ux}$, and $\mathcal{L}_{xq}$ form a right triangle with $\angle \mathcal{L}_{ux}, \mathcal{L}_{xq} \geq 1.5067$. This implies that $d(p, q) \geq \omega_p + d(q, \Pi_{xp}) = \omega_p + \sin(1.5067) d(x, q) \geq \omega_p + 0.99936 \sin(1.5067) d_\gamma(x, q) \geq 1.128\omega_q > \omega_p + \omega_q + 0.06 \min\{\omega_p, \omega_q\} \geq d(p, q)$, a contradiction.

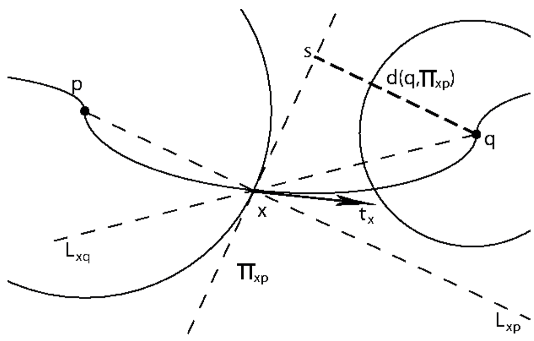

Figure 7: $d(p, \Pi_{xq})$.

(Case 5) Consider the case in which both pairs (p, s) and (v, q) satisfy one of the last two Distance-Weight Relations such that $\omega_s < \omega_p$ and $\omega_v < \omega_q$. Assume that $\omega_p \geq \omega_q$. We claim that either (1) at some point after p and q came into existence but before v and s came into existence there must have existed a ball b_z such that $z \in \mathrm{int}\,(\gamma(p, q))$ and $\omega_z \geq 4\omega_q/5$, or (2) both b_p and b_q are $\mathcal{D}_0$-balls. In the latter case, $d_\gamma(p, q) \geq d(p, q) \geq 3\omega_p \geq 1.5(\omega_p + \omega_q)$, which we know from a previous case results in $d_\gamma(p, q) > d_\gamma(p, q)$, a contradiction. The former case reduces to Case 4 above by the observation that the existence of such a b_z implies $d(p, b_q) \geq 17\omega_q/15$, and this case also ends in contradiction.

Existence of b_z*:* Consider the moment immediately after the call to COVER that generated b_p terminated, and in particular consider the two balls adjacent to b_p. By nature of COVER, these were either placed by the same call to COVER that generated b_p or they are the endballs delimiting the subridge that was covered. in either case, the radii of both balls adjacent to b_p are greater than or equal to $4\omega_p/5 \geq 4\omega_q/5$. A similar proof can be made showing that immediately after b_q is generated both balls adjacent to it have radii greater than or equal to $4\omega_q/5$. So independent of the order in which b_p and b_q are generated there must have existed a ball matching the description of b_z at some point in time after b_p and b_q came into existence.

(Case 6) Consider the case in which both pairs (p, s) and (v, q) satisfy one of the last two Distance-Weight Relations such that $\omega_s < \omega_p$ and $\omega_v > \omega_q$. Here we make the observation that $d_\gamma(q, b_p) \geq d_\gamma(q, v) \geq 37\omega_q/12 \geq 17\omega_q/15$, and so this case also reduces to Case 4 above. ■ ■

Lemma C.7 (C4 Bound) *When all $\mathcal{D}_0$-balls have radii less than $D_{min,0,2}/3$ and all $\mathcal{D}_1$-balls have radii* *less than $D_{min,1,2}/3$ and all triangles have size less than $\min\{D_{min,0,2}/3, D_{min,1,2}/3, D_{min,2,2}/2\}$ and $s_{max}(p) \leq r_{max}(p) \forall p$ C4 must be satisfied.*

Proof: There are two main cases of violation that we may encounter here: (1) a positive-weighted vertex q of $F_\sigma(p)$ lies outside σ, and (2) a zero-weighted vertex of $F_\sigma(p)$ lies outside σ. In the former case, we refine the errant vertex q if its weight is greater than the size of the largest triangle in $F(p)$; In all other cases, we refine the largest triangle of $F(p)$. Assume the aforementioned radius and size conditions are met, and $F_\sigma(p)$ violates C4.

Assume vertex q in $F_\sigma(p)$ is the errant vertex. By satisfaction of C3, we know that at most one of p and q can be positive-weighted. Assume then that $q \in \mathcal{D}_0$. Then $D_{min,0,2} \leq d(p,q) \leq s_{max}(p) + \sqrt{(}s_{max}(p)^2 + r_{max}(p)^2)$. Assuming $s_{max}(p) \leq r_{max}(p)$, this implies that no violations occur when $r_{max}(p) + \sqrt{(}2r_{max}(p)^2) < D_{min,0,2}$, or when $r_{max} < D_{min,0,2}/(1+\sqrt{2})$.

Now assume that $q \in \mathcal{D}_1$. Then $D_{min,1,2} \leq d(p,q) \leq s_{max}(p) + \sqrt{(}s_{max}(p)^2 + r_{max}(p)^2)$. Assuming $s_{max}(p) \leq r_{max}(p)$, this implies that no violations occur when $r_{max}(p) + \sqrt{(}2r_{max}(p)^2) < D_{min,1,2}$, or when $r_{max} < D_{min,1,2}/(1+\sqrt{2})$.

Now assume that $p, q \in \mathcal{D}_2$. Then $D_{min,2,2} \leq d(p,q) \leq 2s_{max}(p)$. This implies that no violations occur when $2s_{max}(p) < D_{min,2,2}$, or when $s_{max} < D_{min,2,2}/2$. ∎

Before proceeding to the proof of lower bounds implied by C5 and C6, we must first establish a few preliminaries. We employ Lemmas 2.1-4.7 of [9] extensively in the proof of bounds for C5 and C6; however, [9] deals with unweighted Delaunay triangulations, and it is not immediately apparent how – or even whether – these lemmas can be applied to weighted triangulations. In the following lemmas, we develop theory that allows us to make use of these lemmas from [9] with some caveats.

Lemma C.8 (Bounded Circumradius) *When a ball would be refined in response to a violation of C5 or C6 by point p, the circumradius of each triangle in $F(p)$ is less than $s_{max}(p) + r_{max}(p)$.*

Proof: The case in which all angles are acute is trivial, as the circumradius must be less than the unweighted distance from a positive-weighted vertex to the orthocenter, which in turn must be less than $s_{max}(p) + r_{max}(p)$. The rest proof proceeds via a case-by-case analysis of triangles in $F(p)$ with respect to how vertex weights are arranged about the triangle's obtuse angle.

The case in which all three vertices of a triangle t in $F(p)$ are protecting balls is precluded by satisfaction of C3. The case in which all three vertices of t are zero-weighted trivially satisfies circumradius$(t) \leq s_{max}(p)$. The case in which exactly one vertex has positive weight can be divided into two subcases: (1) the protecting ball lies on the obtuse angle, and (2) the protecting ball lies on one of the acute angles. In the former case, the circumcenter of t is closer to the positive-weighted vertex than the orthocenter is, and so the circumradius must be less than $s_{max}(p) + r_{max}(p)$. The latter case requires a more rigorous analysis, as the circumradius lies farther from all vertices than the orthocenter. For this case, we proceed by bounding the angle between edge su and the plane bisecting the edge qs, where q is the positive weighted point and $\angle qsu$ is obtuse. This angle bound can then be used to bound the distance between the circumcenter and the orthocenter with respect to $r_{max}(p)$.

