HKUST Theoretical Computer Science Center Research Report HKUST-TCSC-2004-01

Three-Dimensional Delaunay Mesh Generation*†

Siu-Wing Cheng‡ Sheung-Hung Poon‡

February 20, 2004

Abstract

We propose an algorithm to compute a conforming Delaunay mesh of a bounded domain specified by a piecewise linear complex. Arbitrarily small input angles are allowed, and the input complex is not required to be a manifold. Our algorithm encloses the input edges with a small buffer zone, a union of balls whose sizes are proportional to the local feature sizes at their centers. In the output mesh, the radius-edge ratio of the tetrahedra outside the buffer zone is bounded by a constant independent of the domain, while that of the tetrahedra inside the buffer zone is bounded by a constant depending on the smallest input angle. Furthermore, the output mesh is graded. Our work is the first that provides quality guarantees for Delaunay meshes in the presence of small input angles.

1 Introduction

In finite element analysis, a polyhedral domain needs to be partitioned into a cell complex for the purpose of numerical simulation and analysis [11]. The mesh is required to be *conforming*: each input edge appears as the union of some edges in the mesh and each input facet appears as the union of some triangles in the mesh. We assume that the input domain is a bounded volume in $\mathbb{R}^3$ specified by a *piecewise-linear complex*.

Delaunay tetrahedralizations are popular tetrahedral meshes in theory and practice [10, 11]. The geometric quality of a Delaunay mesh is often measured by the shape of the tetrahedra, edge lengths, and the mesh size. A tetrahedron τ is well-shaped if its aspect ratio is upper bounded by a constant. A weaker measure is the *radius-edge ratio* $\rho(\tau)$ which is the ratio of the circumradius of τ to its shortest edge length. Radius-edge ratio is a fairly good indicator of the tetrahedral shape. If the radius-edge ratio is bounded, almost all tetrahedra have bounded aspect ratio except for a class known as *slivers* [4]. Nevertheless, bounded radius-edge ratio works well in some applications [13]. A mesh is *graded* if the shortest edge incident to every mesh vertex v has length at least a constant factor of the local feature size at v. Gradedness is instrumental in proving the optimality of mesh sizes when there is no sliver [3, 12, 16]. If the input domain is convex, it also follows from gradedness and bounded radius-edge ratio that the vertex degree is bounded by a constant, the lengths of two adjacent edges are within a constant factor, and the dual Voronoi complex has bounded aspect ratio [18].

An important challenge in mesh generation is to construct a mesh with good quality. Although such meshing algorithms based on quadtree and octtree are known [1, 14], it remains an open problem how to compute a conforming Delaunay mesh with good quality. We briefly survey the previous results on this problem below. Ruppert [16] proposed the Delaunay refinement algorithm to mesh a two-dimensional polygonal domain. The mesh is graded, every triangle has bounded aspect ratio, and the size of the mesh is asymptotically optimal (i.e., within a constant factor of the size of any mesh with bounded aspect ratio). Shewchuk [17] extended Delaunay refinement to three dimensions for polyhedral

*A preliminary version of this paper appeared in Proc. 14th Annual ACM-SIAM Sympos. Discrete Alg., 2003, 295–304.

†This work was supported by the Research Grant Council, Hong Kong, China (HKUST6088/99E and HKUST6190/02E).

‡Department of Computer Science, HKUST, Clear Water Bay, Hong Kong. Email: {scheng,hung}@cs.ust.hk

domains. A graded conforming Delaunay mesh is obtained but there are two differences. First, the algorithm may not terminate when some input angle is less than $\pi/2$. Second, the radius-edge ratio of the tetrahedra is bounded by a constant, but there may be slivers.

Recently, methods have been discovered to eliminate slivers when every input angle is at least $\pi/2$. Li and Teng [12] improved Delaunay refinement with a random point-placement strategy in line of Chew [6]. Cheng et al. [4] introduced *sliver exudation* to eliminate slivers from a Delaunay mesh of a periodic point set with bounded radius-edge ratio. Cheng and Dey [3] introduced *weighted Delaunay refinement* which extends sliver exudation to handle boundaries. Both algorithms by Li and Teng [12] and Cheng and Dey [3] produce a graded conforming Delaunay mesh with bounded aspect ratio and asymptotically optimal size.

Much less is known about handling polyhedral domains with input angles less than $\pi/2$. Murphy et al. [15] showed the existence of a conforming Delaunay mesh, but their method produces tetrahedra of poor shape and unnecessarily many vertices. Cohen-Steiner et al. [7] proposed an improved method and they experimentally studied the effectiveness of their algorithm. In the above results, gradedness and bounded radius-edge ratio are not guaranteed. It is sometimes unavoidable that the edge lengths and the shape of tetrahedra deteriorate near a small input angle. Thus it is conceivable that there are lower bound on edge lengths and upper bound on radius-edge ratio that use constant factors depending on the input angle. Nevertheless, no such result is known till now.

We present an algorithm MESH that constructs a conforming Delaunay mesh of a bounded domain specified by a piecewise linear complex. Arbitrarily small input angles are allowed, and the input complex is not required to be a manifold. So MESH can handle a wider class of input than polyhedra, for example, domains in which three or more triangles are incident on the same edge. Let $\mu \in (0, 1/7]$ and $\rho_0 > 16$ be two a priori chosen constants. Our algorithm encloses the input edges with a small buffer zone, a union of balls whose sizes are proportional to the local feature sizes at their centers. The constant of proportionality is less than 1 and depends on μ. For every tetrahedron τ in the output mesh, if τ does not lie inside the buffer zone, its radius-edge ratio $\rho(\tau)$ is at most ρ_0; otherwise $\rho(\tau)$ is bounded by a constant that depends on μ and ϕ. The shortest edge incident to a mesh vertex v has length at least a constant factor of the local feature size at v where the constant depends on μ and ϕ. Our work is the first that provides quality guarantees for Delaunay meshes in the presence of small input angles.

The rest of the paper is organized as follows. Section 2 gives the basic definitions. Section 3 describes an overview of our algorithm, and the augmentation of the input complex with the buffer zone. Section 4 describes MESH. Sections 5, 6, and 7 prove that MESH terminates and the output mesh is Delaunay and conforming. Sections 8 prove the gradedness and the radius-edge ratio bound. We conclude in Section 9.

2 Preliminaries

We use $\mathcal{P}$ to denote the input piecewise linear complex. The elements of $\mathcal{P}$ are vertices, edges and facets that intersect properly. That is, the intersection of two elements is either empty or an element of $\mathcal{P}$. The boundary of each facet consists of one or more disjoint simple polygonal cycles. Two elements of $\mathcal{P}$ are *adjacent* if their intersection is non-empty. Two elements of $\mathcal{P}$ are *incident* if one is a boundary element of the other. Since $\mathcal{P}$ represents a bounded domain, a subset of facets form an outer boundary (i.e., a closed 2-manifold) that encloses all other elements of $\mathcal{P}$. Other than the above requirements, $\mathcal{P}$ can be quite arbitrary. For example, we allow isolated vertices, isolated edges, and an arbitrary number of triangles sharing an edge.

There are three types of input angles. First, angles between adjacent edges. Second, take an edge uv and a facet F such that $u \in \partial F$ and uv and F are non-coplanar. Let L be the plane through uv perpendicular to F. The *angle between uv and F* is $\min\{\angle puv : p \neq u,\ p \in L \cap F\}$. Third, take two adjacent and non-coplanar facets F_1 and F_2. Let H_i be the supporting plane of F_i. For each point $u \in H_1 \cap H_2$, let L_u be the plane through u perpendicular to $H_1 \cap H_2$. The *angle between F_1 and F_2* is $\min_{u \in H_1 \cap H_2}\{\angle puq : p \neq u,\ q \neq u,\ p \in L_u \cap F_1,\ q \in L_u \cap F_2\}$. Throughout this paper, ϕ denotes the smallest angle in the domain measured as described above. We assume that

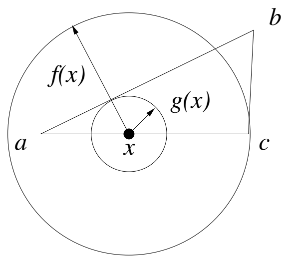

Figure 1: The large and small circles have radii $f(x)$ and $g(x)$ respectively.

$\phi < \pi/2$ as the other case has been solved.

For a point x, the *local feature size* $f(x)$ is the radius of the smallest ball centered at x that intersects two disjoint elements of $\mathcal{P}$. Local feature sizes satisfy the *Lipschitz property*: $f(x) \leq f(y) + \|x - y\|$ for any two points x and y. It is inconvenient to use local feature sizes directly when handling domains with acute angles. So we define the *local gap size* $g(x)$ which is the radius of the smallest ball centered at x that intersects two elements of $\mathcal{P}$, at least one of which does not contain x. Figure 1 illustrates local feature and gap sizes. Clearly, $g(x) \leq f(x)$ and for each vertex v of $\mathcal{P}$, $g(v) = f(v)$. Moreover, we can prove that $g(x) = \Omega(f(x))$ for the vertices of the final mesh. The local gap size is not a continuous function, however, a Lipschitz-like property holds under certain conditions as stated in the following lemma.

LEMMA 2.1 *Let e be an edge of $\mathcal{P}$. If x and y are two points in e such that $x \in \mathrm{int}(e)$, then $g(x) \leq g(y) + \|x - y\|$.*

Proof. Let B be the ball centered at x with radius $g(y) + \|x - y\|$. Thus B intersects the two elements of $\mathcal{P}$ that define $g(y)$. By definition, y does not lie on one of these elements which we denote by E. Since $y \in e$, E is not an incident facet of e. As $x \in \mathrm{int}(e)$, x lies on e and its incident facets only. So $x \notin E$ which implies that $\mathrm{radius}(B) \geq g(x)$. $\square$

We need concepts including *weighted distance* and *orthogonality* that are instrumental to obtaining our results. Let S and S' denote two spheres centered at p and q respectively. The weighted distance $\pi(S, S')$ is defined as $\|p - q\|^2 - \mathrm{radius}(S)^2 - \mathrm{radius}(S')^2$. The weighted distance $\pi(x, S)$ between a point x and S is defined the same way by treating x as a sphere of zero radius. S and S' are *orthogonal* if $\pi(S, S') = 0$. In this case, S and S' intersect and for any point $x \in S \cap S'$, the normal to S at x is tangent to S'. That is, S and S' intersect at right angle. If S and S' are orthogonal, p lies outside S' and q lies outside S. The points at equal weighted distances from S and S' lie on a plane. We call it the *bisector plane* of S and S'. The bisector plane is perpendicular to the line through p and q. If S and S' intersect, their bisector plane is the plane containing the circle $S \cap S'$.

3 Augmenting $\mathcal{P}$

We compute spheres centered at points on edges of $\mathcal{P}$. The spheres are judiciously chosen so that adjacent ones are orthogonal. We use $\mathcal{B}$ to denote the outer boundary of the union of these spheres. The space inside $\mathcal{B}$ is the *buffer zone*. The idea is to mesh the space outside $\mathcal{B}$ such that the tetrahedralization of the space inside $\mathcal{B}$ is automatically induced. $\mathcal{P}$ is augmented with the boundary of $\mathcal{B}$ to yield a new complex $\mathcal{Q}$ for our algorithm to work on. Since adjacent spheres are orthogonal, the space outside $\mathcal{B}$ has non-acute angle, thus allowing the use of Delaunay refinement. Since $\mathcal{Q}$ contains some curved edges and facets, it is impossible to produce a tetrahedral mesh that conforms to these curved elements of $\mathcal{Q}$. Instead, our algorithm will produce a mesh that approximates $\mathcal{Q}$ and conforms to the elements of $\mathcal{P}$. There are two difficulties to overcome. First, we need to guarantee that unnecessarily short edges are not forced when constructing $\mathcal{B}$. Second, we need a method to triangulate the curved boundary of $\mathcal{B}$. This will be addressed in Section 4. This section describes the construction of $\mathcal{B}$ and $\mathcal{Q}$. Note that we need not do anything with isolated vertices. In particular, $\mathcal{B}$ does not enclose isolated vertices. Since isolated vertices are not incident on any input angle, they do not cause any problem for applying Delaunay refinement to the space outside $\mathcal{B}$.

3.1 Protecting spheres

Let μ be some fixed constant chosen from $(0, \frac{1}{7}]$. We create a set of *protecting spheres* with centers lying on the edges of $\mathcal{P}$. First, for each non-isolated vertex v of $\mathcal{P}$, we create a sphere S_v with center v and radius $\mu \cdot g(v)$. Second, for each edge uv of $\mathcal{P}$, we create two protecting spheres S_{u_v} and S_{v_u} with centers u_v and v_u on uv as follows. Let ϕ^u_{uv} be the smallest angle between uv and an edge/facet of $\mathcal{P}$ incident to u. ϕ^v_{uv} is symmetrically defined. Define $\theta^u_{uv} = \min\{\pi/3, \phi^u_{uv}\}$ and $\theta^v_{uv} = \min\{\pi/3, \phi^v_{uv}\}$. The positions of u_v and v_u and the radii of S_{u_v} and S_{v_u} are:

$$\begin{aligned}
\|u - u_v\| &= \mu \sec(\mu\theta^u_{uv}) \cdot g(u) \\
\text{radius}(S_{u_v}) &= \|u - u_v\| \cdot \sin(\mu\theta^u_{uv}) \\
\|v - v_u\| &= \mu \sec(\mu\theta^v_{uv}) \cdot g(v) \\
\text{radius}(S_{v_u}) &= \|v - v_u\| \cdot \sin(\mu\theta^v_{uv})
\end{aligned}$$

Figure 2 illustrates the construction of S_{u_v}. Note that S_u and S_{u_v} are orthogonal and so are S_v and S_{v_u}. Lemma 3.1 bounds the radii of S_{u_v} and S_{v_u}.

LEMMA 3.1 *Let uv be an edge of $\mathcal{P}$. S_{u_v} and S_{v_u} are orthogonal to S_u and S_v respectively. The two ratio* $\text{radius}(S_{u_v})/g(u_v)$ *and* $\text{radius}(S_{v_u})/g(v_u)$ *lie in* $[c_2\mu, c_1\mu]$*, where* $c_1 = 2\pi/(3\sqrt{3})$ *and* $c_2 = \min\{\sqrt{3}/2, \sin\phi\}$.

Proof. S_{u_v} and S_{v_u} are orthogonal to S_u and S_v respectively by construction. Let B be the ball centered at u_v with radius $g(u_v)$. Let E be an element of $\mathcal{P}$ such that $u_v \notin E$ and E touches B. Let d be the minimum distance between u and E. By triangle inequality, $d \leq \|u - u_v\| + g(u_v)$ which is at most $2 \cdot \|u - u_v\|$ as $g(u_v) \leq \|u - u_v\|$. By the definition of $\|u - u_v\|$, we get $d \leq 2\mu \sec(\mu\theta^u_{uv}) \cdot g(u)$. Since $2\mu \leq 2/7 < \cos(\pi/21) \leq \cos(\mu\theta^u_{uv})$, $d < g(u)$ which implies that u lies on E. So either $E = u$ or E is an edge/facet incident to u.

We claim that $\|u - u_v\| \cdot \sin\phi^u_{uv} \leq g(u_v) \leq \|u - u_v\|$. If $E = u$, then $g(u_v) = \|u - u_v\|$ and our claim is true. Otherwise, $g(u_v) = \|u - u_v\| \cdot \sin\psi$, where ψ is the angle between uv and E. So $g(u_v) \leq \|u - u_v\|$. Moreover, as $\psi \geq \phi^u_{uv}$, we have $g(u_v) \geq \|u - u_v\| \cdot \sin\phi^u_{uv}$. This proves our claim. Let $R = \text{radius}(S_{u_v})/g(u_v)$. Our claim implies that

$$R \in [\sin(\mu\theta^u_{uv}), \frac{\sin(\mu\theta^u_{uv})}{\sin\phi^u_{uv}}] \subset [\mu \sin\theta^u_{uv}, \frac{\mu\theta^u_{uv}}{\sin\phi^u_{uv}}].$$

Clearly, $\sin\theta^u_{uv} = \min\{\sin(\pi/3), \sin\phi^u_{uv}\} \geq \min\{\sqrt{3}/2, \sin\phi\} = c_2$. So $R \geq c_2\mu$. If $\phi^u_{uv} \leq \pi/3$, then $\mu\theta^u_{uv}/\sin\phi^u_{uv} = \mu\phi^u_{uv}/\sin\phi^u_{uv}$. Since $\phi^u_{uv}/\sin\phi^u_{uv}$ is increasing within $(0, \pi/2)$ and $\phi^u_{uv} \leq \pi/3$ by assumption, $\phi^u_{uv}/\sin\phi^u_{uv} \leq \pi/(3\sin(\pi/3)) = 2\pi/(3\sqrt{3}) = c_1$. If $\phi^u_{uv} > \pi/3$, then $\theta^u_{uv}/\sin\phi^u_{uv} < \pi/(3\sin(\pi/3)) = 2\pi/(3\sqrt{3}) = c_1$. So $R \leq c_1\mu$. $\square$

After constructing S_{u_v} and S_{v_u}, we call the following algorithm *Split*(u_v, v_u) which returns a sequence of protecting spheres that cover $u_v v_u$. We call two protecting spheres *adjacent* if their centers are neighbors on some edge of $\mathcal{P}$.