Claim C.9 *The circumradius of obtuse triangle t in $F(p)$ with exactly one positive weighted vertex lying on one of its acute angles is at most $s_{max}(p) + r_{max}(p)$.*

Proof: Let the vertices of t be q (the positive-weighted vertex), s (the vertex on the obtuse angle), and u; let x_{su} be point where the bisecting plane of edge su intersects edge su; and let o and c be the ortho- and circumcenter of t respectively. Then $\angle s x_{su} o = \pi/2$. $d(s, x_{su}) \geq d_{min}(p)/2$ because $d(s,u) \geq d_{min}(p)$ by the definition of $d_{min}(p)$. $d(s,o) \leq s_{max}(p)$ by the definition of $s_{max}(p)$, and $s_{max}(p) < d_{min}(p)$ because we would only refine a ball in such a case. This yields $\angle o s x_{su} \leq \arccos(0.5)$. Let x_{sq} be the point where the affine hull of su intersects the plane supporting the Voronoi facet $V_s \cap V_q$. Then $\angle x_{sq}x_{su}o = \pi/2$ and $\angle ox_{sq}x_{su} \leq \angle osx_{su} \leq \arccos(0.5)$ (otherwise either $\angle qsu$ is acute, contradicting assumption, or the plane supporting $V_s \cap V_q$ does not intersect edge sq, which is precluded by the Positive Power Distance Lemma). Now consider a line segment oy parallel to edge qs, with y lying in the plane bisecting edge qs. Triangle yoc is similar to triangle $x_{su}x_{sq}o$, and so $d(c,o) = d(y,o)\sec(\angle ox_{sq}x_{su}) \leq 2d(y,o)$. $d(y,o)$ is merely the distance between the bisecting plane Π_b and the plane supporting the Voronoi facet Π_{qs} of qs. So, $d(y,o) = d(q,\Pi_{qs}) - d(q,\Pi_b)$. $d(q,\Pi_b) = d(q,s)/2 = (\omega_q + d(b_q,s))/2$ and $d(q,\Pi_{qs}) \leq \omega_q + d(b_q,s)/2$. This yields $d(y,o) \leq \omega_q/2 \leq r_{max}(p)/2$ and so $d(c,o) \leq r_{max}(p)$. The circumradius of t is then $d(s,c) \leq d(s,o) + d(o,c) \leq s_{max}(p) + r_{max}(p)$. ∎

The case in which two vertices have strictly positive weights can be divided into three subcases: (1) the larger ball lies on the obtuse angle, (2) the smaller ball lies on the obtuse angle, and (3) the zero-weighted point lies on the obtuse angle. In the first case, the circumcenter is closer to the vertex of the obtuse angle than the orthocenter is, and so the circumradius must be less than $s_{max}(p) + r_{max}(p)$ trivially. In the second case we again bound the angle between edge sq and the plane bisecting the edge su, where q is zero-weighted and $\angle qsu$ is obtuse. For the last case, we can show that the circumcenter of such a triangle must necessarily lie inside the larger of the two protecting balls, thus yielding a circumradius less than $r_{max}(p)$ trivially.

Claim C.10 *The circumradius of obtuse triangle t in $F(p)$ with two positive weighted vertices is at most $s_{max}(p) + r_{max}(p)$ when the smaller protecting ball lies on the obtuse angle.*

Proof: Let triangle t with orthocenter o and circumcenter c have vertices q, s, and u such that $0 = \omega_q < \omega_s \leq \omega_u$, and let Π_{sq} and Π_{su} be the planes supporting the Voronoi facets $V_q \cap V_s$ and $V_u \cap V_s$ respectively. Let $x_{sq} = \Pi_{sq} \cap sq$ and $x_{su} = \Pi_{su} \cap \mathrm{aff}\,(sq)$. Then $\angle x_{su}x_{sq}o = \pi/2$. $d(x_{su}, x_{sq}) > r_{min}(p)$ and $d(o, x_{sq}) \leq s_{max}(p) \leq 0.03r_{min}(p)$, implying $\angle x_{sq}x_{su}o \leq \arctan(0.03) = \beta$. Consider a line segment parallel to edge su with one endpoint at o and the other $y_1 \in \Pi_{b_{su}}$ where $\Pi_{b_{su}}$ is the plane bisecting edge su. Let us construct point c^* as the intersection $\mathrm{aff}\,(t) \cap \Pi_{sq} \cap \Pi_{b_{su}}$. Then the triangle oc^*y_1 is similar to $x_{su}ox_{sq}$ and $d(o,c^*) = d(\Pi_{su}, \Pi_{b_{su}})\sec(\angle x_{sq}x_{su}o) \leq r_{max}(p)\sec(\arctan(0.03))/2$. Now consider a line segment parallel to edge qs with endpoints c and y_2, where y_2 lies on the segment between o and c^*. Then triangle cy_2c^* is similar to $x_{su}x_{sq}o$ and $d(o,c) \leq d(o,c^*) - d(c^*,y_2) + d(y_2,c) \leq r_{max}(p)(\sec(\beta) - tan(\beta) + 1)/2 \leq 0.986r_{max}(p)$. So the circumradius of t is $d(q,c) \leq d(q,o) + d(o,c) \leq s_{max}(p) + r_{max}(p)$. ∎

Claim C.11 *The circumradius of obtuse triangle t in $F(p)$ with two positive weighted vertices is at most $r_{max}(p)$ when the zero-weighted vertex lies on the obtuse angle.*

Proof: Let q, s, and u be the vertices of t such that $\omega_q > 0$ and $\omega_u > 0$. Then by satisfaction of C3 q and u must be adjacent along some ridge $\gamma \in \mathcal{D}_1$. In the proof of the Subridge Covering Lemma, we see that this yields an upper bound on $d(q,u)$ (in terms of ω_q and ω_u) that is dependent on the relationship between b_q and b_u with respect to how and when they were constructed. Given this upper bound, the circumradius of t with edge qu is maximized when s lies as close to qu as possible, and by the Positive Power Distance Lemma this necessarily occurs as s approaches $\mathrm{bd}\,(b_q) \cap \mathrm{bd}\,(b_u)$, giving ω_q and ω_u as the other two edge lengths in the worst case. Below are a list of these relationships and the upper bounds on triangle circumradius that each implies based on the well-known formula for circumradius for a triangle with given edge lengths.

- $\omega_q \geq 4\omega_u$, $d(q,u) \leq \omega_q(1 + 0.25 * 5/12)$. circumradius$(t) \leq 0.582\omega_q$.
- $\omega_q = \omega_u$, $d(q,u) \leq 17\omega_q/12$. circumradius$(t) \leq 0.707\omega_q$.
- $\omega_q = 4\omega_u/5$, $d(q,u) \leq 17\omega_q/12$. circumradius$(t) \leq 0.727\omega_q$.
- $\omega_q = \omega_u$, $d(q,u) \leq 13\omega_q/12$. circumradius$(t) \leq 0.595\omega_q$.

- $\omega_q \geq 16\omega_u/5$, $d(q,u) \leq \omega_q(1 + 5/16 * 1/12)$. circumradius$(t) \leq 0.515\omega_q$.
- $\omega_q \geq 4\omega_u$, $d(q,u) \leq 49\omega_q/48$. circumradius$(t) \leq 0.511\omega_q$.

Each of the above cases has circumradius$(t) \leq \omega_q \leq r_{max}(p)$, thus proving the claim. ■ ■

Lemma C.12 *When $r_{max}(p) \leq 0.01$lfs$_\sigma(p)$ and a ball would be refined in response to a violation of C5 or C6 by point p, the circumradius of each triangle in $F_\sigma(p)$ is less than 0.0103lfs$_\sigma(p)$.*

Proof: This follows somewhat trivially from the conjunction of the previous lemma with the protocol for handling violations of C5 and C6. When a ball would be refined in response to a violation of C5 or C6, $s_{max}(p) < 0.03r_{min}(p) \leq 0.03r_{max}(p)$, and at such a time the circumradius of triangles in $F(p)$ is less than $s_{max}(p) + r_{max}(p)$ by the Bounded Circumradius Lemma. $s_{max}(p) + r_{max}(p) \leq r_{max}(p) + 0.03r_{max}(p)$, which by assumption is less than $0.01\mathrm{lfs}_\sigma(p) + 0.03 * 0.01\mathrm{lfs}_\sigma(p)$. ■

This has certain implications with regards to applying the [9] lemmas mentioned above. In what follows, we will argue that when each protecting ball p has radius less than $0.01\mathrm{lfs}_\sigma(p)$ ($p \in \mathrm{bd}\,(\sigma)$) and each triangle pqs has size less than $\min_{u\in\{p,q,s\}}\{0.03\omega(u)\}$ ($pqs \in F_\sigma(p)$), C5 and C6 cannot be violated. Notice that by the above lemmas this implies triangle circumradii are less than $\max_{u\in\{p,q,s\}}\{0.0103\mathrm{lfs}_\sigma(u)\}$, and by the definition of weighted distance their orthoradii are also less than this value. For Lemmas 2.1-4.2 of [9], the weighted triangulation implies no change because these lemmas do not involve the Voronoi and Delaunay constructs. We find that Lemmas 4.2-4.7 of [9] require some analysis, as they assume that the circumcenter is the orthocenter. These lemmas all rely on the assumption that the distance between a sample point p and the intersection of a Voronoi facet/edge on the boundary of V_p being less than $\alpha\mathrm{lfs}(p)$. Using this assumption and the assumptions that (1) the distance between p and the circumcenter of an incident triangle is less than $\alpha\mathrm{lfs}_\sigma(p)$, and (2) the length of a Delaunay edge incident on p is no more than twice this distance – these two assumptions are direct consequences of the conjunction of the unweighted Delaunay and Voronoi constructs with the previously mentioned assumption. Choosing $\alpha = 0.02$ and adding the assumptions that all protecting balls in $F(p)$ have radii less than $0.01\mathrm{lfs}_\sigma(p)$ and that we are about to refine a ball in $F(p)$, we find that the aforementioned assumptions (1) and (2) will both hold because the circumradius of any triangle in $F(p)$ must be less than $0.0103\mathrm{lfs}_\sigma(p)$ and no Delaunay edge can be longer than twice this distance. This allows us to employ [9] Lemmas 2.1-4.7 as they are.

Lemma C.13 (C5 and C6 Bound) *When each protecting ball b_p has radius less than 0.01lfs$_\sigma(p)$ and each triangle t_{pqs} satisfies size$(t_{pqs}) \leq 0.03\min_{u\in p,q,s}\{\omega_u\}$, C5 and C6 must be satisfied.*

Proof: Satisfaction of C5 and C6 for p implies $F_\sigma(p)$ is a topological disk, and $p \in \mathrm{int}\,(F_\sigma(p))$ iff $p \in \mathrm{int}\,(\sigma)$. Consider how $F_\sigma(p)$ could violate these criteria for $p \in \mathrm{int}\,(\sigma)$, $\sigma \in \mathcal{D}_2$: either a Delaunay edge having p as an endpoint is not incident on exactly two triangles in $F_\sigma(p)$, or $F_\sigma(p) - p$ consists of multiple connected components. In the latter case, $\sigma \cap V_p$ must consist of multiple connected components. In the former case, the errant Delaunay edge is dual to a Voronoi facet F that either has only one Voronoi edge on its boundary intersecting σ, or has more than two Voronoi edges on its boundary intersecting σ. If only one Voronoi edge bounding F intersects σ, then either $\mathrm{bd}\,(\sigma) \cap F \neq \varnothing$ (and therefore $\mathrm{bd}\,(\sigma) \cap V_p \neq \varnothing$), or the Voronoi edge intersects σ in multiple points or tangentially. We can invoke Lemma 4.7 of [9] to contradict a Voronoi edge intersecting σ in multiple points or tangentially when all such intersections lie within distance $0.02\mathrm{lfs}_\sigma(p)$ of p, and this must be the case when all protecting balls satisfy $\omega_p < 0.01\mathrm{lfs}_\sigma(p)$ and restricted triangles satisfy size$(t_{pqs}) < 0.03\omega_p$. Now we show that $Bd(\sigma) \cap V_p = \varnothing$ when $p \in \mathrm{int}\,(\sigma)$, contradicting the possibility of $\mathrm{bd}\,(\sigma)$ intersecting a facet bounding V_p, and allowing us to invoke Lemmas 4.10 (the proof of which requires each connected component of $F \cap \sigma$ to have its two endpoints in distinct edges of V_p, thus implying the requirement $\mathrm{bd}\,(\sigma) \cap \mathrm{bd}\,(V_p) = \varnothing$) and 4.12 (which requires the boundary cycles of $V_p \cap \sigma$ to lie in $\mathrm{bd}\,(V_p)$, thus implying the requirement $\mathrm{bd}\,(\sigma) \cap V_p = \varnothing$) of [9] to contradict σ intersecting more than two edges bounding a Voronoi facet and $\sigma \cap V_p$ having multiple connected components respectively, thus proving that when triangles and protecting balls satisfy the size and radius criteria above C5 and C6 must be satisfied for $p \in \operatorname{int}(\sigma)$.