Algorithm *Split*(x, y)
Input: The segment xy and protecting spheres S_x and S_y.
Output: A sequence of protecting spheres, including S_x and S_y, that cover xy. Every protecting sphere has positive radius. Any two adjacent protecting spheres are orthogonal.

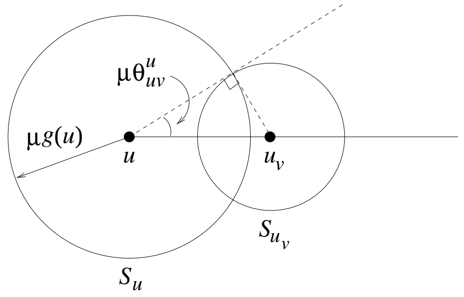

Figure 2: The construction of S_{u_v}.

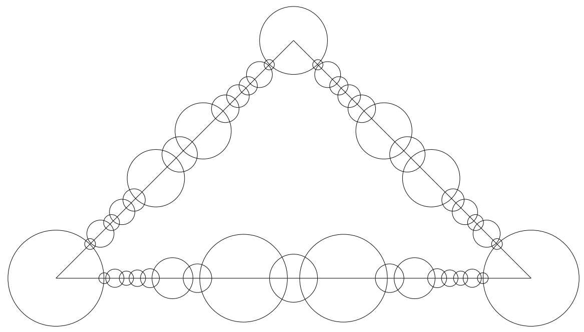

Figure 3: $\mu = 1/7$ and the base angle is $\pi/4$.

1. Compute the point z on xy using the relation

$$\|x - z\| = \frac{\|x - y\|^2 + \mathrm{radius}(S_x)^2 - \mathrm{radius}(S_y)^2}{2 \cdot \|x - y\|}$$

2. Set $Z = \sqrt{\|x - z\|^2 - \mathrm{radius}(S_x)^2}$
3. **if** $Z > 3\mu \cdot g(z)$
4. **then** create a protecting sphere S_z with center z and radius $\mu \cdot g(z)$
5. $Split(x, z)$
6. $Split(z, y)$
7. **else** create a protecting sphere S_z with center z and radius Z

Note that the sphere with center z and radius Z computed in lines 1 and 2 is orthogonal to both S_x and S_y. Figure 3 shows the protecting spheres created for the sides of an isosceles triangle. We will prove that the recursive procedure terminates. We will also bound the radii of the protecting spheres produced. We need the following technical lemma.

LEMMA 3.2 *Let $k = 1.099$. Given a sphere S, let $\overline{S}$ denote the sphere with the same center as S and radius $k \cdot \mathrm{radius}(S)$. Whenever $Split(x, y)$ is called, $\overline{S_x} \cap \overline{S_y} = \emptyset$. Moreover, if $Split(x, y)$ inserts a sphere S_z in line 4, then $\overline{S_x} \cap \overline{S_z} = \overline{S_y} \cap \overline{S_z} = \emptyset$.*

Proof. Let uv be an edge of $\mathcal{P}$. We first show that $\overline{S_{u_v}} \cap \overline{S_{v_u}} = \emptyset$. Since $\mu \le 1/7$, $\theta^u_{uv} \le \pi/3$, and $g(u) \le \|u - v\|$, we have

$$\|u - u_v\| = \mu \sec(\mu\theta^u_{uv}) \cdot g(u) \le \frac{1}{7} \cdot \sec(\pi/21) \cdot \|u - v\| < 0.2 \cdot \|u - v\|.$$

It follows that

$$\mathrm{radius}(S_{u_v}) = \|u - u_v\| \cdot \sin(\mu\theta^u_{uv}) < 0.2 \cdot \|u - v\| \cdot \sin(\pi/21) < 0.05 \cdot \|u - v\|.$$

Thus $\|u - u_v\| + \mathrm{radius}(S_{u_v}) < \|u - v\|/2$ which means that $\overline{S_{u_v}}$ does not reach the midpoint of uv. The same holds for $\overline{S_{v_u}}$. So $\overline{S_{u_v}} \cap \overline{S_{v_u}} = \emptyset$.

Consider the creation of a protecting sphere S_z in line 4 of *Split*(x, y), assuming that $\overline{S_x} \cap \overline{S_y} = \emptyset$. Since z lies outside S_x and line 3 of *Split* is satisfied, we have

$$\|x - z\| > Z > 3\mu \cdot g(z). \tag{1}$$

Assume to the contrary that $\overline{S_x}$ intersects $\overline{S_z}$. Then $\mu \cdot g(z) \geq \|x - z\|/k - \mathrm{radius}(S_x)$. Substituting this into (1), we get $\|x - z\| > (3/k) \cdot \|x - z\| - 3 \cdot \mathrm{radius}(S_x)$, so

$$\|x - z\| < \frac{3k}{3 - k} \cdot \mathrm{radius}(S_x). \tag{2}$$

Let E be an element of $\mathcal{P}$ such that z does not lie on E and E touches the ball centered at z with radius $g(z)$. Let d be the distance between x and E. Starting with the triangle inequality, we get

$$\begin{aligned} d &\leq \|x - z\| + g(z) \\ &\overset{(1)}{<} \frac{1 + 3\mu}{3\mu} \cdot \|x - z\| \\ &\overset{(2)}{<} \frac{k(1 + 3\mu)}{\mu(3 - k)} \cdot \mathrm{radius}(S_x). \end{aligned}$$

Observe that $\mathrm{radius}(S_x) \leq c_1\mu \cdot g(x)$: if $x = u_v$, then $\mathrm{radius}(S_x) \leq c_1\mu \cdot g(x)$ by Lemma 3.1, otherwise line 4 of *Split* enforces that $\mathrm{radius}(S_x) = \mu \cdot g(x)$. This implies that $d < (c_1k(1 + 3\mu)/(3 - k)) \cdot g(x)$. By our choices of k, c_1 and μ, one can verify that $c_1k(1 + 3\mu)/(3 - k) < 1$ and so $d < g(x)$. However, since $x, z \in \mathrm{int}(uv)$ and z does not lie on E, x does not lie on E too. But this implies that $d \geq g(x)$, a contradiction. So $\overline{S_x} \cap \overline{S_z} = \emptyset$. Similarly, $\overline{S_y} \cap \overline{S_z} = \emptyset$. Then the lemma follows by an inductive argument. ◻

We are ready to analyze the recursive procedure *Split*.

LEMMA 3.3 *There exist a constant $c_3 < c_2$ such that for each edge uv of $\mathcal{P}$, the following hold.*

(i) Split(u_v, v_u) terminates and returns a sequence $\mathcal{S}$ of protecting spheres covering $u_v v_u$. Any two adjacent protecting spheres in $\mathcal{S}$ are orthogonal.

(ii) For any $S_z \in \mathcal{S} - \{S_{u_v}, S_{v_u}\}$, the ratio $\frac{\mathrm{radius}(S_z)}{g(z)}$ lies in $[c_3\mu, 3\mu]$.

Proof. If *Split*(u_v, v_u) does not terminate, Lemma 3.2 implies that infinitely many non-intersecting protecting spheres are created in line 4 of *Split*. Each such sphere S_z has radius $\mu \cdot g(z)$. This is impossible as there is a constant $\epsilon > 0$ such that $g(z) \geq \epsilon$ for any point $z \in u_v v_u$. Lines 1, 2 and 7 of *Split* guarantee that any two adjacent protecting spheres created are orthogonal and hence overlapping. Thus, the spheres in $\mathcal{S}$ cover $u_v v_u$. This proves (i).

Take a sphere $S_z \in \mathcal{S} - \{S_{u_v}, S_{v_u}\}$. By lines 3, 4 and 7, $\mathrm{radius}(S_z)/g(z) \leq 3\mu$. Next, we lower bound $\mathrm{radius}(S_z)/g(z)$. If S_z was created in line 4, then $\mathrm{radius}(S_z) = \mu \cdot g(z)$, otherwise $\mathrm{radius}(S_z) = Z$. So it suffices to prove that $Z \geq c_3\mu \cdot g(z)$ when S_z was created in line 7. Lemma 3.2 implies that z is at distance at least $(k - 1) \cdot \mathrm{radius}(S_x)$ from S_x or at least $(k - 1) \cdot \mathrm{radius}(S_y)$ from S_y, say the former is true. Since S_x intersects S_z,

$$Z \geq (k - 1) \cdot \mathrm{radius}(S_x). \tag{3}$$

It follows that $\|x - z\| \leq Z + \text{radius}(S_x) \leq kZ/(k-1)$. Using this and Lemma 2.1, we get

$$g(z) \leq g(x) + \|x - z\| \leq g(x) + kZ/(k-1). \tag{4}$$

Observe that $\text{radius}(S_x) \geq c_2\mu \cdot g(x)$: if $x = u_v$, then $\text{radius}(S_x) \geq c_2\mu \cdot g(x)$ by Lemma 3.1, otherwise $\text{radius}(S_x) = \mu \cdot g(x)$ (note that $c_2 < 1$). Substituting this into (3) yields $Z \geq c_2\mu(k-1) \cdot g(x)$. Substituting this into (4) yields $g(z) \leq Z(1 + c_2\mu k)/(c_2\mu(k-1))$ or equivalently, $Z \geq (c_2\mu(k-1)/(1+c_2\mu k)) \cdot g(z)$. So we can prove (ii) by setting $c_3 = c_2(k-1)/(1+c_2 k)$. ◻

3.2 The new complex $\mathcal{Q}$

Buffer zone. Given a set $\mathcal{S}$ of spheres, we use $\text{Bd}(\bigcup_{S\in\mathcal{S}} S)$ to denote the outer boundary of $\bigcup_{S\in\mathcal{S}} S$. Let $\mathcal{B} = \text{Bd}(\bigcup S_x)$, where S_x runs over all protecting spheres created. The space inside $\mathcal{B}$ is the buffer zone. For each edge uv of $\mathcal{P}$, let $\mathcal{S}_{uv}$ be the sequence of protecting spheres whose centers lie on uv (including u and v). $\mathcal{B} \cap \bigcup_{S_x \in \mathcal{S}_{uv}} S_x$ consists of a sequence of *rings* delimited by two spheres with holes. This decomposition is obtained by cutting $\mathcal{B} \cap \bigcup_{S_x \in \mathcal{S}_{uv}} S_x$ with the bisector planes of adjacent protecting spheres. The two delimiting spheres with holes are $\mathcal{B} \cap S_u$ and $\mathcal{B} \cap S_v$. For each $S_z \in \mathcal{S}_{uv} - \{S_u, S_v\}$, S_z contributes exactly one ring to $\mathcal{B} \cap S_z$. By Lemma 3.3(ii), $\text{radius}(S_x) \leq 3\mu \cdot g(x) < g(x)/2$. If E is a vertex, edge, or facet of $\mathcal{P}$ disjoint from x, the distance between x and E is at least $g(x)$. This implies that $\mathcal{B}$ encloses the edges of $\mathcal{P}$ without causing any unwanted self-intersection or intersection with $\mathcal{P}$.

Structure of $\mathcal{Q}$. We merge $\mathcal{B}$ with $\mathcal{P}$ to produce a new complex $\mathcal{Q}$. The merging of $\mathcal{B}$ and $\mathcal{P}$ means that we will mesh the union of the input domain and the buffer zone. Since we will guarantee that the output mesh conforms to $\mathcal{P}$, it is easy to remove the tetrahedra that lie outside the input domain afterwards.

$\mathcal{B}$ splits each facet of $\mathcal{P}$ into two smaller facets, one inside $\mathcal{B}$ and one outside $\mathcal{B}$. These facets are the *flat facets* of $\mathcal{Q}$. For each edge uv of $\mathcal{P}$, each ring $\mathcal{B} \cap S_x$, where $x \in \text{int}(uv)$, is divided by the facets of $\mathcal{P}$ incident to uv into curved quadrilateral patches. For each non-isolated vertex v of $\mathcal{P}$, $\mathcal{B} \cap S_v$ is divided by the facets of $\mathcal{P}$ incident to v into spherical patches. All the above curved patches are the *curved facets* of $\mathcal{Q}$. The centers of protecting spheres split the edges of $\mathcal{P}$ into the *linear edges* of $\mathcal{Q}$. The circular arcs on the boundaries of curved and flat facets are the *curved edges* of $\mathcal{Q}$. The vertices of $\mathcal{Q}$ consists of the endpoints of the linear and curved edges, as well as the isolated vertices of $\mathcal{P}$. Two elements of $\mathcal{Q}$ are *adjacent* if their intersection is non-empty. Two elements of $\mathcal{Q}$ are *incident* if one is the boundary element of the other. Figures 4 and 5 show some examples.

We elaborate on the topology of curved facets and the geometry of their boundaries. If an edge uv of $\mathcal{P}$ has zero or one incident facet, any curved edge between two adjacent protecting spheres centered on uv is a full circle: if uv has no incident facet, the curved edge has no endpoint; otherwise, it has exactly one endpoint. Let E be a curved facet on $\mathcal{B} \cap S_x$ for some protecting sphere S_x. E may not be a topological disk. For example, if x is the common endpoint of three edges of $\mathcal{P}$ with no incident facet, then $\mathcal{B} \cap S_x$ is a single curved facet with three disjoint boundary curved edges (which are full circles). The above example also shows that ∂E may consist of more than one connected component. Let ∂E_i be a connected component of ∂E. ∂E_i may not be a simple closed curve, see Figure 5 for an example. Any curved edge in ∂E_i has one of the following two types:

Type 1: the curved edge lies at the intersection between S_x and a facet of $\mathcal{P}$.

Type 2: the curved edge lies at the intersection between S_x and an adjacent protecting sphere.

∂E_i may be a single type 2 edge which must then be a full circle. If ∂E_i contains a type 1 curved edge, this edge has two distinct endpoints that are incident on two distinct type 2 curved edges in ∂E_i. In the case where E is a topological disk, ∂E is a simple cycle and the type 1 and type 2 curved edges alternate in ∂E. See Figure 4 for an illustration.

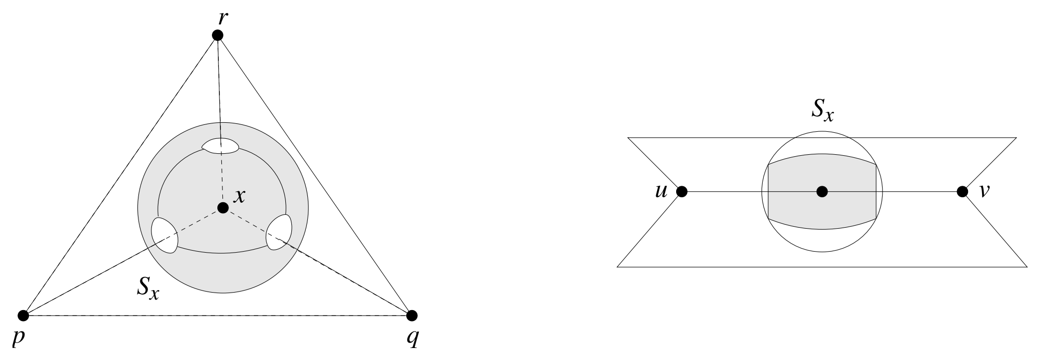

Figure 4: In the left figure, $\mathcal{P}$ consists of the boundary triangles of the tetrahedron $pqrx$. The three facets incident to x divide $\mathcal{B} \cap S_x$ into two curved facets and each is a topological disk. In the right figure, the two facets incident to the edge uv divide the ring $\mathcal{B} \cap S_x$ into two quadrilateral curved facets and one is shown shaded. In both figures, the type 1 and type 2 curved edges alternate in the boundary of the curved facets.

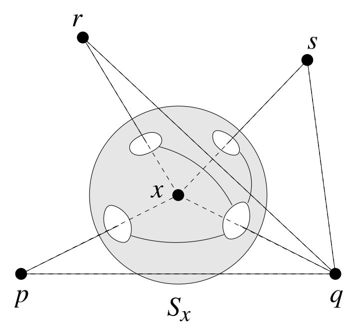

Figure 5: $\mathcal{P}$ consists of three triangles incident to the edge qx. $\mathcal{B} \cap S_x$ is a single curved facet. The hole on $\mathcal{B} \cap S_x$ around the edge qx contains three curved edges.