Claim C.14 *If $p \in \operatorname{int}(\sigma)$, then $Bd(\sigma) \cap V_p = \varnothing$.*

Proof: Consider p violating C5 or C6 such that $p \in \operatorname{int}\sigma$. We know that $\omega_p = 0$ because positive-weighted points are only inserted by COVER and PROTECT, which insert points only in $\mathcal{D}_{\leq 1}$, which is the union of boundaries of the elements of $\mathcal{D}_2$. We claim that $V_p \cap \operatorname{bd}(\sigma) = \varnothing$. Assume the contrary: that $V_p \cap \operatorname{bd}(\sigma) \neq \varnothing$. Then $p =$argmin$_{y \in P}\{d_w(y,x)\}$ for some $x \in \operatorname{bd}(\sigma)$. If $x \in P$ and x is positive-weighted, then $d_w(p,x)^2 = d(p,x)^2 - \omega_x^2 > d(x,x)^2 - 2\omega_x^2 = -2\omega_x^2 = d_w(x,x)^2$, contradicting $p =$argmin$_{y \in P}\{d_w(y,x)\}$; if $x \in P$ and x is zero-weighted, then $d_w(p,x) = d(p,x) > 0 = d(x,x) = d_w(x,x)$, again contradicting $p =$argmin$_{y \in P}\{d_w(y,x)\}$. The only other possibility is $x \notin P$. We know that the ridge $\gamma \ni x$ is covered by protecting balls by the Subridge Covering Lemma, and so $\exists q \in P \cap \gamma$ such that $b_q \ni x$, but then $d_w(q,x) < 0 < d(p,x) = d_w(p,x)$, contradicting $p =$argmin$_{y \in P}\{d_w(y,x)\}$, so it must be that $V_p \cap \operatorname{bd}(\sigma) = \varnothing$, proving the claim. ■

We can now apply Lemma 4.10 from [9], which says that if a smooth surface without boundary intersects a Voronoi facet on the boundary of V_p in multiple disjoint segments then a Voronoi edge on the boundary of the facet must intersect σ at a distance greater than $0.02\text{lfs}_\sigma(p)$ from p, and Lemma 4.12 from [9], which says that when there are multiple boundary cycles of $\sigma \cap V_p$ such that each cycle lies in $\operatorname{bd}(V_p)$ and does not lie fully in a single facet of V_p then there must be an edge of V_p intersecting σ at a distance greater than $0.02\text{lfs}_\sigma(p)$ from p. So when triangles and protecting balls satisfy the size and radius criteria above C5 and C6 must be satisfied for $p \in \operatorname{int}(\sigma)$, because under these conditions an edge of V_p cannot intersect σ at a distance greater than $0.02\text{lfs}_\sigma(p)$ from p.

Consider p violating C5 or C6 such that $p \notin \operatorname{int}\sigma$. We again employ Lemmas 2.1-4.7 of [9] in proving that C5 and C6 must be satisfied when protecting balls and triangles are sufficiently small. Here, satisfaction of C5 and C6 for p implies $F_\sigma(p)$ is a topological disk with p on its boundary. In order for this to be violated, either more than two triangles in $F_\sigma(p)$ must be incident on the same edge or there exist multiple connected components of $F_\sigma(p) - p$ or p lies on the interior of $F_\sigma(p)$. In the case where multiple triangles are incident on the same Delaunay edge e, consider how σ must intersect the Voronoi facet F dual to e. At least three Voronoi edges bounding F must intersect σ, and this implies that there are at least two disjoint curve segments I and I' in $F \cap \sigma$. When at least two such segments exist such that neither is bounded by a point in $\operatorname{bd}(\sigma)$, then we can again employ Lemma 4.10 from [9]. We then have only to show that $F \cap \sigma$ cannot contain multiple disjoint segments when one or more of these segments is bounded by $\operatorname{bd}(\sigma) \cap F$.

Claim C.15 *$F \cap \sigma$ cannot contain multiple disjoint segments when one or more of these segments is bounded by $\operatorname{bd}(\sigma) \cap F$.*

Proof: Let $q = \operatorname{bd}(\sigma) \cap F$. We know that this is a single point when the protecting balls are small enough by the Ridge-Facet Lemma, so let us assume q is a single point. Let us also assume that all points in $\operatorname{bd}(F) \cap \sigma$ lie within 0.02lfs(p) distance of p. Let q be an endpoint of I, let z be the other endpoint of I, and let x and y be the endpoints of I'. By assumption, q is the only point in $\operatorname{bd}(\sigma) \cap F$, so x, y, z must all lie in $\operatorname{bd}(F) \cap \sigma$ and must therefore lie within 0.02lfs(p) of p. By Dey06 Lemma 1.3, the feature sizes lfs(x), lfs(y), lfs$(z) \geq 0.98$lfs(p). We know that the open ball whose boundary lies tangent to σ at x (or y or z) with radius lfs(x) (lfs(y), lfs(z), respectively) must be empty of points in σ, so let us construct two such balls (one on each side of the surface σ) B_x, B'_x for x and two B_y, B'_y for y. Let $C_x = B_x \cap F$, $C'_x = B'_x \cap F$, $C_y = B_y \cap F$, and $C'_y = B'_y \cap F$. By [9] Lemma 4.4(ii) we know that the acute angle between the plane Π containing F and n_x (or n_y or n_z) $\angle_a F, n_x \leq$arcsin(0.02), so radius$(C_x) \geq 0.98\text{lfs}(p)\sqrt{(}1 - 0.02^2) >$ $0.9798\text{lfs}(p)$, and similarly radius(C'_x), radius(C_y), and radius$(C'_y) > 0.9798\text{lfs}(p)$. Let B_x and B_y lie on the same side of σ, and let B'_x and B'_y lie on the opposite side. Because $d(x,p), d(y,p) \leq 0.02\mathrm{lfs}(p)$, we have $d(x,y) \leq 0.04\mathrm{lfs}(p) \leq 0.0409\mathrm{lfs}(x)$, which implies $\angle n_x, n_y \leq 0.0466$ by [9] Lemma 2.1. The normalized vector v_x pointing from x to the center of C_x is merely the normalized projection of n_x onto F along the direction normal to F, and similarly for v_y and n_y, so $\angle v_x, v_y \leq \angle_a F, n_x + \angle n_x, n_y + \angle_a F, n_y \leq 0.0466 + 2(\arcsin 0.02 + 0.0213) \leq 0.13$. This implies that the centers of C_x and C_y lie within $(0.04 + 2 * 1.02\sin(0.13/2))\mathrm{lfs}_\sigma(p) < 0.173\mathrm{lfs}_\sigma(p)$ of one another, and so $C_x \cap C_y \neq \varnothing$. By the same argument, $C'_x \cap C'_y \neq \varnothing$, and so we see that $x \cup C_x \cup C'_x \cup C'_y \cup C_y \cup y$ divides Π into two disjoint connected components: one bounded Π_b and one unbounded Π_u. Because σ has no self-intersections by definition and C_x, C_y, C'_x, C'_y must be empty of points in σ by construction, the entirety of I must lie in Π_b or Π_u, as must the entirety of I'. First consider I'. If I' lies in Π_u then the Voronoi edge e_x containing x must be perpendicular to v_x, and since v_x is nothing more than n_x projected along the normal to Π it follows that $\angle n_x, e_x = \pi/2$ as well, but this contradicts [9] Lemma 4.3, so I' must lie in Π_b. Now consider z. Let z lie in

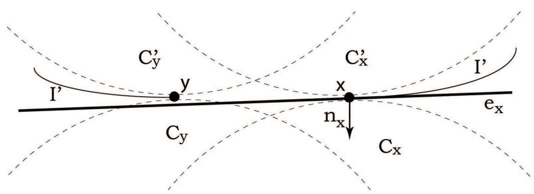

Figure 8: If I' lies in Π_u then e_x is perpendicular to n_x.