How many (type 1) curved edges can a facet F of $\mathcal{P}$, where $x \in \partial F$, contribute to ∂E? If $\mathcal{B} \cap S_x$ is a ring, the answer is clearly one. Suppose that x is a vertex of $\mathcal{P}$. Observe that x lies on exactly one simple cycle in ∂F. Moreover, S_x is too small to intersect more than one cycle in ∂F or intersect the same cycle more than twice. Thus, $S_x \cap F$ is connected. It follows that F contributes exactly one edge to ∂E. How many (type 2) curved edges in ∂E may lie on the same hole on $\mathcal{B} \cap S_x$? If $\mathcal{B} \cap S_x$ is a ring, the answer is clearly one. Otherwise, there may be more than one, see Figure 5.

Angles. By design, all angles in the space outside $\mathcal{B}$ are equal to $\pi/2$. The next lemma gives a precise statement.

LEMMA 3.4

(i) Let F be a curved facet. Let F' be a curved/flat facet adjacent to F. If F and F' do not lie on the same sphere, the normal to F' at any point in $F \cap F'$ is tangent to F.

(ii) Let e and e' be two adjacent curved edges that do not lie on the same circle. Let ℓ (resp. ℓ') be the line through $e \cap e'$ that is tangent to and coplanar with e (resp. e'). Then ℓ is perpendicular to ℓ'.

(iii) Let F be a curved/flat facet. Let e be a curved edge that is adjacent to F but not incident on F. If e and F do not lie on the same plane or sphere, the normal to F at $e \cap F$ is tangent to and coplanar with e.

4 Algorithm MESH

Since all angles outside $\mathcal{B}$ are equal to $\pi/2$ by Lemma 3.4, Delaunay refinement can be applied in the space outside $\mathcal{B}$. Of course, it has to be enhanced in order to deal with the curved elements of $\mathcal{Q}$. In essence, we compute a mesh that approximates $\mathcal{Q}$ and conforms to the elements of $\mathcal{P}$. Our algorithm inserts points incrementally and maintains a set $\mathcal{V}$ of vertices. $\mathcal{V}$ is initialized to be the set of vertices of $\mathcal{Q}$. The points to be inserted are related to three types of geometric objects: helper arcs, helper triangles, and subfacets. In the following, we first provide their definitions and then describe our algorithm.

Notations. Given a circle C on a sphere S, the *orthogonal sphere of S at C* is the sphere orthogonal to S that passes through C. We use $\widehat{pq}$ to denote a circular arc α with endpoints p and q. Note that if $p = q$, α is a full circle. Given a cap K on a sphere S, if the angular diameter of K is less than π, we use $K^\perp$ to denote the orthogonal sphere of S at ∂K. In this case, K lies inside $K^\perp$. If S is a plane (infinite sphere), then $K^\perp$ is the equatorial sphere of K.

Helper arcs. Each curved edge e of $\mathcal{Q}$ is split by the vertices in $\mathcal{V}$ into *helper arcs*. Let S be the equatorial sphere of e, i.e., e lies on an equator of S. Note that S is a protecting sphere iff e is a type 1 edge. Let α be a helper arc on e. The *circumcap* of α is the smallest cap on S that contains α. It is denoted by K_α. If the angular width of α is less than π, the *normal sphere* of α is $K_\alpha^\perp$. The helper arc α is *encroached* by a point v if v lies on or inside $K_\alpha^\perp$. (This is stronger than disallowing v from lying inside $K_\alpha^\perp$. The stronger definition makes it easier to achieve conformity.) If the angular width of α is larger than $\pi/3$, α is *wide*. In the special case where the curved edge e contains zero or one vertex in $\mathcal{V}$, e is one helper arc and e is wide. In these cases, if e contains only one vertex in $\mathcal{V}$, we define its midpoint to be the point diametrically opposite to this vertex; if e contains no vertex in $\mathcal{V}$, we fix an arbitrary point on e to be its midpoint. The notions of circumcap, normal sphere, and wideness as well as their notations can be generalized to any arc on a curved edge.

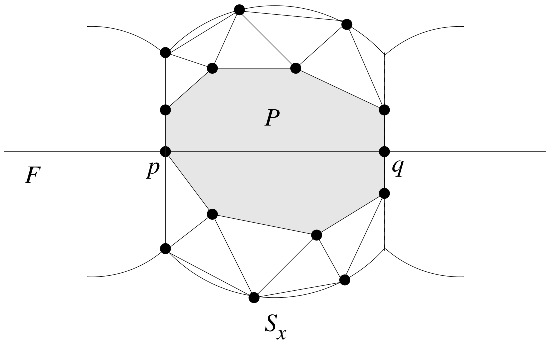

Figure 6: The horizontal line in the middle is the side view of a facet F of $\mathcal{P}$ cutting the three protecting spheres shown. The shaded polygon P is a facet of *CH*$_x$. F cuts S_x and generates a helper arc $\widehat{pq}$.

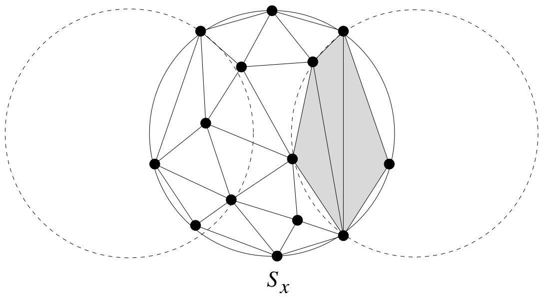

Figure 7: The figure shows S_x and two protecting spheres adjacent to S_x. Some boundary triangles of *CH*$_x$ are shown. The non-shaded triangles are helper triangles. The shaded ones are not as the vertices of each shaded triangle lie on the boundary of the same hole on $\mathcal{B} \cap S_x$.

Helper triangles. Helper triangles are defined when no helper arc is wide or encroached by a vertex in $\mathcal{V}$. Let *CH*$_x$ denote the convex hull of $\mathcal{V} \cap \mathcal{B} \cap S_x$ for a protecting sphere S_x. Note that *CH*$_x$ has at least one vertex: since x is not an isolated vertex of $\mathcal{P}$, $\mathcal{B} \cap S_x$ contains some helper arc(s) and they are not wide. We first deal with the general case where *CH*$_x$ is three-dimensional.

If a convex polygon P with more than three vertices appears as a boundary facet of *CH*$_x$, we arbitrarily triangulate P. Note that P cannot stab an input facet as in Figure 6, otherwise since the vertices of P are cocircular, the helper arc $\widehat{pq}$ would be encroached by some vertex of P, a contradiction. (It may happen that the vertices of P lie on the boundary of the circumcap of $\widehat{pq}$. This is where we need the stronger definition of encroachment for helper arcs.) Therefore, the arbitrary triangulation of P does not cause any concern for conformity. A boundary triangle t of *CH*$_x$ is a *helper triangle* if no hole on $\mathcal{B} \cap S_x$ contains all vertices of t on its boundary. See Figure 7. Let H be the plane containing a helper triangle t. The *circumcap* of t is the cap on S_x that is bounded by $H \cap S_x$ and protrudes in the outward normal direction of t. The circumcap of t is denoted by K_t. If the angular diameter of K_t is less than π, the *normal sphere* of t is $K_t^{\perp}$. The helper triangle t is *encroached* by a point v if v lies inside $K_t^{\perp}$. If the angular diameter of K_t is larger than $\pi/3$, t is *wide*.

It may happen that *CH*$_x$ has dimension less than three. As x is not an isolated vertex of $\mathcal{P}$, $\mathcal{B} \cap S_x$ has at least one hole. Since no helper arc is wide by assumption, there are at least six vertices on each hole boundary. So the dimension of *CH*$_x$ is at least two. It is exactly two when $\mathcal{B} \cap S_x$ has only one hole and all vertices in $\mathcal{V} \cap \mathcal{B} \cap S_x$ lie on the boundary of this hole. In this case, we arbitrarily triangulate *CH*$_x$. We duplicate each resulting triangle so as to treat *CH*$_x$ as a three-dimensional body with zero volume. We assign opposite outward normals to the duplicates of each triangle. We only take the copies of triangles facing x as helper triangles. Their circumcaps are defined as in the three-dimensional case. All helper triangles are wide in this case, and their normal spheres are undefined.

Subfacets. Subfacets are defined when no helper arc is wide or encroached by a vertex in $\mathcal{V}$. For every facet F of $\mathcal{P}$, a *subfacet* is a triangle on F in the two-dimensional Delaunay triangulation of $\mathcal{V} \cap F$, that does not lie completely inside $\mathcal{B}$. We define subfacet using facets of $\mathcal{P}$ instead of flat facets of $\mathcal{Q}$ because MESH only approximates $\mathcal{Q}$ and it does not respect the boundary curved edges of flat facets.

Recall that F is divided into two flat facets by $\mathcal{B}$. Let ξ denote the curved boundary between these two flat facets. For each helper arc $\widehat{pq}$ on ξ, since $\widehat{pq}$ is not encroached, the edge pq appears in the two-dimensional Delaunay triangulation. Thus, ξ is approximated by a polygonal closed curve with vertices on ξ. This implies that the vertices of any subfacet on F must lie on the flat facet that lies outside $\mathcal{B}$.

The *circumcap* of a subfacet σ is the disk bounded by the circumcircle of σ. It is denoted by K_σ. The *normal sphere* of σ, denoted by $K_\sigma^\perp$, is the equatorial sphere of σ. If a point v lies inside $K_\sigma^\perp$, σ is *encroached* by v.

Algorithm. We are ready to describe MESH. Starting with $\mathcal{V}$ as the set of vertices of $\mathcal{Q}$, MESH repeatedly invoke the applicable rule of the least index in the following list until no rule is applicable. Recall that $\rho_0 > 16$ is an a priori chosen constant.

Rule 1: Pick a helper arc α that is wide or encroached by a vertex in $\mathcal{V}$. Preference is given to wide helper arcs. Insert the midpoint of α.

Rule 2: Pick a helper triangle t that is wide or encroached by a vertex in $\mathcal{V}$. Preference is given to wide helper triangles, and if there are more than one wide helper triangle, preference is given to those with angular diameter π or more. Let v be the center of K_t. If v does not encroach upon any helper arc, insert v. Otherwise, reject v and apply rule 1 to split one helper arc encroached by v.

Rule 3: Pick a subfacet σ that is encroached by a vertex in $\mathcal{V}$. Let v be the center of K_σ. If v does not encroach upon any helper arc, insert v. Otherwise, reject v and apply rule 1 to split one helper arc encroached by v.

Rule 4: Let τ be a tetrahedron such that $\rho(\tau) > \rho_0$, τ lies inside the domain, and τ does not lie inside $\mathcal{B}$. Let v be the circumcenter of τ. If v does not encroach upon any helper arc, helper triangle, or subfacet, then insert v. Otherwise, reject v and apply one of the following:

- If v encroaches some helper arc(s), use rule 1 to split one.
- Otherwise, v encroaches upon some helper triangle(s) or subfacet(s). Use rule 2 or 3 correspondingly to split one.

When the above loop terminates, we triangulate the inside of $\mathcal{B}$ using the following two types of tetrahedra. First, for each protecting sphere S_x, we construct the convex hull of x and each helper triangle on CH_x. Second, for each linear edge xy and each helper arc $\widehat{pq} \subseteq S_x \cap S_y$, we construct the tetrahedron $pqxy$. Finally, the subcomplex covering the input domain is the final mesh.

Running time. Assuming that MESH terminates, we derive the running time of the algorithm in terms of N, the number of output vertices. Unfortunately, the running time cannot be polynomial in the input size. This is impossible even in two dimensions if bounded radius-edge ratio is achieved [16]. Note that N is also an upper bound on the number of input vertices. MESH has to construct the protecting spheres, maintain CH_x for each S_x, the two-dimensional Delaunay triangulation for each facet of $\mathcal{P}$, and the three-dimensional Delaunay triangulation $\mathrm{Del}\,\mathcal{V}$.

In constructing each protecting sphere S_x, we need to compute $g(x)$. This can be done by checking every vertex, edge, and facet of $\mathcal{P}$. Observe that the edges and facets of $\mathcal{P}$ contain some vertices of $\mathcal{Q}$ in their interior. As such interior vertices of $\mathcal{Q}$ cannot be shared, there are $O(N)$ edges and facets in $\mathcal{P}$. In all, constructing all protecting spheres takes $O(N^2)$ time.

When a vertex p is inserted on $\mathcal{B} \cap S_x$, the existing triangles on CH_x that are visible from p are deleted and they can be identified by a linear-time search. Then the resulting polygonal hole is connected to p to form the new triangles. Thus the number of new triangles created is proportional to the number of triangles deleted. It follows that the total time for maintaining the convex hulls for all protecting spheres is $O(N^2)$.

When a vertex p is inserted on a facet F of $\mathcal{P}$, we invoke a linear-time search to find the triangles in the two-dimensional Delaunay triangulation whose circumcircles contain p. We delete these triangles and connect the resulting polygonal hole to p to form the new two-dimensional Delaunay triangulation. Thus the number of new triangles created is proportional to the number of triangles deleted. It follows that the total time for maintaining the two-dimensional Delaunay triangulations for all facets of $\mathcal{P}$ is $O(N^2)$.

Consider the maintenance of $\mathrm{Del}\,\mathcal{V}$. Let $\mu : \mathbb{R}^3 \rightarrow \mathbb{R}^4$ be the map that sends a point $x = (x_1, x_2, x_3) \in \mathbb{R}^3$ to a point $\mu(x) = (x_1, x_2, x_3, x_1^2 + x_2^2 + x_3^2) \in \mathbb{R}^4$. For a point set $\mathcal{V}$ in three dimensions, let $\mu(\mathcal{V}) = \{\mu(v), v \in \mathcal{V}\}$. The Delaunay triangulation of $\mathcal{V}$ is the projection of the convex hull of $\mu(\mathcal{V})$ [9]. So the first $\mathrm{Del}\,\mathcal{V}$ can be done in $O(N^2)$ time using Chazelle's convex hull algorithm [2]. MESH iterates less than N times and a vertex is inserted in each iteration. After the insertion of a vertex p, we can no longer afford a linear-time search to find the tetrahedra whose circumspheres contain p. If we do so, we may need to check more than $O(N)$ tetrahedra. Instead, we first locate a Delaunay tetrahedron that is destroyed by the insertion of p. This can easily be done in $O(1)$ time. If p is the circumcenter of a skinny tetrahedron, we take this tetrahedron. If p splits a helper triangle or subfacet, we take any tetrahedron incident to the helper triangle or subfacet. If p splits a helper arc $\widehat{rs}$, we take any any tetrahedron incident to the edge rs. Afterwards, starting from this tetrahedron, we explore the Delaunay triangulation in a depth first manner to collect all tetrahedra that are destroyed with the insertion of p. Once these tetrahedra are identified, p is connected to the boundary of the union of them to update the Delaunay triangulation. If D_p is the number of deleted tetrahedra, the complexity of this update is $O(D_p)$. We argue that the total number of deleted tetrahedra in the entire algorithm is $O(N^2)$.

In the lifted diagram in four dimensions, the insertion of p can be viewed as follows. The point $\mu(p)$ is below the convex hull of $\mu(\mathcal{V})$ and let T be the set of tetrahedra on this convex hull visible to $\mu(p)$. Insertion of $\mu(p)$ creates new tetrahedra on the updated convex hull by connecting $\mu(p)$ to the boundary of the union of tetrahedra in T. The space between these new tetrahedra and T can be triangulated by connecting $\mu(p)$ to each tetrahedron in T. Thus, assuming that the convex hull of the initial point set is triangulated, one can maintain a triangulation in the lifted diagram after each insertion, which contains the lifted deleted tetrahedra. Therefore, all tetrahedra deleted by MESH can be mapped to tetrahedra in the triangulation of N points in four dimensions. Since the size of any triangulation of N points in four dimensions is only $O(N^2)$ (Theorem 1.2 [8]), the same bound applies to the number of deleted tetrahedra.

Each insertion is preceded by a search of an encroached helper arc, encroached helper triangle, encroached subfacet, or a skinny tetrahedron. We argue that this search can also be done in $O(N^2)$ total time. We maintain a stack of all skinny tetrahedra, which means that a skinny tetrahedra can be accessed in $O(1)$ time. Next, we need to account for searching for the encroached objects. This encroachment may occur by an inserted or rejected point. Since each rejected point leads to an insertion, the total number of inserted and rejected points is $O(N)$. For each such point we can scan all helper arcs, helper triangles, and subfacets to determine the encroachments. The helper triangles on each curved facet form a planar graph and so do the subfacets on each flat facet. Also, since all angles outside $\mathcal{B}$ are at least $\pi/2$, a helper arc is incident on only a constant number of curved facets and flat facets. It follows that the total number of helper arcs, helper triangles, and subfacets is $O(N)$ at any time of the algorithm. Therefore, counting over all points, all encroachments can be determined in $O(N^2)$ time.

LEMMA 4.1 MESH *runs in $O(N^2)$ time, where N is the number of output vertices.*

5 Geometric lemmas

We prove some geometry properties of orthogonal spheres and planes, which will be needed in Sections 6 and 7. We have defined orthogonality between two spheres before. Given a sphere S_1 and a plane S_2, we say that S_1 and S_2 are orthogonal if the center of S_1 lies on S_2.

LEMMA 5.1 *Let S_1 be a sphere and let S_0 be a sphere or plane. Assume that S_0 and S_1 are orthogonal. Then a sphere S_2 is orthogonal to S_0 if and only if the bisector plane of S_1 and S_2 is orthogonal to S_0.*

Proof. Suppose that S_0 is a sphere. Let π_i be the weighted distance function for S_i for i = 1 and 2. Let x be the center of S_0. Since S_1 and S_0 are orthogonal, $\pi_1(x) = \mathrm{radius}(S_0)^2$. S_0 and S_2 are orthogonal if and only if $\pi_2(x) = \mathrm{radius}(S_0)^2$, which is equivalent to $\pi_1(x) = \pi_2(x)$.

Suppose that S_0 is a plane. Let ℓ be the line through the centers of S_1 and S_2, which is orthogonal to their bisector plane. Since S_1 and S_0 are orthogonal, the center of S_1 lies on S_0. So S_0 and S_2 are orthogonal if and only if ℓ lies on S_0, which is equivalent to the bisector plane of S_1 and S_2 being orthogonal to S_0. $\square$

LEMMA 5.2 *Let S_1 be a sphere and let S_2 be a sphere or plane. Assume that S_1 and S_2 are orthogonal. Let a and b be two points on $S_1 \cap S_2$ that are not diametrally opposite. There is a unique sphere S_3 that passes through a and b, and is orthogonal to S_1 and S_2.*

Proof. Suppose that S_2 is a sphere. Orient space so that ab and the line through the centers of S_1 and S_2 are horizontal. For i = 1 or 2, let H_a^i and H_b^i be the tangent planes of S_i at a and b, respectively. Observe that the line $H_a^i \cap H_b^i$ is the locus of the centers of spheres that pass through a and b and are orthogonal to S_i. Note that the orientation of $H_a^i \cap H_b^i$ is perpendicular to the orientation of ab. There is another way to define the locus of such sphere centers. Let H_{ab} be the vertical plane that passes through the midpoint of ab. Then we have $H_a^i \cap H_b^i = H_a^i \cap H_{ab}$. Since ab are not diametrally opposite, the lines $H_a^1 \cap H_{ab}$ and $H_a^2 \cap H_{ab}$ are not parallel. So $H_a^1 \cap H_{ab}$ and $H_a^2 \cap H_{ab}$ intersect at exactly one point, which is the center of S_3.

Suppose that S_2 is a plane. Since a and b lie on $S_1 \cap S_2$, H_a^1 and H_b^1 are perpendicular to S_2. Since a and b are not diametrally opposite, H_a^1 and H_b^1 are not parallel and they intersect. So the line $H_a^1 \cap H_b^1$ is perpendicular to S_2 and intersects S_2 at one point. This point is the center of S_3. $\square$

LEMMA 5.3 *Let S be a sphere and let S' be a sphere or plane. Assume that S and S' are orthogonal. If α is a circular arc on $S \cap S'$ with angular width less than π, then $K_\alpha^\perp$ is orthogonal to both S and S'.*

Proof. Let S'' be the equatorial sphere of α. By construction, the centers of S'' and $K_\alpha^\perp$ lie on the bisector plane of S and S'. So the bisector plane of S and S' is orthogonal to both S'' and $K_\alpha^\perp$. Note that the bisector plane of S and S' is also the bisector plane of S and S''. So we can invoke Lemma 5.1 with $K_\alpha^\perp$, S'', and S taking the roles of S_0, S_1, and S_2, respectively. It follows that S is orthogonal to $K_\alpha^\perp$. Similarly, we can show that S' is orthogonal to $K_\alpha^\perp$. $\square$

Let α be a non-wide helper arc α at the intersection of two protecting spheres S_x and S_y. Lemma 5.3 implies that $K_\alpha^\perp$ is orthogonal to S_x and S_y. Moreover, Lemmas 5.2 and 5.3 imply that $K_\alpha^\perp$ is the only sphere that passes through the endpoints of α, and is orthogonal to S_x and S_y.

LEMMA 5.4 *Let S be a sphere or plane. Let K_1, K_2, and K_3 be three caps on S. Assume that the angular widths of K_1, K_2, and K_3 are less than π. If $K_3 \subseteq K_1 \cup K_2$, then $\mathrm{Bd}(K_1^\perp \cup K_2^\perp)$ encloses $K_3^\perp$.*

Proof. Let H_{ij} be the bisector plane of $K_i^\perp$ and $K_j^\perp$. By Lemma 5.1, H_{ij} is orthogonal to S. We claim that if H is a plane orthogonal to S, then $H \cap S$ and $H \cap K_i^\perp$ intersect at right angle. If S is a plane, the center of the circle $H \cap K_i^\perp$ clearly lies on the line $H \cap S$. If S is a sphere, the claim has been proved in Claim 1 in [4].

Consider the case in which K_3 lies within K_1 or K_2, say K_1. We show that $K_1^\perp \cap K_3^\perp$ do not cross. Suppose not. Then H_{13} passes through the circle $K_1^\perp \cap K_3^\perp$. By our previous claim, $H_{13} \cap S$ and $K_1^\perp \cap K_3^\perp$ intersect at right angle. So $K_1^\perp \cap K_3^\perp$ intersects S at two distinct points. But then ∂K_1 and ∂K_3 cross at these two points, a contradiction. So $K_1^\perp$ encloses $K_3^\perp$.

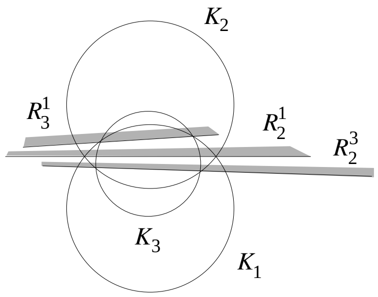

Figure 8:

Consider the case in which $K_3 \not\subseteq K_1$ and $K_3 \not\subseteq K_2$. Let B_i denote the ball bound by $K_i^\perp$. Let R_j^i be the halfspace bound by H_{ij} such that $B_i \cap R_j^i \subseteq B_j$. Refer to Figure 8. Some R_j^i's are shown as shaded sides of the bisector planes. The line $H_{12} \cap H_{13}$ does not stab $K_1 \cup K_2$ because the planes H_{12} and H_{13} intersect the boundary of K_1 only twice. Similarly, the line $H_{12} \cap H_{23}$ does not stab $K_1 \cup K_2$. Therefore, $B_3 \cap R_3^1 \subseteq B_3 \cap R_2^1 \subseteq B_3 \cap R_2^3$. Since $B_3 \cap R_2^3 \subseteq B_2$ by the definition of R_2^3, we obtain $B_3 \cap R_3^1 \subseteq B_2$. Also, $B_3 \cap R_1^3 \subseteq B_1$ by the definition of R_1^3. So $B_3 \subseteq B_1 \cup B_2$. $\square$

6 Buffer zone clearance, conformity, and Delaunayhood

In this section, we prove that MESH never inserts a vertex inside $\mathcal{B}$, and $\mathrm{Del}\,\mathcal{V}$ conforms to the elements of $\mathcal{P}$ whenever no helper arc is wide or encroached and no subfacet is encroached. Thus, assuming termination, MESH indeed returns a conforming Delaunay mesh. Termination will be proved in Section 7.

6.1 Attachment and locations of centers

We first study the location of the circumcap centers of helper triangles and subfacets because they are some of the vertices inserted by MESH. Let σ be a helper triangle or subfacet. We say that σ *attaches* to an element E of $\mathcal{Q}$ if the center of K_σ lies in $\mathrm{int}(E)$. We show that attachment is well-defined under the right conditions.

LEMMA 6.1 *Let σ be a helper triangle (resp. subfacet). Assume that there is no wide or encroached helper arc. Then σ attaches to exactly one curved facet of $\mathcal{Q}$ (resp. flat facet of $\mathcal{Q}$ outside $\mathcal{B}$).*

Proof. If σ is a helper triangle, let S denote the protecting sphere that contains the vertices of σ, and let $R = S \cap \mathcal{B}$. If σ is a subfacet, let S denote the support plane of σ, and let R be the flat facet containing the vertices of σ (R lies outside $\mathcal{B}$ by the definition of subfacets). In both cases, the center v of K_σ lies on S and the vertices of σ lie on $R \subseteq S$.

We first show that $v \in \mathrm{int}(R)$. Assume to the contrary that $v \notin \mathrm{int}(R)$. We claim that v lies on or inside some protecting sphere S_x adjacent to R such that K_σ crosses a curved edge in $S_x \cap R$. Suppose that S is a protecting sphere. Then v lies on or inside a hole $S_x \cap S$ on S for some protecting sphere S_x adjacent to S. $S_x \cap S$ consists of curved edges. $S_x \cap S$ cannot lie inside K_σ; otherwise there would be a wide helper arc on $S_x \cap S$ by the emptiness of K_σ. Conversely, K_σ cannot lie inside the hole $S_x \cap S$ because the vertices of σ lie on R. Hence K_σ crosses a curved edge in $S_x \cap S$. Suppose that S is a plane. So some curved edges in ∂R intersect K_σ, and they cross K_σ completely by the emptiness of K_σ. Exactly one such curved edge e separates v from $\mathrm{int}(R)$ within K_σ. We prove the claim by setting S_x to be the protecting sphere that e lies on.

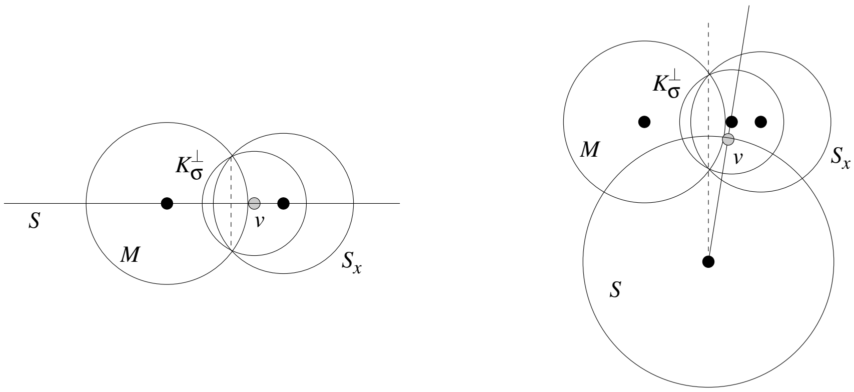

Figure 9: The left and right figures represent the cases in which S is a plane and a sphere, respectively. The shaded dot is v. The black dots are sphere centers. In the left figure, the side view of S is shown. The center of $K_\sigma^\perp$ happens to be v in the left figure.

Let α be the arc $K_\sigma \cap (R \cap S_x)$. By our claim, α lies on a curved edge. Thus α also lies on a helper arc by the emptiness of K_σ. The angular width of α is at most $\pi/3$ as there is no wide helper arc. Let M be the orthogonal sphere of S_x at the circle $K_\sigma^\perp \cap S_x$. Figure 9 shows the situations. The centers of M, $K_\sigma^\perp$, and S_x are collinear. Since the center v of K_σ lies on or inside S_x, $K_\sigma^\perp$ and S_x are closer than orthogonal. Therefore, as the angular width of α is less than π, the center of $K_\sigma^\perp$ lies between the centers of M and S_x. It follows that $\mathrm{Bd}(M \cup S_x)$ encloses $K_\sigma^\perp$. Since S_x is adjacent to R, S_x and S are orthogonal. By Lemma 5.1, the bisector plane of $K_\sigma^\perp$ and S_x is orthogonal to S. Observe that this plane is also the bisector plane of M and S_x. So Lemma 5.1 implies that M is also orthogonal to S. As M passes through the endpoints of α, Lemmas 5.2 and 5.3 imply that $M = K_\alpha^\perp$. We conclude that $\mathrm{Bd}(K_\alpha^\perp \cup S_x)$ encloses $K_\sigma^\perp$. Since some vertex of σ lies outside S_x, it must lie inside $K_\alpha^\perp$ then. By Lemma 5.4, the normal sphere of the helper arc containing α also contains this vertex of σ. But then it is encroached, a contradiction.

We have shown that $v \in \mathrm{int}(R)$. If R is a flat facet, we are done. Otherwise, S is a protecting sphere and the lemma can only be violated when v lies on a curved edge e shared by two curved facets on R. In this case, e is a great circular arc of S and so is $\beta = K_\sigma \cap e$. The vertices of σ lie on the boundary of $K_\sigma = K_\beta$. But then the helper arc containing β is encroached by the vertices of σ, a contradiction. ◻

6.2 Clearance of buffer zone

We show that $\mathcal{B}$ contains no vertex other than the endpoints of linear edges.

LEMMA 6.2 MESH *never inserts any vertex inside $\mathcal{B}$.*

Proof. Assume to the contrary that MESH wants to insert a vertex p inside $\mathcal{B}$ for the first time. MESH is not applying rule 1 since rule 1 never inserts a vertex inside $\mathcal{B}$. It follows that there is no wide or encroached helper arc. By Lemma 6.1, p is not inserted by rules 2 and 3. Thus p is inserted by rule 4 inside some protecting sphere S_x. In this case, there is no wide or encroached helper arc/triangle, and p is the circumcenter of some tetrahedron τ. By rule 4, τ has a vertex outside S_x, so the circumball of τ intersects S_x.

We consider two kinds of caps on S_x. The boundary of CH_x consists of helper triangles and convex polygons that lie between S_x and adjacent protecting spheres. Note that the vertices of such a convex polygon lie on the boundary of a hole on $S_x \cap \mathcal{B}$. The caps of the first kind are the circumcaps of helper triangles. The caps of the second kind are the circumcaps of the convex polygons (i.e., the caps separated from x by the support planes of the polygons). We number the caps K_1, K_2, $\ldots$, in an arbitrary order. Observe that the angular diameters of the above caps are less than π. Also, for each cap K_i of the second kind, $K_i^{\perp}$ is a protecting sphere adjacent to S_x.

The intersection of S_x and the circumball of τ is a cap, which we denote by K. Let K_i be a cap on S_x intersected by K. Note that K does not contain K_i as K is empty. Let S_τ denote the circumsphere of τ.

Case 1: $K \subseteq K_i$. As K_i has angular diameter less than π, so does K. The center of $K^{\perp}$, x, and p are collinear. As p lies inside S_x and x does not lie inside S_τ, p lies between x and the center of $K^{\perp}$. This implies that $\mathrm{Bd}(K^{\perp} \cup S_x)$ encloses S_τ. By Lemma 5.4, $K_i^{\perp}$ encloses $K^{\perp}$. So $\mathrm{Bd}(K_i^{\perp} \cup S_x)$ encloses S_τ. Since τ has a vertex u outside S_x, u lies inside $K_i^{\perp}$. This means that u lies inside some protecting sphere adjacent to S_x or u encroaches some helper triangle, a contradiction.

Case 2: $K \nsubseteq K_i$. Then $K - K_i$ intersects a cap K_j such that the helper triangles/convex polygons corresponding to K_i and K_j share an edge. Since the endpoints of this edge do not lie inside K, the arc $K \cap \partial K_i$ lies within the arc $K_j \cap \partial K_i$. If $K \subseteq K_i \cup K_j$, Lemma 5.4 implies that $\mathrm{Bd}(K_i^{\perp} \cup K_j^{\perp})$ encloses $K^{\perp}$. So $\mathrm{Bd}(K_i^{\perp} \cup K_j^{\perp} \cup S_x)$ encloses S_τ. Since τ has a vertex u outside S_x, we conclude as in case 1 that u lies inside some protecting sphere adjacent to S_x or u encroaches some helper triangle, a contradiction. If $K \nsubseteq K_i \cup K_j$, then $K \cap K_i \subset K \cap K_j$. Thus we can repeat case 1 or case 2 with K_i replaced by K_j, and we will never return to K_i again. Hence we will reach a contradiction eventually.

$\square$

6.3 Conformity and Delaunayhood

We show that whenever rules 1, 2, and 3 do not apply, the current mesh is Delaunay and conforming to $\mathcal{P}$. Also, MESH never inserts any vertex outside the union of the domain and the buffer zone.