Π_b. For every location $u \in \Pi_b \backslash I'$, $\exists w \in I'$ such that $\angle \mathcal{L}_{uw}, v_w = 0$, where $\mathcal{L}_{uw}$ is the line containing both u and w, and v_w is the projection of n_w onto Π. This implies that for some $w \in I'$, $\angle \mathcal{L}_{zw}, v_w = 0$ contradicting Lemma 2.3 of [9] (because we know $\angle n_w, v_w$ is small) and thereby implying $z \notin \Pi_b$. Then let z lie in Π_u. Let us assume that z lies closer to x than to y, and now consider the angle that must be formed by e_x and

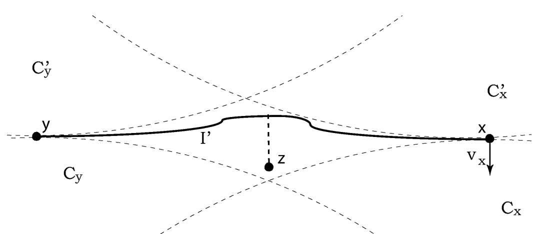

Figure 9: If z lies in Π_b then $\angle \mathcal{L}_{zw}, v_w = 0$ for some $w \in I'$.

n_x. The line supporting e_x divides Π into two half-planes, and x, y, and z must all lie in the closure of one of these half-planes. Since $d(x,y), d(x,z) \leq 0.0409\mathrm{lfs}_\sigma(x)$, $\angle \mathcal{L}_{xy}, n_x, \angle \mathcal{L}_{xz}, n_x \geq 1.55$, and for all three of these points to lie on one side of the supporting line of e_x, $\angle e_x, n_x \geq \min\{\angle \mathcal{L}_{xy}, n_x, \angle \mathcal{L}_{xz}, n_x\} \geq 1.55$. This contradicts [9] Lemma 4.3, which states $\angle e_x, n_x \leq \pi/3$. So $z \notin \Pi$ and therefore $I \not\subseteq F$. This contradicts the assumption that both I and I' lie in F, and so we see that if all points where a Voronoi edge of V_p intersects σ lie within distance $0.02\mathrm{lfs}_\sigma(p)$ of p then there is at most one connected component of $F \cap \sigma$ which also intersects $\mathrm{bd}\,(F)$. ■

One corollary of the above claim is this: an edge of $F_\sigma(p)$ dual to a Voronoi facet intersecting $\mathrm{bd}\,(\sigma)$ will be incident on exactly one triangle in $F_\sigma(p)$, and an edge of $F_\sigma(p)$ dual to a Voronoi facet not intersecting $\mathrm{bd}\,(\sigma)$ will be incident on exactly two triangles in $F_\sigma(p)$. This in conjunction with the Ridge-Facet Lemma imply that, unless there are multiple connected components of $F_\sigma(p) - p$, $F_\sigma(p)$ must be a half disk about p when $p \in \mathrm{bd}\,(\sigma)$ and must be a full disk about p when $p \in \mathrm{int}\,(\sigma)$ – essentially, if C5 is satisfied then C6 must be satisfied. So if we can show that there cannot be multiple connected components of $F_\sigma(p) - p$, then we are done.

Now we must consider the case where there exist multiple connected components of $F_\sigma(p) - p$. This implies that there are multiple connected components of $\sigma \cap V_p$ having non-empty intersections with the edges of V_p. We have already shown results that preclude this possibility when $p \in \operatorname{int}(\sigma)$, so we now consider the case when $p \in \operatorname{bd}(\sigma)$. Let D and D' be two connected components of $V_p \cap \sigma$ and let $p \in D$. This implies that $D \cap \operatorname{bd}(\sigma) \neq \varnothing$. Then let $\gamma_D \ni p$, where γ_D is a connected component of $\operatorname{bd}(\sigma) \cap V_p$. If γ_D were a cycle, then PROTECT would have added multiple weighted points on γ_D, and the centers of all of these balls would lie inside V_p, implying that they also lie in b_p. If b_p is centered on a point in $\mathcal{D}_0$, then it cannot contain the centers of any other protecting balls by satisfaction of C1 and C2. If b_p is not centered on a point in $\mathcal{D}_0$, then it must contain the centers of at least two $\mathcal{D}_0$-balls in γ_D, but this is not possible because we know that, for $p \in \gamma(u,v)$ and $u, v \in \mathcal{D}_0 \cap \gamma_D$, radius$(b_p) \leq 5\min\{\text{radius}_0(b_u),\text{radius}_0(b_v)\}/16 \leq 5d(u,v)/48$, where radius$_0(*)$ is the initial radius of a $\mathcal{D}_0$-ball, and so b_p can contain at most only one such center. So then γ_D is a curve segment bounded by two points in $\sigma \cap \operatorname{bd}(V_p)$. Then by the Ridge-Facet Lemma these two endpoints of γ_D can be the only points in $\operatorname{bd}(\sigma) \cap \operatorname{bd}(V_p)$ for sufficiently small protecting balls. This implies that if $D' \cap \operatorname{bd}(\sigma) \neq \varnothing$, then $D' \cap \operatorname{bd}(\sigma)$ is a cycle or a set of cycles, but then PROTECT would have added multiple weighted points on a given connected component γ'_D of $D' \cap \operatorname{bd}(\sigma)$ and this violates satisfaction of C2, so it must be that $D' \cap \operatorname{bd}(\sigma) = \varnothing$.

So now we must show that D' cannot exist when $D \ni p$, $p \in \operatorname{bd}(\sigma)$, and $D' \cap \operatorname{bd}(\sigma) = \varnothing$. In this case, $\operatorname{bd}(D') = D' \cap \operatorname{bd}(V_p)$. Because all intersections of σ with the edges of V_p lie within distance $0.02\text{lfs}_\sigma(p)$ of p, we can apply Lemma 4.11 of [9] here to show all points in $\operatorname{bd}(D')$ lie within distance $0.02\text{lfs}_\sigma(p)$ of p (since $D' \cap \operatorname{bd}(\sigma) = \varnothing$, no adaptation of the proof of Lemma 4.11 is necessary). Let x be the point in $\operatorname{bd}(D')$ closest to p, and let us construct a plane Π_{xp} containing p and the line $\mathcal{L}_x$, where $\mathcal{L}_x$ is the line through x in the direction of n_x, the normal vector at x. Note that $p \cap \mathcal{L}_x = \varnothing$; otherwise, $\angle n_x, \mathcal{L}_{xp} = 0$, where $\mathcal{L}_{xp}$ is the line containing both x and p, and this contradicts Lemma 2.3 of [9]. Note that $\Pi_{xp} \cap \operatorname{bd}(D')$ cannot contain a curve segment, as this implies that Π_{xp} contains a facet of V_p, and this further implies that p lies on $\operatorname{bd}(V_p)$ or outside of V_p, both of which are impossible when no sample point lies inside the protecting ball about another sample point. The proof below can be summarized as follows: we show that if x is not an endpoint bounding a curve segment of $\operatorname{int}(D') \cap \Pi_{xp}$, then x must lie on a Voronoi edge, but then there must be some curve segment of $\operatorname{int}(D') \cap \Pi_{xp}$ bounded by x and one other point, a contradiction; if, on the other hand, x is on the boundary of some curve segment of $\operatorname{int}(D') \cap \Pi_{xp}$, then either x is not the closest point in $\operatorname{bd}(D')$ to p, contradicting the selection of x, or $\angle n_p, \mathcal{L}_{yp}$ is very small for some y in this curve segment, contradicting Lemma 2.3 of [9].