LEMMA 6.3 *Assume that there is no wide or encroached helper arc/triangle, and there is no encroached subfacet. Then the current mesh conforms to $\mathcal{P}$ and it is Delaunay.*

Proof. Since no subfacet is encroached, they all appear in the current mesh. So we only need to study what happens inside $\mathcal{B}$. Let S_x and S_y be two adjacent protecting spheres. Let F be a facet of $\mathcal{P}$ intersecting S_x. Note that $x \in \partial F$. By Lemma 6.2, the only vertices inside $\mathcal{B}$ are the endpoints of linear edges of $\mathcal{Q}$.

Let $\alpha = \widehat{pq}$ be a helper arc on $S_x \cap F$. Let S be the equatorial sphere of pqx. Observe that the center of S lies between the centers of S_x and $K_\alpha^{\perp}$ on a straight line. Thus, $\mathrm{Bd}(S_x \cup K_\alpha^{\perp})$ encloses S. Since x is the only vertex inside $\mathrm{Bd}(S_x \cup K_\alpha^{\perp})$, S is empty.

Let p be a helper arc endpoint lying on $S_x \cap S_y$. Let S be the equatorial sphere of pxy. Since S_x and S_y intersect at right angle, $\angle xpy$ in triangle pxy is equal to $\pi/2$. Thus, xy is the diameter of S which implies that $\mathrm{Bd}(S_x \cup S_y)$ encloses S. Since x and y are the only vertices inside $\mathrm{Bd}(S_x \cup S_y)$, S is empty.

The above implies that $\mathcal{P}$ is covered by Delaunay edges and triangles in the current. Hence the current mesh is conforming. MESH guarantees that all tetrahedra, that do not lie completely inside $\mathcal{B}$, are Delaunay. So we only need to study the tetrahedra inside $\mathcal{B}$.

Let $\widehat{pq}$ be a helper arc on $S_x \cap S_y$. The circumsphere of $pqxy$ is the equatorial sphere of pxy which has been proved to be empty in the above. Let pqr be a helper triangle on the convex hull of points on S_x. Let S be the circumsphere of $pqrx$. Observe that the center of S lies between the centers of S_x and $K_{pqr}^{\perp}$ on a straight line. Thus, $\mathrm{Bd}(S_x \cup K_{pqr}^{\perp})$ encloses S. Since x is the only vertex inside $\mathrm{Bd}(S_x \cup K_{pqr}^{\perp})$, S is empty. ◻

LEMMA 6.4 MESH *never inserts any vertex outside the union of the domain and the buffer zone.*

Proof. Assume to the contrary that MESH inserts a vertex p outside the union of the domain and the buffer zone for the first time. Clearly, p cannot be inserted by rule 1. By Lemma 6.1, p cannot be inserted by rules 2 and 3. So p is inserted by rule 4, and p is the circumcenter of a tetrahedron τ inside the domain. In this case, there is no wide or encroached helper arc/triangle and there is no encroached subfacet. The proof of Lemma 6.3 shows that the diametral spheres of all linear edges of $\mathcal{Q}$ are empty at this point. Moreover, each facet F of $\mathcal{P}$ is the union of some triangles in the mesh, whose equatorial spheres are empty. But then it is known that the circumcenter of τ must lie inside the domain under these conditions (Lemma 3.2 in [3]), a contradiction. ◻

7 Termination of MESH

In this section, we prove a lower bound on the inter-vertex distance in terms of the local feature size with respect to $\mathcal{Q}$. This implies the termination of MESH by a packing argument.

7.1 Adjacent and non-incident elements

The major reason that the conventional Delaunay refinement strategy fails to handle small angles is that given two adjacent and non-incident elements E_1 and E_2, a vertex on E_2 encroaches upon some mesh element on E_1 and causes it to be split. But then this may repeat indefinitely. We prove that this phenomenon cannot happen in $\mathcal{Q}$ outside $\mathcal{B}$ because all angles are equal to $\pi/2$. We have three lemmas corresponding to the cases where E_1 is a curved edge, a flat facet outside $\mathcal{B}$, or a curved facet. Recall that there are two types of curved edges. A type 1 edge lies at the intersection between a protecting sphere and a facet of $\mathcal{P}$. A type 2 edge lies at the intersection between two adjacent protecting spheres.

LEMMA 7.1 *Let E_1 be a curved edge. Let E_2 be an element of $\mathcal{Q}$ on or outside $\mathcal{B}$. Assume that E_1 and E_2 are adjacent and non-incident. Then for any arc β on E_1 with angular width less than π, no vertex on $E_2 - \beta$ lies on or inside $K_{\beta}^{\perp}$.*

Proof. Suppose that E_1 is a type 1 edge on a protecting sphere S_x. Then E_2 must be a curved element lying on a protecting sphere S_y adjacent to S_x (including elements on $S_x \cap S_y$). Let R_2 be the union of rays that emits from x through S_y. Consider the smallest cap containing E_1 on S_x. Let R_1 be the union of rays that emits from x through this cap. Observe that E_1 lies outside R_2 (not even on the boundary of R_2). It follows that R_1 is a subset of the closure of $\mathbb{R}^3 - R_2$. Either $K_{\beta}^{\perp}$ lies strictly inside R_1 or $K_{\beta}^{\perp}$ touches the boundary of R_1 at an endpoint of β. Thus either $K_{\beta}^{\perp}$ lies strictly outside R_2 or $K_{\beta}^{\perp}$ touches the boundary of R_2 at an endpoint of β. Hence no vertex on $E_2 - \beta$ lies on or inside $K_{\beta}^{\perp}$.

Suppose that E_1 is a type 2 edge. The center of the support circle of E_1 lies in the interior of an edge h of $\mathcal{P}$. For each facet F of $\mathcal{P}$ incident to h, $K_{\beta}^{\perp}$ either avoids F or touches F at an endpoint of β. If E_1 has two distinct endpoints, they lie on two facets F_1 and F_2 of $\mathcal{P}$ that are incident to h. Note that either E_2 lies on F_1 or F_2, or E_2 is separated from E_1 by F_1 and F_2. Thus no vertex on $E_2 - \beta$ lies on or inside $K_{\beta}^{\perp}$. If E_1 has one endpoint only, E_2

must lie on the facet F_3 of $\mathcal{P}$ that is incident to h and passes through the endpoint of E_1. Thus no vertex on $E_2 - \beta$ lies on or inside $K_\beta^\perp$. If E_1 has no endpoint, E_2 cannot exist. ◻

LEMMA 7.2 *Let E_1 be a flat facet outside $\mathcal{B}$. Let E_2 be an element of $\mathcal{Q}$ on or outside $\mathcal{B}$ such that E_1 and E_2 are adjacent and non-incident. Assume that there is no wide or encroached helper arc. For any subfacet σ attached to E_1, no vertex on E_2 lies inside $K_\sigma^\perp$.*

Proof. Let H be the support plane of E_1. Let D be the intersection of H and the three-dimensional weighted Voronoi diagram of the normal spheres of subfacets attached to E_1, the normal spheres of helper arcs in ∂E_1, and the protecting spheres adjacent to E_1. D can also be viewed as the two-dimensional weighted Voronoi diagram of the circles at the intersections of H and the above spheres. Observe that the subfacets attached to E_1 are Voronoi regions in D (D may contain more Voronoi regions). The normal sphere of each subfacet owns the subfacet as a Voronoi region, and the normal sphere of each helper arc $\widehat{pq}$ owns the Voronoi edge pq. Assume to the contrary that a vertex $u \in E_2$ lies inside $K_\sigma^\perp$. Let u' be the orthogonal projection of u onto H. Let $\vec{s}$ be the directed line segment from u' to a vertex of σ. $\vec{s}$ crosses an ordered sequence of Voronoi cells. Let Σ be the corresponding sequence of spheres in the same order. Observe that the weighted distance of u from the spheres in Σ increases monotonically along Σ. Since all angles outside $\mathcal{B}$ are equal to $\pi/2$, u' lies on ∂E_1 or outside E_1. Thus at or before reaching σ, $\vec{s}$ intersects some Voronoi edge pq such that $\alpha = \widehat{pq}$ is a helper arc in ∂E_1. We have $\pi(u, K_\alpha^\perp) \leq \pi(u, K_\sigma^\perp) < 0$ as u lies inside $K_\sigma^\perp$. But then α is encroached, a contradiction. ◻

We proceed to the last case in which E_1 is a curved facet. It turns out that a vertex on E_2 may indeed encroach upon a helper triangle attached to E_1. Fortunately, they are still separated by a chain of edges of $\mathcal{P}$ in some sense, and so they are at distance at least some local feature size away. This will be sufficient to avoid indefinite splitting of mesh elements. To handle wide helper triangles whose normal spheres are undefined, we prove a slightly more general result. Note that the choice of the helper triangle in the statement of Lemma 7.3 is consistent with the rule 2 of MESH.

LEMMA 7.3 *Let S_x be a protecting sphere. Pick a helper triangle t on CH_x with preference for one such that K_t has angular diameter at least π. Let E_1 be the curved facet that t attaches to. Let E_2 be an element of $\mathcal{Q}$ on or outside $\mathcal{B}$ such that E_1 and E_2 are adjacent and non-incident. Let $K \subseteq K_t$ be a cap with angular diameter less than π and centered at the center v of K_t. Assume that there is no wide or encroached helper arc. If a vertex w on E_2 lies inside $K^\perp$, then vw intersects some edge of $\mathcal{P}$.*

Proof. Consider the case in which K_t has angular diameter less than π. By our choice of t, it follows that the circumcaps of all helper triangles on CH_x have angular diameter less than π. So their normal spheres are defined. Let D be the intersection of S_x and the three-dimensional weighted Voronoi diagram of the normal spheres of helper triangles on CH_x, the normal spheres of helper arcs on S_x, and the protecting spheres adjacent to S_x. Observe that each helper triangle projects radially onto a Voronoi cell in D. The rest of the proof proceeds as in the proof of Lemma 7.2. Any vertex u on E_2 projects to a point u' on S_x that lies on ∂E_1 or outside E_1. Assume to the contrary that u lies inside $K_t^\perp$. Then we can walk along the great circular arc from u' to a vertex of t to obtain the contradiction that some helper arc on the boundary of E_1 is encroached. Thus u lies outside $K_t^\perp$ and hence outside $K^\perp$ as well by Lemma 5.4.

Consider the case that K_t has angular diameter π or more. Suppose that E_2 is a curved element. If E_2 lies on S_x, no vertex on E_2 lies inside K by its emptiness. So no vertex on E_2 lies inside $K^\perp$. The other possibility is that E_2 lies on a protecting sphere S_y adjacent to S_x. Assume to the contrary that a vertex w on E_2 lies inside $K^\perp$. Then K must intersect $S_x \cap S_y$ at some empty arc α. As the helper arc containing α is not wide, the angular width of α is less than π. By Lemmas 5.2 and 5.3, $K_\alpha^\perp$ is the orthogonal sphere of S_y at $K^\perp \cap S_y$. So the centers of $K^\perp$, $K_\alpha^\perp$, and S_y are collinear. Since v lies outside S_y, y does not lie between the centers of $K^\perp$ and $K_\alpha^\perp$. So $K_\alpha^\perp$ encloses the part of

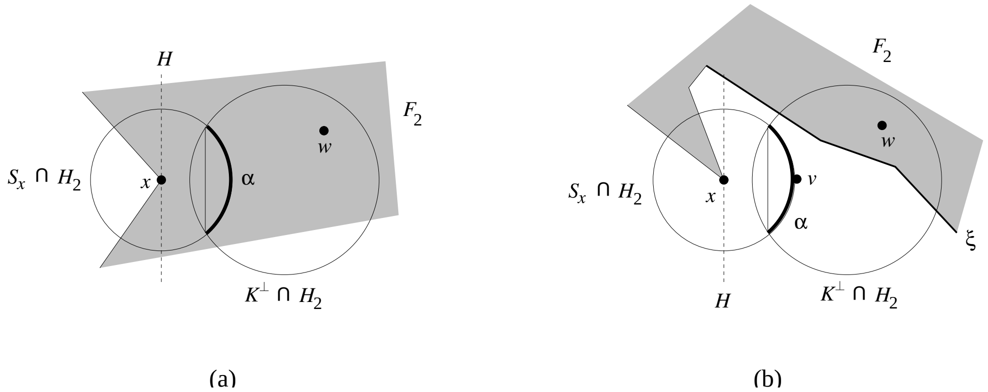

Figure 10: The figures show the cross-sections on H_2. The shaded region represents F_2.

S_y that lies inside $K^\perp$, which implies that w lies inside $K_\alpha^\perp$. But then the helper arc containing α is encroached, a contradiction.

Suppose that E_2 is a flat facet outside $\mathcal{B}$. Let H be the plane that passes through x and is parallel to the plane containing ∂K_t. Let H^+ denote the side of H that contains the center of K_t. So there is no vertex on $S_x \cap H^+$. Consider a linear edge e incident to x. It goes through a hole C on S_x. So if e does not lie strictly outside H^+, $C \cap H^+$ is a half circle or more. But then the helper arc containing $C \cap H^+$ is wide, a contradiction. So all linear edges incident to x lie strictly outside H^+, which implies that x is a vertex of $\mathcal{P}$. Let $K \subseteq K_t$ be a cap with angular diameter less than π and centered at the center v of K_t. Let F_2 be the facet of $\mathcal{P}$ containing E_2. Note that $x \in \partial F_2$. Let H_2 be the support plane of F_2. By Claim 1 in [4], the two circles $H_2 \cap K^\perp$ and $H_2 \cap S_x$ intersect at right angle. The arc $\alpha = K \cap H_2$ contains no vertex by the emptiness of K. Since all linear edges incident to x lie outside H^+, either α lies on a type 1 curved edge on F_2 or α lies outside F_2. See Figure 10 for the two situations. If no vertex on E_2 lies inside $K^\perp$, we are done. Assume that a vertex w on E_2 lies inside $K^\perp$. We have $w \in H^+$ as $K^\perp \subseteq H^+$. In Figure 10(a), w lies inside $K^\perp \cap H_2$. Since x is a vertex of $\mathcal{P}$ and $H_2 \cap K^\perp$ intersects $H_2 \cap S_x$ at right angle, $K^\perp \cap H_2 = K_\alpha^\perp \cap H_2$. But then the helper arc containing α is encroached by w, a contradiction. In Figure 10(b), $K^\perp \cap H_2$ intersects a chain ξ of boundary edges of F_2 in H^+ in order to enclose w. Since x does not lie on any edge on ξ, ξ lies outside S_x, which implies that ξ separates v and w within $K^\perp \cap H_2$. Hence vw intersects some edge on ξ. ◻

7.2 Notations

We need some notations to prove the lower bound on the inter-vertex distances. For each vertex $v \in \mathcal{V}$, we define the *insertion radius* of v as follows. If v is a vertex of $\mathcal{Q}$, r_v is the minimum distance from v to another vertex of $\mathcal{Q}$. If v is inserted or rejected by MESH, r_v is the minimum distance to a vertex in $\mathcal{V}$ at the time when v is inserted or rejected.

Consider the time when MESH inserts or rejects a vertex v using rule i, $1 \leq i \leq 4$. We say that v *has type* i and we define the *parent* of v as follows.

- If v is the midpoint of a wide helper arc or the circumcap center of a wide helper triangle, the parent of v is undefined.

- Suppose that v is the midpoint of an encroached helper arc or v is the circumcap center of an encroached helper triangle or subfacet. If $\mathcal{V}$ has a vertex encroaching upon σ, the parent of v is its nearest encroaching vertex in $\mathcal{V}$. Otherwise, $K_\sigma^\perp$ is empty. What happens is that MESH rejected a vertex p for encroaching upon σ and this also prompted MESH to consider v. The parent of v is p in this case.

- Suppose that v is the circumcenter of a tetrahedron τ. Let e be the shortest edge of τ. The parent of v is the endpoint of e that appeared in $\mathcal{V}$ the latest.

Finally, the parents of vertices of $\mathcal{Q}$ are undefined.

For any point $x \in \mathbb{R}^3$, the local feature size $\widehat{f}(x)$ at x with respect to $\mathcal{Q}$ is the radius of the smallest ball that intersects two disjoint elements of $\mathcal{Q}$.

7.3 Lower bound on insertion radii

The proof consists of two steps. We first show a recurrence relation between the insertion radii of a vertex and its parent. Then we apply induction to lower bound the insertion radius in terms of $\widehat{f}$. We need the following technical lemma.

LEMMA 7.4 *Let K be a cap with angular diameter at most $\pi/3$. Let v be the center of K. For any point p inside $K^{\perp}$ and any point q on or outside $K^{\perp}$, $\|q-v\| > (1/4) \cdot \max\{\|p-v\|, \|p-q\|\}$.*

Proof. Let z be the center of $K^{\perp}$. Refer to Figure 11. Since the angular diameter of K is at most $\pi/3$, $\|v-z\| <$

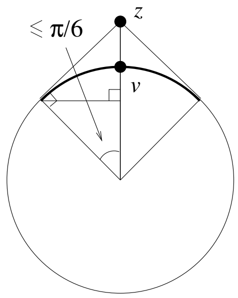

Figure 11: The bold arc represents K.

$\text{radius}(K^{\perp}) \cdot \sin(\pi/6) \leq \|q-z\| \cdot (1/2)$. Substituting this into the triangle inequality $\|q-v\| \geq \|q-z\| - \|v-z\|$, we obtain

$$\|q-v\| > \|q-z\|/2. \tag{5}$$

Since p and v lie inside $K^{\perp}$, $\|q-z\| \geq \|p-v\|/2$. Substituting this into (5) yields $\|q-v\| > \|p-v\|/4$, which proves part of the lemma. Starting with the triangle inequality, we have $\|p-q\| \leq \|p-z\| + \|q-z\| \leq 2 \cdot \|q-z\|$. Substituting this into (5) yields $\|q-v\| > \|p-q\|/4$, which proves the other part of the lemma. ◻

We are ready to develop the recurrence involving the insertion radii of a vertex and its parent.

LEMMA 7.5 *Let v be a vertex of $\mathcal{Q}$ or a vertex inserted or rejected by* MESH. *Let p be the parent of v.*

(i) If p is undefined, $r_v \geq \widehat{f}(v)/2$.

(ii) Otherwise, $r_v \geq \|p-v\|/4$. If $r_v < \widehat{f}(v)/4$, then

- *(a) if v has type 1, then p has type 2, 3 or 4 and $r_v \geq r_p/4$;*
- *(b) if v has type 2 or 3, then p has type 4 and $r_v \geq r_p/4$; and*
- *(c) if v has type 4, then $r_v \geq \rho_0 \cdot r_p$.*

Proof. Go back to the time when MESH considered v. If v is a vertex of $\mathcal{Q}$, p is undefined and $r_v \geq \widehat{f}(v)$ by definition. We analyze the other cases below.

Case 1: v is the midpoint of a wide helper arc α. p is undefined and we are to show that (i) holds. Let β be the subarc of α with midpoint v and angular width $\pi/3$. Let B be the smallest ball centered at v that contains β. Let x be the center of the support circle of α. We have $\mathrm{radius}(B) = \|v - x\| \cdot 2\sin(\pi/12) > \|v - x\|/2$. Observe that x lies on some linear edge of $\mathcal{Q}$ which is disjoint from the curved edge containing α. Thus $\|v - x\| \geq \widehat{f}(v)$ which implies that $\mathrm{radius}(B) \geq \widehat{f}(v)/2$.

If B does not contain any vertex in $\mathcal{V}$, then $r_v \geq \mathrm{radius}(B) \geq \widehat{f}(v)/2$. Assume that B contains a vertex in $\mathcal{V}$. Let w be the vertex inside B closest to v. By definition, $r_v = \|v - w\|$. If β lies on a type 1 curved edge, then x is a vertex of $\mathcal{P}$ and the equatorial sphere of β is the protecting sphere S_x. As the angular width of β is $\pi/3$, B does not contain x, which is the only vertex inside S_x. If β lies on a type 2 curved edge, the equatorial sphere of β contains no vertex. Therefore, we conclude in both cases that w lies inside $K_\beta^\perp$. Since α is a wide helper arc, MESH has split helper arcs only so far. Thus either w is a vertex of $\mathcal{Q}$ disjoint from the curved edge containing β, or w lies on some curved edge E of $\mathcal{Q}$. In the latter case, it follows from Lemma 7.1 that E is disjoint from the curved edge containing first appeared in $\mathcal{V}$g v. Hence $r_v = \|v - w\| \geq \widehat{f}(v)$.

Case 2: v is the center of K_t for a wide helper triangle t. p is undefined and we are to show that (i) holds. Let K be the cap inside K_t centered at v with angular diameter $\pi/3$. Let B be the smallest ball centered at v that contains K. If B does not contain any vertex in $\mathcal{V}$, we can show that $r_v \geq \widehat{f}(v)/2$ as in case 1. Assume that B contains a vertex in $\mathcal{V}$. Let w be the vertex inside B closest to v. By definition, $r_v = \|v - w\|$. K lies on some protecting sphere S_x, and x is the only vertex inside S_x. As the angular diameter of K is $\pi/3$, B does not contain x. Thus w lies inside $K^\perp$. Since t is a wide helper triangle, MESH has split helper arcs and helper triangles only so far. Thus either w is a vertex of $\mathcal{Q}$ not on S_x, or w lies on some curved element E' of $\mathcal{Q}$. Clearly, in the first case, $r_v = \|v - w\| \geq \widehat{f}(v)$. In the second case, recall that v lies in the interior of the curved facet E that t attaches to. E' cannot be a boundary curved edge of E as the emptiness of K would preclude the presence of w inside $K^\perp$. By Lemma 7.3, either E and E' are disjoint, or vw intersects some edge of $\mathcal{P}$. In either case, $r_v = \|v - w\| \geq \widehat{f}(v)$.

Case 3: v is the midpoint of an encroached helper arc α. v has type 1. p lies inside $K_\alpha^\perp$ as it encroaches upon α. Let q be the vertex in $\mathcal{V}$ such that $r_v = \|q - v\|$. If q lies on or inside $K_\alpha^\perp$, then q encroaches upon α, which implies that $p = q$ and $\|q - v\| = \|p - v\|$. Otherwise, by Lemma 7.4, $\|q - v\| > \|p - v\|/4$. This proves that $r_v \geq \|p - v\|/4$. Next we relate r_v to $\widehat{f}(v)$ and r_p. Let e be the curved edge containing α. If p is a vertex of $\mathcal{Q}$, for p to lie inside $K_\alpha^\perp$, p cannot be an endpoint of e. So $\|p - v\| \geq \widehat{f}(v)$. Hence $r_v \geq \|p - v\|/4 \geq \widehat{f}(v)/4$ and we are done. Assume that p is not a vertex of $\mathcal{Q}$.

- Case 3.1: p lies on an element E of $\mathcal{Q}$ that is non-incident to e. Then Lemma 7.1 implies that e and E are disjoint. So $\|p - v\| \geq \widehat{f}(v)$. Hence $r_v \geq \|p - v\|/4 \geq \widehat{f}(v)/4$.
- Case 3.2: Either a facet of $\mathcal{Q}$ incident to e contains p, or p has type 4. So p has type 2, 3, or 4, which implies that MESH rejects p. Since q is a vertex in $\mathcal{V}$ at this time, $q \neq p$ and $r_p \leq \|p - q\|$. As $q \neq p$, q lies outside $K_\alpha^\perp$ and Lemma 7.4 implies that $\|q - v\| > \|p - q\|/4 \geq r_p/4$. Hence $r_v = \|q - v\| \geq r_p/4$. This proves (ii)(a).

Case 4: v is the center of K_σ where σ is an encroached helper triangle or subfacet. v has type 2 or 3. We can show that $r_v \geq \|p - v\|/4$ as in case 3. Since σ is encroached by p, p lies inside $K_\sigma^\perp$. Let E be the curved facet or flat facet that σ attaches to. If p is a vertex of $\mathcal{Q}$ or p has type 1, 2 or 3, then $p \in \mathcal{V}$ as MESH cannot reject p for encroaching σ. In this case, p lies on some element E' of $\mathcal{Q}$. E' cannot be a boundary curved edge of E as the emptiness of K_σ would preclude the presence of p inside $K_\sigma^\perp$. Then by Lemmas 7.2 and 7.3, either E and E' are disjoint, or vp intersects some edge of $\mathcal{P}$. In either case, $\|p - v\| \geq \widehat{f}(v)$. Hence $r_v \geq \|p - v\|/4 \geq \widehat{f}(v)/4$ and we are done. The remaining case is that p has type 4. In this case, MESH rejects p for encroaching σ. Then we can show that $r_v \geq r_p/4$ as in case 3.2. This proves (ii)(b).

Case 5: v is the circumcenter of a tetrahedron τ. The shortest edge of τ is pq for some vertex $q \in \mathcal{V}$. If p is a vertex of $\mathcal{Q}$, q is also a vertex of $\mathcal{Q}$ as q appears in $\mathcal{V}$ no later than p. Thus $r_v = \|p - v\| = \|q - v\| = \widehat{f}(v)$ and we are done. If p is not a vertex of Q, then since $\rho(\tau) > \rho_0$, $r_v = \|p - v\| > \rho_0 \cdot \|p - q\| \geq \rho_0 r_p$. This proves (ii)(c). ∎

Next, we lower bound r_v in terms of $\widehat{f}(v)$ by induction. We define four constants $C_1 = 84\rho_0/(\rho_0 - 16)$, $C_2 = C_3 = (20\rho_0 + 16)/(\rho_0 - 16)$, and $C_4 = (4\rho_0 + 20)/(\rho_0 - 16)$. Note that $C_1 > C_2 = C_3 > C_4 > 4$ for $\rho_0 > 16$.

LEMMA 7.6 *Let v be a vertex. If v is a vertex of $\mathcal{Q}$, then $r_v \geq \widehat{f}(v)$. Otherwise, if v has type i, then $r_v > \widehat{f}(v)/C_i$.*

Proof. We prove the lemma by induction on the order of vertex insertions. At the beginning, $r_v \geq \widehat{f}(v)$ for each vertex v of $\mathcal{Q}$. In the induction step, if $r_v > \widehat{f}(v)/4$, we are done as $C_4 > 4$. Otherwise, Lemma 7.5 implies that the parent p of v is defined and $r_v \geq \|p - v\|/4$. Then the Lipschitz property implies that

$$\widehat{f}(v) \leq \widehat{f}(p) + \|p - v\| < \widehat{f}(p) + 4r_v \leq \widehat{f}(p) \cdot \frac{r_v}{c \cdot r_p} + 4r_v, \tag{6}$$

provided that $r_v \geq c \cdot r_p$ for some constant c.

If v has type 1, by Lemma 7.5(a), p has type 2, 3 or 4 and $r_v \geq r_p/4$. By induction assumption, $\widehat{f}(p) \leq C_2 r_p$. Substituting these into (6) yields $\widehat{f}(v) \leq 4C_2 r_v + 4r_v = C_1 r_v$. If v has type 2 or 3, by Lemma 7.5(b), p has type 4 and $r_v \geq r_p/4$. By induction assumption, $\widehat{f}(p) \leq C_4 r_p$. Substituting these into (6) yields $\widehat{f}(v) \leq 4C_4 r_v + 4r_v = C_2 r_v$. If v has type 4, then $r_v \geq \rho_0 r_p$ by Lemma 7.5(c). By induction assumption, $\widehat{f}(p) \leq C_1 r_p$ regardless of the type of p. Substituting these into (6) yields $\widehat{f}(v) \leq C_1 r_v/\rho_0 + 4r_v = C_4 r_v$. ∎

7.4 Termination

We use Lemma 7.6 and the Lipschitz property to lower bound the inter-vertex distances. Then the termination of MESH follows by a packing argument.

LEMMA 7.7 MESH *terminates. For each output vertex v, its shortest incident edge has length at least $\widehat{f}(v)/(1+C_1)$.*

Proof. Let vw be any edge incident to v. If w appeared in $\mathcal{V}$ no later than v, then $\|v - w\| \geq r_v \geq \widehat{f}(v)/C_1$ by Lemma 7.6. If v appeared in $\mathcal{V}$ before w, then $\|v - w\| \geq r_w \geq \widehat{f}(w)/C_1$ by Lemma 7.6. Using the Lipschitz condition, we get $\widehat{f}(v) \leq \widehat{f}(w) + \|v - w\| \leq (1 + C_1) \cdot \|v - w\|$. The edge length bound implies that we can center disjoint balls at the output vertices with radii $\widehat{f}_{\min}/(2 + 2C_1)$, where $\widehat{f}_{\min}$ is the minimum local feature size with respect to $\mathcal{Q}$. By Lemma 6.4, all vertices inserted by MESH do not lie outside the union of the domain and the buffer zone. Since $\widehat{f}_{\min} > 0$ and the union of the domain and the buffer zone has bounded volume, there is a finite number of output vertices. It follows that MESH terminates. ∎

8 Mesh quality

By Lemma 7.7 and rule 4, all tetrahedra that do not lie inside $\mathcal{B}$ have radius-edge ratio bounded by ρ_0. In this section, we prove the gradedness and bound the radius-edge ratio of tetrahedra inside $\mathcal{B}$. Our proof consists of three steps. First, we analyze some lengths concerning linear edges and $\mathcal{B}$. Second, we use these results to prove that for all output vertices p, $g(p) = \Omega(f(p))$. Third, we show that $\widehat{f}(p) = \Omega(f(p))$. Then the gradedness result follows. Bounding the radius-edge ratio of tetrahedra inside $\mathcal{B}$ only requires a little extra effort.

8.1 Length properties of linear edges and $\mathcal{B}$

We first recall a few definitions in Sections 3.1 and 3.2. μ is a constant chosen from $(0, \frac{1}{7}]$. ϕ denotes the smallest angle in $\mathcal{P}$. Let uv be an edge of $\mathcal{P}$. The edge uv is recursively split into linear edges of $\mathcal{Q}$. If x is a linear edge endpoint in $\mathrm{int}(uv)$, $S_x \cap \mathcal{B}$ is called a ring, and it has two parallel holes. The *width* of a ring is equal to the distance between the two planes containing the holes.

We define ϕ_{uv}^u to be the smallest angle between uv and an edge/facet of $\mathcal{P}$ incident to u, and $\theta_{uv}^u = \min\{\pi/3, \phi_{uv}^u\}$. ϕ_{uv}^v and θ_{uv}^v are symmetrically defined. The recursive splitting of uv starts after placing the points u_v and v_u, and constructing the protecting spheres S_{u_v} and S_{v_u}:

$$\begin{aligned}
\|u - u_v\| &= \mu \sec(\mu\theta_{uv}^u) \cdot g(u) \\
\mathrm{radius}(S_{u_v}) &= \|u - u_v\| \cdot \sin(\mu\theta_{uv}^u) \\
\|v - v_u\| &= \mu \sec(\mu\theta_{uv}^v) \cdot g(v) \\
\mathrm{radius}(S_{v_u}) &= \|v - v_u\| \cdot \sin(\mu\theta_{uv}^v)
\end{aligned}$$

LEMMA 8.1 *Let c_3 be the constant in Lemma 3.3. There exists a constant $c_4 < 1$ such that for any linear edge xy of $\mathcal{Q}$, $\|x - y\| > c_3\mu \cdot \max\{g(x), g(y)\}$ and $g(y) \geq c_4\mu \cdot g(x)$.*

Proof. First, we claim that for any linear edge endpoint z,

$$c_3\mu \cdot g(z) \leq \mathrm{radius}(S_z) \leq 3\mu \cdot g(z). \tag{7}$$

If z is a vertex of $\mathcal{P}$, then $g(z) = f(z)$ and $\mathrm{radius}(S_z) = \mu \cdot f(z)$ by construction. If z is not a vertex of $\mathcal{P}$, then $c_3\mu \cdot g(z) \leq \mathrm{radius}(S_z) \leq 3\mu \cdot g(z)$ by Lemmas 3.1 and 3.3(ii).

Since S_x and S_y are orthogonal, $\|x - y\| > \max\{\mathrm{radius}(S_x), \mathrm{radius}(S_y)\} \geq c_3\mu \cdot \max\{g(x), g(y)\}$ by (7). This proves the first part of the lemma. We prove the other part of the lemma for $c_4 = (1/c_1)\min\{\sqrt{3}/2, \sin\phi\}$, where $c_1 > 1$ is the constant in Lemma 3.1. There are two cases to analyze.

- $x = u$ or v. Assume that $x = u$. The case where $x = v$ can be handled similarly. By elementary trigonometry and the definition of θ_{uv}^u, $\tan(\mu\theta_{uv}^u) \geq \mu \sin\theta_{uv}^u \geq \mu \cdot \min\{\sqrt{3}/2, \sin\phi\} = c_1c_4\mu$. Since $y = u_v$, $\mathrm{radius}(S_y) = \mu\tan(\mu\theta_{uv}^u) \cdot g(x) \geq c_1c_4\mu^2 \cdot g(x)$. By Lemma 3.1, $\mathrm{radius}(S_y) \leq c_1\mu \cdot g(y)$. Hence $g(y) \geq c_4\mu \cdot g(x)$.