Claim C.16 *x cannot be isolated in $\Pi_{xp} \cap D'$.*

Proof: Assume that x is isolated in $\Pi_{xp} \cap D'$, meaning that it does not bound any curve in $\Pi_{xp} \cap \operatorname{int}(D')$. This implies that $\operatorname{bd}(D')$ intersects Π_{xp} tangentially. Assume $x \in \operatorname{int}(F_x)$, where F_x is a facet of V_p, and let Π_{Fx} be the plane containing F_x. $\angle_a \Pi_{Fx}, n_x \leq 0.0412$ by Lemma 4.4 of [9], so the projection $P_{\Pi_{Fx}}(\mathcal{L}_x)$ of $\mathcal{L}_x$ onto Π_{Fx} along the normal to Π_{Fx} satisfies $\angle P_{\Pi_{Fx}}(\mathcal{L}_x), n_x \leq 0.0412$. We proceed by showing that $\angle \Pi_{Fx} \cap \Pi_{xp}, P_{\Pi_{Fx}}(\mathcal{L}_x)$ is small. Since $\operatorname{bd}(D')$ intersects Π_{xp} tangentially, there exist two points $y, y' \in \operatorname{bd}(D') \cap F_x$ such that x lies on the path from y to y' in $\operatorname{bd}(D') \cap F_x$, $d(x,y) \leq 0.01d(x,p)$, $d(x,y') \leq 0.01d(x,p)$, and y and y' lie on a line parallel to $\Pi_{Fx} \cap \Pi_{xp}$. This will imply that $\angle \mathcal{L}_{yy'}, n_y$ is small, which is contradicted by [9] Lemma 2.3. We first find the minimal $\angle \Pi_{Fx} \cap \Pi_{xp}, \mathcal{L}_{xy}$, and then find the maximal $\angle \mathcal{L}_{xy}, P_{\Pi_{Fx}}(\mathcal{L}_x)$; because $\mathcal{L}_{xy}$, $P_{\Pi_{Fx}}(\mathcal{L}_x)$, and $\Pi_{Fx} \cap \Pi_{xp}$ all lie in Π_{Fx}, we can conclude that $\angle \Pi_{Fx} \cap \Pi_{xp}, P_{\Pi_{Fx}}(\mathcal{L}_x)$ is the difference between these two angles. Before proceeding, we must establish some angle bounds between certain lines of interest: $\mathcal{L}_{xp}$, $\mathcal{L}_{xy}$, $\mathcal{L}_x$. $\text{lfs}_\sigma(x) \geq 0.98\text{lfs}_\sigma(p)$ because $d(x,p) \leq 0.02\text{lfs}_\sigma(p)$, so by [9] Lemma 2.3 we have $\angle \mathcal{L}_x, \mathcal{L}_{xp} \geq \arccos(0.02/(2 * 0.98))$. From $d(y,p) \geq d(x,p)$

and $d(x, y) \leq 0.01d(x, p)$, we have $\angle\mathcal{L}_{xp}, \mathcal{L}_{xy} \geq \arccos(0.01/2)$. Again, using [9] Lemma 2.3, we find $\angle\mathcal{L}_x, \mathcal{L}_{xy} \geq \arccos(0.01 * 0.02/(2 * 0.98))$. To find $\angle\Pi_{Fx} \cap \Pi_{xp}, \mathcal{L}_{xy}$, we compute the angle between $\mathcal{L}_{xy}$ and the normal to the plane Π_{xp}: $\angle_a\mathcal{L}_{xy}, \Pi_{xp} = \pi/2 - \angle_{max}\vec{xy}, \vec{n_x} \times \vec{xp} \geq 1.559$, where $\vec{xy}$, $\vec{n_x}$, and $\vec{xp}$ are unit vectors in the directions of $\mathcal{L}_{xy}$, $\mathcal{L}_x$, and $\mathcal{L}_{xp}$ respectively, found using the angle constraints between these lines discussed above. The obtuse angle, $\angle\mathcal{L}_{xy}, P_{\Pi_{Fx}}(\mathcal{L}_x) \geq (\pi - \angle\mathcal{L}_{xy}, \mathcal{L}_x) + \angle_a\mathcal{L}_x, \Pi_{Fx} \leq 1.62$, yields $\angle\Pi_{Fx} \cap \Pi_{xp}, P_{\Pi_{Fx}}(\mathcal{L}_x) \leq 1.62 - 1.559 = 0.061$. By [9] Lemma 2.1, $\angle n_x, n_y \leq 0.015$, implying $\angle n_y, \mathcal{L}_{yy'} = \angle n_y, \Pi_{Fx} \cap \Pi_{xp} \leq \angle n_x, n_y + \angle P_{\Pi_{Fx}}(\mathcal{L}_x), n_x + \angle P_{\Pi_{Fx}}(\mathcal{L}_x), \Pi_{Fx} \cap \Pi_{xp} \leq 0.12$, where $\mathcal{L}_{yy'}$ is the line passing through y and y', but this contradicts [9] Lemma 2.3. So if $x \in \mathrm{int}\,(F_x)$ then $\mathrm{bd}\,(D')$ intersects Π_{xp} transversely, but this contradicts assumption, so $x \notin \mathrm{int}\,(F_x)$, implying that it lies on some Voronoi edge e_x, and by [9] Lemma 4.7 $\mathrm{bd}\,(D')$ cannot meet e_x tangentially. Then the only way for x to be isolated is for x to lie on a Voronoi edge e_x such that $\mathrm{bd}\,(D')$ meets Π_{xp} tangentially at x, but traverses e_x at x (meaning that there are two facets $F_z, F_{z'}$ incident on e_x such that for any segment of $\mathrm{bd}\,(D')$ containing x on its interior, some part of that segment must lie in F_z and another part must lie in F'_z, and the union of these two "subsegments" forms a single connected component containing x). The remainder of the proof of this claim shows that this cannot happen.