- $x \in \mathrm{int}(uv)$. We have $\|x - y\| \leq \mathrm{radius}(S_x) + \mathrm{radius}(S_y)$, which is at most $3\mu(g(x) + g(y))$ by (7). By Lemma 2.1, $g(x) \leq g(y) + \|x - y\|$. So $g(x) \leq (1+3\mu) \cdot g(y) + 3\mu \cdot g(x)$, which implies that $g(y) \geq \frac{1-3\mu}{1+3\mu} \cdot g(x)$. It can be verified that $\frac{1-3\mu}{1+3\mu} > c_4\mu$.

$\square$

LEMMA 8.2 *Let S_x be a protecting sphere. There exist constants $c_7 < c_6 < c_5 < c_4$ such that:*

(i) The radius of any hole on $\mathcal{B} \cap S_x$ is at least $c_5\mu^2 \cdot g(x)$.

(ii) If $\mathcal{B} \cap S_x$ is a ring, its width is at least $c_6\mu^2 \cdot g(x)$.

(iii) If E is a vertex, edge or facet of $\mathcal{P}$ disjoint from x, the minimum distance between S_x and E is at least $(1 - 3\mu) \cdot g(x)$.

(iv) Let S_y be a protecting sphere that is not adjacent to S_x. The minimum distance between $\mathcal{B} \cap S_x$ and $\mathcal{B} \cap S_y$ is at least $c_7\mu^3 \cdot g(x)$.

Proof. We prove the lemma for the constants $c_5 = c_3c_4/\sqrt{2}$, $c_6 = c_3^2c_4/(3+3c_4)$, and $c_7 = c_4c_6$.

Consider (i). A hole on $\mathcal{B} \cap S_x$ is equal to $S_x \cap S_z$ for some protecting sphere S_z adjacent to S_x. By (7), we have $\min\{\mathrm{radius}(S_x), \mathrm{radius}(S_z)\} \geq c_3\mu \cdot \min\{g(x), g(z)\}$, which is at least $c_3c_4\mu^2 \cdot g(x)$ by Lemma 8.1. Since S_x intersects S_z at right angle, $\mathrm{radius}(S_x \cap S_z) \geq \min\{\mathrm{radius}(S_x), \mathrm{radius}(S_z)\}/\sqrt{2} \geq (c_3c_4\mu^2/\sqrt{2}) \cdot g(x)$.

Consider(ii). Let S_z be a protecting sphere adjacent to S_x. Let d be the distance between x and the bisector plane of S_x and S_z. The width of $\mathcal{B} \cap S_x$ is at least d. Since S_x and S_z are orthogonal, we have

$$\begin{aligned} d &= \mathrm{radius}(S_x)^2/\|x-z\| &(8)\\ &\overset{(7)}{\geq} (c_3\mu \cdot g(x))^2/\|x-z\|. &(9) \end{aligned}$$

By (7), $\|x-z\| \leq \mathrm{radius}(S_x) + \mathrm{radius}(S_z) \leq 3\mu \cdot g(x) + 3\mu \cdot g(z)$. By Lemma 8.1, $g(x) \geq c_4\mu \cdot g(z)$. So $\|x-z\| \leq \frac{3+3\mu c_4}{c_4} \cdot g(x)$. Plugging into (9), we get $d \geq \frac{c_3^2c_4\mu^2}{3+3\mu c_4} \cdot g(x) \geq c_6\mu^2 \cdot g(x)$.

Consider (iii). The distance between S_x and E is at least $g(x) - \mathrm{radius}(S_x)$, which is at least $(1-3\mu) \cdot g(x)$ by (7).

Consider (iv). Let d be the minimum distance between $\mathcal{B} \cap S_x$ and $\mathcal{B} \cap S_y$. Suppose that x and y do not lie on the same edge of $\mathcal{P}$. Then $\|x-y\| \geq \max\{g(x), g(y)\} \geq (g(x)+g(y))/2$. Thus

$$\begin{aligned} d &\geq \|x-y\| - \mathrm{radius}(S_x) - \mathrm{radius}(S_y)\\ &\overset{(7)}{\geq} \|x-y\| - 3\mu(g(x)+g(y))\\ &\geq (1-6\mu) \cdot \|x-y\|\\ &\geq (1-6\mu) \cdot g(x). \end{aligned}$$

Suppose that x and y lie on the same edge of $\mathcal{P}$. There is a ring $\mathcal{B} \cap S_z$ between $\mathcal{B} \cap S_x$ and $\mathcal{B} \cap S_y$ that is adjacent to $\mathcal{B} \cap S_x$. The distance between $\mathcal{B} \cap S_x$ and $\mathcal{B} \cap S_y$ is at least $\mathrm{width}(\mathcal{B} \cap S_z)$. By Lemma 8.1, $g(z) \geq c_4\mu \cdot g(x)$. Then by (ii), $\mathrm{width}(\mathcal{B} \cap S_z) \geq c_6\mu^2 \cdot g(z) \geq c_4c_6\mu^3 \cdot g(x)$. ◻

8.2 Lower bound g in terms of f

We first lower bound $g(p)$ for the special case in which p lies in the middle of the edges of $\mathcal{P}$.

LEMMA 8.3 *Let p be a point on an edge uv of $\mathcal{P}$. If $\|p-u\| \geq \frac{\mu}{2} \cdot f(u)$ and $\|p-v\| \geq \frac{\mu}{2} \cdot f(v)$, then $g(p) \geq k_1\mu \cdot f(p)$, where $k_1 = \sin\phi/3$.*

Proof. Let B be the ball centered at p with radius $g(p)$. If B intersects two disjoint elements of $\mathcal{P}$, $g(p) = f(p)$. Otherwise, we can assume that B touches u or the interior of an edge/facet of $\mathcal{P}$ incident to u. So

$$g(p) \geq \|p-u\| \cdot \sin\phi. \tag{10}$$

By $\|p-u\| \geq \frac{\mu}{2} \cdot f(u)$ and the Lipschitz condition $f(p) \leq f(u) + \|p-u\|$, we get $f(p) \leq \frac{2+\mu}{\mu} \cdot \|p-u\|$. Plugging this into (10) yields $g(p) \geq \frac{\mu\sin\phi}{2+\mu} \cdot f(p) \geq k_1\mu \cdot f(p)$. ◻

We need a technical result on the distance between p and the edges of $\mathcal{P}$ when p lies on or outside $\mathcal{B}$.

LEMMA 8.4 *Let p be a point on or outside $\mathcal{B}$. Let q be the closest point to p on an edge of $\mathcal{P}$. Then $\|p-q\| \geq k_2\mu^3 \cdot f(p)$, where $k_2 = k_1c_5/(1+3\mu+k_1c_5\mu^3)$.*

Proof. Assume that q lies on the edge uv of $\mathcal{P}$. Since p lies on or outside $\mathcal{B}$, we have

$$\|p-u\| \geq \mathrm{radius}(S_u) = \mu \cdot f(u). \tag{11}$$

If $q = u$ or v, say u, then by plugging (11) into the Lipschitz condition $f(p) \leq f(u) + \|p-u\|$, we get $\|p-q\| = \|p-u\| \geq \frac{\mu}{1+\mu} \cdot f(p) \geq k_2\mu^3 \cdot f(p)$. Suppose that $q \in \mathrm{int}(uv)$. There are two cases.

- $\|q-u\| < \frac{\mu}{2} \cdot f(u)$. The case where $\|q-v\| < \frac{\mu}{2} \cdot f(v)$ can be handled similarly. By (11), we have $\|q-u\| \leq \|p-u\|/2$. Thus $\|p-u\| \leq \|p-q\| + \|q-u\| \leq \|p-q\| + \|p-u\|/2$, which implies that

$$\|p-q\| \geq \|p-u\|/2. \tag{12}$$

  Plugging (11) into the Lipschitz condition $f(p) \leq f(u) + \|p-u\|$, we get $f(p) \leq \frac{1+\mu}{\mu} \cdot \|p-u\|$. Plugging this into (12) yields $\|p-q\| \geq \frac{\mu}{2+2\mu} \cdot f(p) \geq k_2\mu^3 \cdot f(p)$.

- $\|q-u\| \geq \frac{\mu}{2} \cdot f(u)$ and $\|q-v\| \geq \frac{\mu}{2} \cdot f(v)$. Let H be the plane orthogonal to uv and passing through p and q. H cuts $S_x \cap \mathcal{B}$ for some $x \in uv$. Since q lies inside S_x, $\|q-x\| \leq 3\mu \cdot g(x)$ by (7). By Lemma 2.1,

$$g(q) \leq g(x) + \|q-x\| \leq (1+3\mu) \cdot g(x). \tag{13}$$

  $\|p-q\|$ is no less than the radius of the circle $H \cap \mathcal{B} \cap S_x$, which is in turn no less than the radius of a hole on $\mathcal{B} \cap S_x$ that uv goes through. Therefore, (13) and Lemma 8.2(i) yield $\|p-q\| \geq \frac{c_5\mu^2}{1+3\mu} \cdot g(q)$. By Lemma 8.3, we get $\|p-q\| \geq \frac{k_1c_5\mu^3}{1+3\mu} \cdot f(q)$. Substituting this into the Lipschitz condition $f(p) \leq f(q) + \|p-q\|$, we get $\|p-q\| \geq \frac{k_1c_5\mu^3}{1+3\mu+k_1c_5\mu^3} \cdot f(p)$.

$\square$

We are ready to lower bound g in terms of f for points of interest to us.

LEMMA 8.5 *For each point p, if p is an output vertex inside $\mathcal{B}$ or p lies on or outside $\mathcal{B}$, then $g(p) \geq k_3\mu^3 \cdot f(p)$, where $k_3 = k_2 \sin(\phi/2)$.*

Proof. If p is a vertex of $\mathcal{P}$, then $g(p) = f(p)$ and we are done. Suppose that p lies inside $\mathcal{B}$. Then p is a linear edge endpoint in the interior of some edge uv of $\mathcal{P}$. Observe that p lies outside S_u and S_v, and so $\|p-u\| \geq \mu \cdot f(u)$ and $\|p-v\| \geq \mu \cdot f(v)$. By Lemma 8.3, we get $g(p) \geq k_1\mu \cdot f(p)$.

Suppose that p lies on or outside $\mathcal{B}$. Let B be the ball centered at p with radius $g(p)$. If B intersects two disjoint elements of $\mathcal{P}$, then $g(p) = f(p)$. Suppose not. If B intersects an edge uv of $\mathcal{P}$, then let q be the point on uv closest to p. Using Lemma 8.4, we get $g(p) \geq \|p-q\| \geq k_2\mu^3 \cdot f(p)$. The remaining case is that B intersects the interior of two adjacent facets F_1 and F_2 of $\mathcal{P}$. Let H_i be the plane containing F_i. Let r be the point in the line $H_1 \cap H_2$ closest to p. Since the angle between H_1 and H_2 is at least ϕ, pr makes an angle at least $\phi/2$ with H_1 or H_2. Thus

$$g(p) \geq \|p-r\| \cdot \sin(\phi/2). \tag{14}$$

Project pr orthogonally onto H_1 and H_2. Since the interior of F_1 and F_2 do not intersect, the projections must intersect a boundary edge e of F_1 or F_2 at a point s. Since $\|p-s\| \leq \|p-r\|$, (14) yields $g(p) \geq \|p-s\| \cdot \sin(\phi/2)$. Clearly, the distance from p to the edge e is no greater than $\|p-s\|$. Hence Lemma 8.4 implies that $g(p) \geq k_2\mu^3 \sin(\phi/2) \cdot f(p)$. $\square$

8.3 Lower bound $\widehat{f}$ in terms of f

We are to show that $\widehat{f}(p) = \Omega(f(p))$ for all points p on or outside $\mathcal{B}$. The analysis involves considering how a ball B centered at p with radius $\widehat{f}(p)$ intersects the elements of $\mathcal{Q}$.

Technical lemma. We need a technical lemma stating that if B intersects a protecting sphere S_x such that $\mathrm{radius}(B) = \Omega(g(x))$, then $\mathrm{radius}(B) = \Omega(f(p))$.

LEMMA 8.6 *Let B be a ball centered at a point p on or outside $\mathcal{B}$. Assume that B intersects a protecting sphere S_x such that $\mathrm{radius}(B) \geq c \cdot g(x)$ for some constant c. Then $\mathrm{radius}(B) \geq k\mu^3 \cdot f(p)$, where $k = ck_3/(1+c+3\mu)$.*

Proof. Let A be the ball centered at p with radius $g(x) + \|p-x\|$. A intersects the two elements of $\mathcal{P}$ that define $g(x)$. Since p lies on or outside $\mathcal{B}$, p does not lie on any edge of $\mathcal{P}$. Thus at most one of the elements of $\mathcal{P}$ that intersect A contains p. It follows that

$$g(p) \leq \mathrm{radius}(A) = g(x) + \|p - x\| \tag{15}$$

As B intersects S_x, $\|p-x\| \leq \mathrm{radius}(B) + \mathrm{radius}(S_x)$. By (7), we get $\|p-x\| \leq \mathrm{radius}(B) + 3\mu \cdot g(x)$. Plugging this into (15) yields $g(p) \leq \mathrm{radius}(B) + (1+3\mu) \cdot g(x)$. By the assumption that $\mathrm{radius}(B) \geq c \cdot g(x)$, we get $g(p) \leq \frac{1+c+3\mu}{c} \cdot \mathrm{radius}(B)$. Finally, Lemma 8.5 implies that $\mathrm{radius}(B) \geq \frac{ck_3\mu^3}{1+c+3\mu} \cdot f(p)$. ◻

Two critical cases. There are two critical cases in our analysis which we deal with separately. One critical case is when B intersects two disjoint curved elements E and E' on two protecting spheres S_x and S_y such that x is a vertex of $\mathcal{P}$, and S_x and S_y are identical or adjacent. The other case is similar except that E' is a flat facet. We prove lower bound on the distance between E and E' in these two cases. We use $d(X,Y)$ to denote the distance between two objects X and Y.

LEMMA 8.7 *Let E and E' be two disjoint curved elements of $\mathcal{Q}$. Let S_x and S_y be two protecting spheres that contain E and E', respectively. Assume that x is a vertex of $\mathcal{P}$, and S_x and S_y are identical or adjacent. Then $d(E,E') = \Omega(g(x))$.*

Proof. Let $p \in E$ and $q \in E'$ be points such that $\|p-q\| = d(E,E')$. If E is a curved facet, we claim that $p \notin \mathrm{int}(E)$. Otherwise, pq is normal to S_x. If S_x and S_y are orthogonal, then q cannot lie on $S_y \cap \mathcal{B}$, a contradiction. If $S_x = S_y$, pq is a diameter of S_x. But then we can rotate pq at q to decrease its length, a contradiction. Thus p lies on a curved edge $e \subseteq E$ ($e = E$ if E is a curved edge). Similarly, q lies on a curved edge $e' \subseteq E'$. Recall that there are two types of curved edges. A type 1 edge lies at the intersection between a protecting sphere and a facet of $\mathcal{P}$. A type 2 edge lies at the intersection between two adjacent protecting spheres. There are two cases to consider.

Case 1: $S_x \neq S_y$. So $\mathcal{B} \cap S_y$ is a ring, and y lies on an edge h of $\mathcal{P}$ incident to x. Suppose that e is a type 2 edge. Then e lies on $S_x \cap S_z$ for some protecting sphere S_z adjacent to S_x. If $S_y \neq S_z$, then by Lemma 8.2(iv), $d(\mathcal{B}\cap S_y, \mathcal{B} \cap S_z) \geq c_7\mu^3 \cdot g(y)$. This is a lower bound for $d(E,E')$ too. By Lemma 8.1, we get $d(E,E') \geq c_4c_7\mu^4 \cdot g(x)$. Suppose that $S_y = S_z$. That is, e lies on $S_x \cap S_y$. Since E and E' are disjoint, there are two possibilities. First, e and e' lie on different holes of $S_y \cap \mathcal{B}$. Second, e and e' are separated by two facets of $\mathcal{P}$ incident to h. In the first case, $d(E,E') \geq \mathrm{width}(S_y \cap \mathcal{B})$, which is at least $c_6\mu^2 \cdot g(y) \geq c_4c_6\mu^3 \cdot g(x)$ by Lemmas 8.2(ii) and 8.1. In the second case, since the angle between these facets at h is at least ϕ, $d(E,E') \geq 2\sin(\phi/2) \cdot \mathrm{radius}(S_x \cap S_y)$. By Lemma 8.2(i), $\mathrm{radius}(S_x \cap S_y) \geq c_5\mu^2 \cdot g(x)$. Therefore, $d(E,E') \geq 2c_5\mu^2 \sin(\phi/2) \cdot g(x)$.