Now consider the two facets $F_z, F_{z'}$ of V_p incident on e_x such that $x \in \mathrm{bd}\,(D' \cap F_z)$ and $x \in \mathrm{bd}\,(D' \cap F_{z'})$. Also, let us choose $z \in F_z \cap D'$ and $z' \in F_{z'} \cap D'$ such that $d(x, z) = d(x, z') \leq 0.01d(x, p)$. $z, z' \in \mathrm{bd}\,(D')$, and so by construction $d(p, x) \leq d(p, z)$ and $d(p, x) \leq d(p, z')$. This yields $\angle\mathcal{L}_{xz}, \mathcal{L}_{xp} \geq \arccos(0.5d(x, z)/d(x, p)) = \arccos(0.005) > 1.57$ and $\angle\mathcal{L}_{xz}, \mathcal{L}_x \geq \arccos(0.01 * 0.02/(2 * 0.98)) > 1.57$, which further implies $\angle\Pi_{Fz}, \Pi_{xp} > 1.57$ because $\mathcal{L}_{xz} \subset \Pi_{Fz}$. We now want to constrain the angle between $\mathcal{L}_x$ and the intersection line $\Pi_{Fz} \cap \Pi_{xp}$. Π_{Fz} contains $\mathcal{L}_{e_x}$, and $\angle\mathcal{L}_{e_x}, \mathcal{L}_x < 0.133$, implying $\angle\mathcal{L}_{e_x}, \Pi_{xp} < 0.133$. This yields $\angle\mathcal{L}_x, \Pi_{Fz} \cap \Pi_{xp} < 0.1542$. The same will be true of $\angle\mathcal{L}_x, \Pi_{Fz'} \cap \Pi_{xp}$. These two intersection lines, $\Pi_{Fz} \cap \Pi_{xp}$ and $\Pi_{Fz'} \cap \Pi_{xp}$, divide Π_{xp} unevenly into four quadrants Q_1, Q_2, Q_3, Q_4, as shown in Figure 10. If z and z' lie on the same side of Π_{xp}, then $\mathrm{int}\,(V_p) \cap \Pi_{xp}$ lies in $Q_1 \cup Q_3$. $\angle\mathcal{L}_{xp}, \mathcal{L}_x > 1.56$ by Lemma 2.3 of [9], but all points $y \in V_p \cap \Pi_{xp} \subseteq Q_1 \cup Q_3$ have $\angle\mathcal{L}_{xy}, \mathcal{L}_x < 0.1542$, so V_p must be a subset of $Q_2 \cup Q_4$, but then z and z' lie on opposite sides of Π_{xp} and so $\mathrm{bd}\,(D')$ intersects Π_{xp} transversely at x, a contradiction. So x cannot be isolated in $\Pi_{xp} \cap D'$. ■

Claim C.17 *x cannot bound a curve segment in $\Pi_{xp} \cap \mathrm{int}\,(D')$.*

Proof: Assume x is on the boundary of curve segment $s_x \subset \Pi_{xp} \cap \mathrm{int}\,(D')$. Then we must have one of the following: s_x is a non-degenerate cycle, s_x has another endpoint $z \neq x$. If s_x is a cycle, then $\exists y, y' \in s_x$ such that $d(x, y) \leq 0.01\mathrm{lfs}_\sigma(x)$, $d(x, y') \leq 0.01\mathrm{lfs}_\sigma(x)$, and y and y' lie on a line parallel to $\mathcal{L}_x$. Then, as previously, we see that this cannot satisfy both Lemma 2.3 and Lemma 2.1 of [9], so s_x cannot be a cycle. Consider then the case when s_x has another point z in its boundary. Then $z \in \mathrm{bd}\,(D')$ and so $d(x, z) \leq 0.04\mathrm{lfs}_\sigma(p) < 0.042\mathrm{lfs}_\sigma(x)$. By Lemma 2.3 of [9], $\angle\mathcal{L}_x, \mathcal{L}_{xz} > 1.54$; the region whose constituent points meet this angle requirement (so, the region containing all candidate points for z) is depicted in Figure ?? as $T_L \cup T_R$. As above, we can show that any Voronoi facet F_x of V_p containing x satisfies $\angle\mathcal{L}_x, \Pi_{Fx} \cap \Pi_{xp} < 0.1542$, and the region containing all such lines of intersection is labeled as T_F. One may now observe that any line in T_F divides Π_{xp} into two half-spaces, one of which contains T_R and the other of which contains T_L. This implies that only one of T_R and T_L can lie in V_p. Assume $T_L \subset V_p$. Then $p, z \in T_L$. Let us now make a smaller triangle $T'_L \subset T_L$ by cutting off T_L at a line passing through p and parallel to $\mathcal{L}_x$. Any point in T'_L is closer to p than x, contradicting our choice of x if we allow $z \in T'_L$, so let $z \in T_L \setminus T'_L$. Then s_x traverses the line through p parallel to $\mathcal{L}_x$. Let us call the point at which this traversal occurs y. Then we have $\angle\mathcal{L}_{py}, n_p = \angle\mathcal{L}_x, n_p = \angle n_x, n_p \leq 0.022$. On the other hand, $d(p, y) < 0.0\mathrm{lfs}_\sigma(p)$, so by Lemma 2.3 of [9] we have $\angle n_p, \mathcal{L}_{py} > 1.54$, but this contradicts $\angle\mathcal{L}_{py}, n_p \leq 0.022$ that we just found. So x does not bound a curve segment in $\Pi_{xp} \cap \mathrm{int}\,(D')$. ■

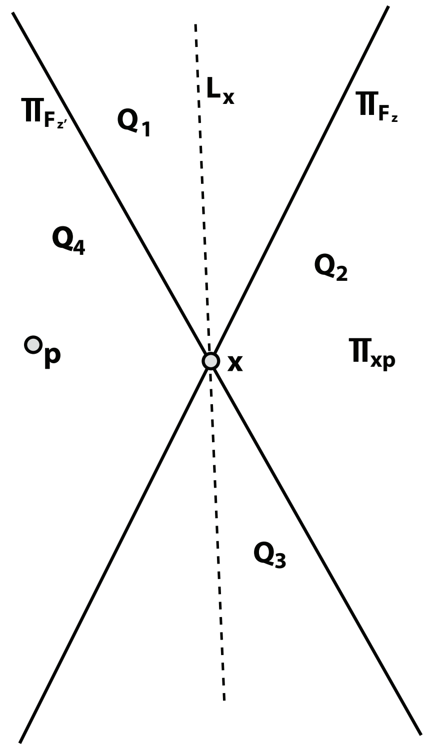

Figure 10: Depiction of the four quadrants dividing Π_{xp} in Claim C.16.

Claims C.17 and C.16 together imply that x cannot exist, meaning that there is no closest point $x \in D'$ to $p \in D$ and therefore that there is no D'. This further implies that $F_\sigma(p) - p$ consists of only one connected component, forcing $F_\sigma(p)$ to be a disk.

Note that although all of the above claims rely on the circumradius of each triangle t_{pqs} being less than size$(t_{pqs}) + \max_{u \in p,q,s}\{\omega_u\}$, this condition is necessarily satisfied before any of the protecting balls on t_{pqs} can be refined (as discussed in the Bounded Circumradius Lemma). Since size$(t_{pqs}) < 0.03\omega_p$ is also necessarily satisfied at this same time, then for any triangle with incident protecting ball p of radius less than $0.01\text{lfs}_\sigma(p)$, all of the above claims must be satisfied before b_p can be refined and so b_p is not refined because C5 and C6 are satisfied. It follows that t_{pqs} is not refined because size$(t_{pqs}) < 0.03\omega_p$. This completes the proof that C5 and C6 must be satisfied for sufficiently small triangles and protecting balls. ■

Lemma C.18 *When all triangles have size less than λ, C7 must be satisfied.*

Proof: Trivially, one may observe that a triangle with size less than λ will not be refined by C7. ■

D Ball and Triangle Size Bounds

Lemma D.1 *$\mathcal{D}_0$-ball size depends only on surface geometry.*

Proof: In the C1 Bound proof, we find that a $\mathcal{D}_0$-ball with radius less than $\min\{0.06\text{lfs}_{min}(\gamma), D_{min,0,1}\}$ will not be refined, as they cannot violate C1. In the C2 Bound proof, we find that $\mathcal{D}_0$-balls with radii less than $D_{min,0,1}/2$ will not be refined, as they cannot violate C2.

In the C3 Bound proof, we find a $\mathcal{D}_0$-ball b_p with radius ω_p less than $\min\{0.06\text{lfs}_{min}(\gamma), D_{min,0,1}/4\}$ will not be refined, as it cannot violate C3 when incident triangles have orthoradii less than $0.03\omega_p$, and this point must be reached before we refine b_p.

In the C4 Bound proof, we find a $\mathcal{D}_0$-ball b_p with radius ω_p less than $D_{min,0,2}/3$ will not be refined, as it cannot violate C4 when its incident triangles have orthoradii less than ω_p, and this must necessarily be the case for b_p to be refined.

In the C5 and C6 Bound proof, we find a $\mathcal{D}_0$-ball b_p with radius less than $0.01\text{lfs}_\sigma(p)$, $\sigma \ni p$, will not be refined, as it cannot violate C5 or C6 when incident triangles have orthoradii less than $0.03\omega_p$, and this point must be reached before we refine b_p.

The reader should note that each of the terms above is a constant dependent only on surface geometry; thus, $\mathcal{D}_0$-ball size depends only on surface geometry. ■

Lemma D.2 *$\mathcal{D}_1$-ball size depends only on surface geometry.*

Proof: In the C1 Bound Lemma, we observe that there is no relationship between C1 satisfaction and $\mathcal{D}_1$-ball size. In the C2 Bound Lemma, we find that $\mathcal{D}_1$-balls with radii less than $\min\{D_{min,0,1}, D_{min,1,1}\}/2$ will not be refined, as they cannot violate C2.

In the C3 Bound Lemma, we find a $\mathcal{D}_1$-ball b_p with radius ω_p less than $\min\{0.06\text{lfs}_{min}(\gamma), D_{min,0,1}, D_{min,1,1}\}/4$ will not be refined, as it cannot violate C3 when incident triangles have orthoradii less than $0.03\omega_p$, and this point must be reached before we refine b_p.

In the C4 Bound Lemma, we find a $\mathcal{D}_1$-ball b_p with radius ω_p less than $D_{min,1,2}/3$ will not be refined, as it cannot violate C4 when its incident triangles have orthoradii less than ω_p, and this must necessarily be the case for b_p to be refined.

In the C5 and C6 Bound Lemma, we find a $\mathcal{D}_1$-ball b_p with radius less than $0.01\text{lfs}_\sigma(p)$, $\sigma \ni p$, will not be refined, as it cannot violate C5 or C6 when incident triangles have orthoradii less than $0.03\omega_p$, and this point must be reached before we refine b_p.

The reader should note that each of the terms above is dependent on either surface geometry or $\mathcal{D}_0$-ball size, and recall that minimum $\mathcal{D}_0$-ball size is dependent on only surface geometry; ergo, $\mathcal{D}_1$-ball size depends only on surface geometry. ■

Lemma D.3 *Triangle size depends only on surface geometry.*

Proof: In the C1 Bound and C2 Bound Lemmas, we observe that there is no relationship between C1/C2 satisfaction and triangle size.

By our violation handling protocol for C3, we find a triangle t with size less than $0.03\omega_p$ (where t is incident to $\mathcal{D}_{\leq 1}$ point p) will not be refined. By our violation handling protocol for C4, we find a triangle t with size less than $0.03\omega_p$ (where t is incident to $\mathcal{D}_{\leq 1}$ point p) will not be refined.

By our violation handling protocol for C6, we find that a triangle t with size less than $\min\{0.03\omega_p, d_{min}(p)\}$ (where t is incident on $\mathcal{D}_{\leq 1}$ point p) will not be refined. As the minimum ω_p depends only on surface geometry, it is only necessary to show that at some point size$(t) < d_{min}(p)$. Assume otherwise; then the refinement of triangles in $F(p)$ always yields new triangles with size greater than $d_{min}(p)$. The moment after a triangle refinement, it must be that the new $d_{min}(p) = \min\{d'_{min}(p),\text{size}(t_p)\}$, where $d'_{min}(p)$ is the value of $d_{min}(p)$ just before the refinement and t_p is the triangle in $F(p)$ that was just refined. This implies that every triangle in $F(p)$ that we refine has size greater than or equal to a constant value $d'_{min}(p)$, and this value is not diminished by any number of iterations of refinement. Then we can insert an infinite number of points into our sample set by this method, each of which is at least distance $d'_{min}(p)$ from all other points, but this is impossible for a bounded domain such as $\mathcal{D}$, and so we have a contradiction. So at some point it must be that size$(t) < d_{min}(p)$, and we will not refine t in response to a violation of C5 or C6.

The reader should note that the minimum ω_p is dependent on only surface geometry. Furthermore, the point at which size$(t) < d_{min}(p)$ is also dependent on surface geometry. So triangle size depends only on surface geometry. ■

E Extra Results

model	λ	**V**ersion	#Vertices (thousand)	**m**em (MB)	**T**ime (sec.)
Lucy	0.004	LocPSC: $\kappa = 1k$	371	114	776
	0.004	CGAL	366	712	1383
Bimba	0.0025	LocPSC: $\kappa = 1k$	438	211	894
	0.0025	CGAL	433	910	994
Fertility	0.005	LocPSC: $\kappa = 1k$	302	215	712
	0.005	CGAL	296	673	725
Fertility	0.0035	LocPSC: $\kappa = 1k$	611	276	1340
	0.0035	CGAL	604	1253	1496
3Holes	0.005	LocPSC: $\kappa = 1k$	400	86	633
	0.005	CGAL	389	714	1408
3Holes	0.003	LocPSC: $\kappa = 1k$	1106	246	1973
	0.003	CGAL	1080*	1922*	4369*
3Holes	0.0012	LocPSC: $\kappa = 1k$	6935	1335	16495
	0.0012	CGAL	NA	NA	NA

Table 2: Time and memory usage for different models for CGAL (CGAL 3.8, release mode, -O3 optimization) and LocPSC results for $\kappa = 1000$. Number of vertices expressed in thousand unit($\times 1000$); NA indicates that an experiment could not be completed due to memory constraints; starred results indicate significant memory thrashing.