Suppose that e is a type 1 edge. Let F be the facet of $\mathcal{P}$ that contains e. F is incident to x. If F is not incident to h, then by Lemma 8.2(iii), $d(S_y, F) \geq (1-3\mu) \cdot g(y)$. This is a lower bound for $d(E,E')$ too. By

Lemma 8.1, we get $d(E,E') \geq c_4\mu(1-3\mu)\cdot g(x)$. Suppose that F is incident to h. Then exactly one endpoint of e lies on $S_x \cap S_y$. There are two possibilities. First, e' lies on the hole on $S_y \cap \mathcal{B}$ opposite $S_x \cap S_y$. Second, e is separated from e' by a facet of $\mathcal{P}$ incident to h. Thus we can show that $d(E,E') = \Omega(g(x))$ as before.

Case 2: $S_x = S_y$. If e or e' is a type 2 edge, it also lies on S_z for some protecting sphere S_z adjacent to S_x. Thus we can reduce this case to case 1. Assume that both e and e' are type 1 edges. So both are great circular arcs of S_x. Let F and F' be the facets of $\mathcal{P}$ that contain e and e', respectively.

We claim that p or q is a curved edge endpoint. Assume to the contrary that $p \in \mathrm{int}(e)$ and $q \in \mathrm{int}(e')$. Let H be the plane through p, q, and x. By the minimality of $\|p-q\|$, pq intersects both e and e' at right angle. So H intersects e and e' at right angles. But then if we translate H slightly away from x, H would intersect e and e' at two points closer than p and q, a contradiction. By our claim, p or q is a curved edge endpoint, say p. So p lies on $S_x \cap S_a$ for some protecting sphere S_a adjacent to S_x. By construction, xa is a linear edge of $\mathcal{Q}$ lying on the boundary of F. Figure 12 shows the three possibilities.

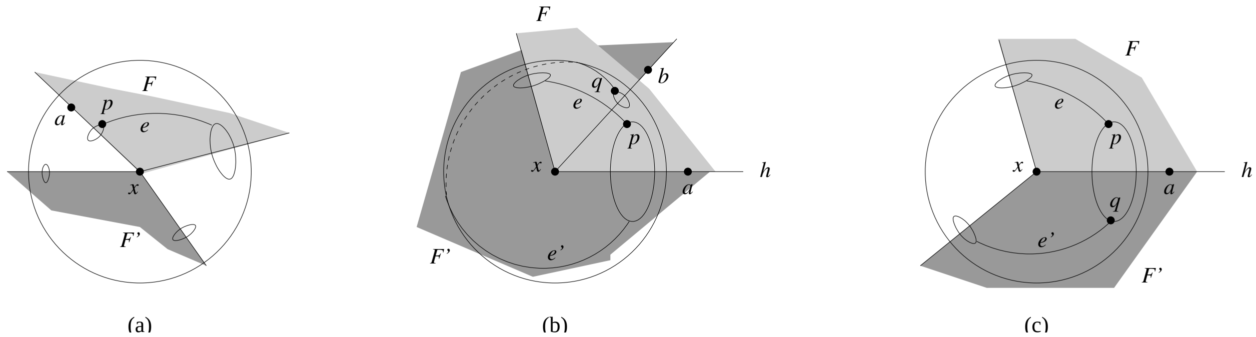

Figure 12:

If a does not lie on the boundary of F' (Figure 12(a)), then by Lemma 8.2(iii), $d(S_a, F') \geq (1-3\mu)\cdot g(a)$. This is a lower bound for $d(E,E')$ too. By Lemma 8.1, $d(E,E') \geq c_4\mu(1-3\mu)\cdot g(x)$. The remaining case is that a lies on the boundary of F'. Thus xa lies on an edge h of $\mathcal{P}$ shared by F and F'. h goes through the hole $S_x \cap S_a$. Take any point $r \in e'$. Observe that as r moves from one endpoint of e' to the other endpoint, $\angle pxr$ increases and then decreases monotonically. Since $\|p-q\| = 2\sin\frac{\angle pxq}{2}\cdot \mathrm{radius}(S_x)$, we conclude that $\|p-q\|$ is minimized when q is an endpoint of e'. So q lies on $S_x \cap S_b$ for some protecting sphere S_b adjacent to S_x. If $S_b \neq S_a$ (Figure 12(b)), then by Lemma 8.2(iv), $d(\mathcal{B}\cap S_a, \mathcal{B}\cap S_b) \geq c_7\mu^3\cdot g(a)$. This is a lower bound for $d(E,E')$ too. By Lemma 8.1, we get $d(E,E') \geq c_4c_7\mu^4\cdot g(x)$. If $S_b = S_a$ (Figure 12(c)), then since the angle between F and F' at h is at least ϕ, $d(E,E') \geq 2\sin(\phi/2)\cdot \mathrm{radius}(S_x \cap S_a)$. By Lemma 8.2(i), $\mathrm{radius}(S_x \cap S_a) \geq c_5\mu^2\cdot g(x)$. Hence $d(E,E') \geq 2c_5\mu^2\sin(\phi/2)\cdot g(x)$.

$\square$

LEMMA 8.8 *Let E and E' be two disjoint curved element and flat facet of $\mathcal{Q}$, respectively. Let S_x be a protecting sphere that contains E. Let F' be the facet of $\mathcal{P}$ that contains E'. Assume that x is a vertex of $\mathcal{P}$, and F' is incident to x. Then $d(E,E') = \Omega(g(x))$.*

Proof. Let $p \in E$ and $q \in E'$ be points such that $\|p-q\| = d(E,E')$. Let p' be the orthogonal projection of p onto the support plane of F'.

First, we claim that we can assume that E' lies inside $\mathcal{B}$. F' contains two flat facets. One is E' and we denote the other by E''. If E' lies inside $\mathcal{B}$, we are done. Suppose not. Since E' and E'' meet $\mathcal{B}\cap S_x$ at the same curved edge, E and E'' are also disjoint. If $p' \in E''$, then $d(E, E'') \leq d(E, E')$. So it suffices to lower bound $d(E, E'')$. Suppose that $p' \notin E''$. Then $p' \notin F'$. So $p'q$ must intersect some boundary edge h of F' at a point q'. Observe that $\|p - q\| \geq \|p - q'\|$. If h is not incident to x, by Lemma 8.2(iii), $\|p - q'\| \geq d(S_x, h) \geq (1 - 3\mu) \cdot g(x)$. So $\|p - q\| \geq \|p - q'\| \geq (1 - 3\mu) \cdot g(x)$ and the lemma is proved. Thus we can assume that h is incident to x. In this case, h bounds E''. Thus $d(E, E'') \leq d(E, E')$, and it suffices to lower bound $d(E, E'')$.

Second, given the first claim, we show that q lies strictly inside S_x. If $p' \in E'$, then $q = p'$ and it lies strictly inside S_x. Otherwise, $p' \notin F'$ and so $p'q$ intersects an edge h of F' bounding E'. As in the proof of the first claim, we can assume that h is incident to x; otherwise the lemma holds already. Then as $\|p - q\|$ is minimized, we have $q = x$ or pq is perpendicular to h. In both cases, q lies strictly inside S_x.

Based on the two claims, there are three cases to consider depending on the locations of q.

Case 1: $q = x$. Then $\|p - q\| = \mathrm{radius}(S_x) = \mu \cdot g(x)$.

Case 2: $q \neq x$ and q lies on a boundary linear edge e' of E'. Since q lies strictly inside S_x, e' is incident to x. Observe that pq is perpendicular to e'. So $\|p - q\|$ is at least the radius of the hole on $\mathcal{B} \cap S_x$ that e' goes through. By Lemma 8.2(i), $\|p - q\| \geq c_5\mu^2 \cdot g(x)$.

Case 3: $q \in \mathrm{int}(E')$. We claim that $p \notin \mathrm{int}(E)$ if E is a curved facet. Otherwise, pq is normal to E and hence S_x. But this implies that $q = x$, a contradiction. By our claim, p lies on a curved edge $e \subseteq E$ ($e = E$ if E is a curved edge). There are two cases to consider.

Case 3.1: e is a type 2 edge. Then e lies on $S_x \cap S_y$ for some protecting sphere S_y adjacent to S_x. If y does not lie on the boundary of F', then by Lemma 8.2(iii), $d(S_y, F') \geq (1 - 3\mu) \cdot g(y)$. This is a lower bound for $d(E, E')$ too. By Lemma 8.1, $d(E, E') \geq c_4\mu(1 - 3\mu) \cdot g(x)$. Suppose that y lies on the boundary of F'. Then xy lies on a boundary edge h of F', and h goes through the hole $S_x \cap S_y$. Since $q \in \mathrm{int}(E')$, pq is a perpendicular from e to E'. Since e and E' are disjoint, pq must intersect another facet G of $\mathcal{P}$ incident to h. Since the angle between F' and G at h is at least ϕ, $d(E, E') = \|p - q\| \geq \sin\phi \cdot \mathrm{radius}(S_x \cap S_y)$, which is at least $c_5\mu^2 \sin\phi \cdot g(x)$ by Lemma 8.2(i).

Case 3.2: e is a type 1 edge. Then e lies on some facet F incident to x. Note that e is a great circular arc of S_x. Since $q \in \mathrm{int}(E')$, pq is normal to E'. Then p must be an endpoint of e; otherwise we can translate pq slightly to decrease the length of pq, a contradiction. So p lies on $S_x \cap S_y$ for some protecting sphere S_y adjacent to S_x. Figure 13 shows the two possibilities.

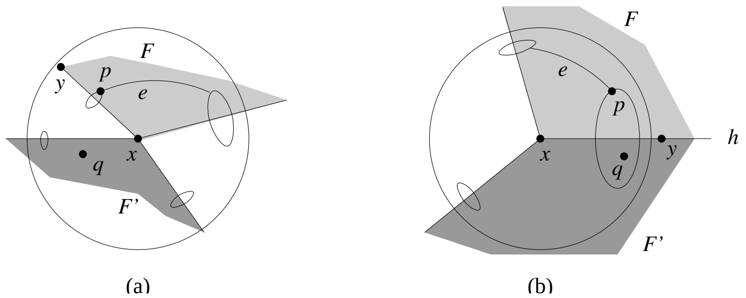

Figure 13:

If y does not lie on the boundary of F' (Figure 13(a)), then by Lemma 8.2(iii), $d(S_y, F') \geq (1 - 3\mu) \cdot g(y)$. This is a lower bound for $d(E, E')$ too. By Lemma 8.1, we get $d(E, E') \geq c_4\mu(1 - 3\mu) \cdot g(x)$. The remaining case is that y lies on the boundary of F' (Figure 13(b)). Then xy lies on an edge h of $\mathcal{P}$ shared by F and F'. h goes through the hole $S_x \cap S_y$. Since the angle between F and F' at h is at least ϕ,

$d(E, E') = \|p - q\| \geq \sin\phi \cdot \mathrm{radius}(S_x \cap S_y)$. By Lemma 8.2(i), $\mathrm{radius}(S_x \cap S_y) \geq c_5\mu^2 \cdot g(x)$. Hence $d(E, E') \geq c_5\mu^2 \sin\phi \cdot g(x)$.

□

Lower bound. We are ready to combine the previous results to prove that $\widehat{f}(p) = \Omega(f(p))$ whenever p lies on or outside $\mathcal{B}$.

LEMMA 8.9 *For each point p on or outside $\mathcal{B}$, $\widehat{f}(p) = \Omega(f(p))$.*

Proof. Let B be the ball centered at p with radius $\widehat{f}(p)$. Let E and E' be two disjoint elements of $\mathcal{Q}$ intersected by B. We conduct a case analysis. In the first two cases, we directly show that $\mathrm{radius}(B) = \Omega(f(p))$. In the other cases, we show that B intersects a protecting sphere S_x such that $\mathrm{radius}(B) = \Omega(g(x))$. Then Lemma 8.6 implies $\mathrm{radius}(B) = \Omega(f(p))$. Let $d(E, E')$ denote the minimum distance between E and E'.

Case 1: E or E' is a linear edge, say E. $\mathrm{radius}(B)$ is at least the distance from p to E, which is $\Omega(f(p))$ by Lemma 8.4.

Case 2: E and E' are flat facets. Since p lies on or outside $\mathcal{B}$, p lies on at most one of the facets of $\mathcal{P}$ containing E and E'. Thus $\mathrm{radius}(B) \geq g(p)$, which is $\Omega(f(p))$ by Lemma 8.5.

Case 3: E and E' are curved elements. Let S_x and S_y be two protecting spheres containing E and E', respectively.

Case 3.1: S_x and S_y are neither identical nor adjacent. By Lemma 8.2(iv), $d(E, E') \geq c_7\mu^3 \cdot g(x)$. Thus $\mathrm{radius}(B) \geq \frac{c_7\mu^3}{2} \cdot g(x)$.

Case 3.2: Both $\mathcal{B} \cap S_x$ and $\mathcal{B} \cap S_y$ are rings (adjacent or identical). Then x and y lie in the interior of some edge h of $\mathcal{P}$. If E and E' do not intersect the same hole, then one of them is a curved edge lying on a hole that the other does not intersect. In this case, we can change our choices of S_x and S_y to invoke case 3.1. Suppose that E and E' intersect the same hole, say a hole on $\mathcal{B} \cap S_x$. Since E and E' are disjoint, they are separated by two facets of $\mathcal{P}$ incident to h. The angle between these two facets at h is at least ϕ. So $d(E, E')$ is at least $2\sin(\phi/2)$ times the radius of the smallest hole on $\mathcal{B} \cap S_x$. Thus, by Lemma 8.2(i), $d(E, E') \geq 2c_5\mu^2 \sin(\phi/2) \cdot g(x)$. It follows that $\mathrm{radius}(B) \geq c_5\mu^2 \sin(\phi/2) \cdot g(x)$.

Case 3.3: $\mathcal{B} \cap S_x$ or $\mathcal{B} \cap S_y$ is not a ring, and S_x and S_y are identical or adjacent. Assume that $\mathcal{B} \cap S_x$ is not a ring. So x is a vertex of $\mathcal{P}$. By Lemma 8.7, $d(E, E') = \Omega(g(x))$. Hence $\mathrm{radius}(B) \geq d(E, E')/2 = \Omega(g(x))$.

Case 4: E is a curved element and E' is a flat facet. Let S_x be a protecting sphere that contains E. Let F' be the facet of $\mathcal{P}$ that contains E'. If $x \notin \partial F'$, Lemma 8.2(iii) implies that $d(E, E') \geq (1 - 3\mu) \cdot g(x)$. Thus $\mathrm{radius}(B) \geq \frac{1-3\mu}{2} \cdot g(x)$. Consider the case in which $x \in \partial F'$.

Case 4.1: x is not a vertex of $\mathcal{P}$. So $\mathcal{B} \cap S_x$ is a ring and x lies in the interior of an edge h of $\mathcal{P}$. Since E and E' are disjoint, E and F' are separated by a facet G of $\mathcal{P}$ incident to h. Since the angle between F' and G at h is at least ϕ, $d(E, E')$ is at least $\sin\phi$ times the radius of the smallest hole on $\mathcal{B} \cap S_x$. By Lemma 8.2(i), $d(E, E') \geq c_5\mu^2 \sin\phi \cdot g(x)$. Thus $\mathrm{radius}(B) \geq \frac{c_5\mu^2 \sin\phi}{2} \cdot g(x)$.

Case 4.2: x is a vertex of $\mathcal{P}$. By Lemma 8.8, $d(E, E') = \Omega(g(x))$. Thus $\mathrm{radius}(B) \geq d(E, E')/2 = \Omega(g(x))$.

□

8.4 Quality guarantees

THEOREM 8.1 *Let $\rho_0 > 16$ and $\mu \in (0, 1/7]$ be two constants. Given a bounded domain represented by a piecewise linear complex with smallest angle ϕ,* MESH *terminates and produces a conforming Delaunay mesh $\mathcal{M}$ in $O(N^2)$ time, where N is the number of vertices in $\mathcal{M}$.*

(i) For each vertex v of $\mathcal{M}$, the shortest incident edge of v has length $\Omega(f(v))$, where the constant depends on μ and ϕ.

(ii) Let τ be a tetrahedron in $\mathcal{M}$. If τ does not lie inside $\mathcal{B}$, then $\rho(\tau) \leq \rho_0$. Otherwise, $\rho(\tau)$ is bounded by a constant depending on μ and ϕ.

Proof. The termination of MESH has been proved in Lemma 7.7. Since MESH terminates, Lemma 6.3 implies that $\mathcal{M}$ is Delaunay and conforms to $\mathcal{P}$. The running time follows from Lemma 4.1.

Consider (i). Let e be an incident edge of v. Suppose that v is a linear edge endpoint. Either e is a linear edge or $\mathrm{length}(e) = \mathrm{radius}(S_v)$. In the first case, $\mathrm{length}(e) \geq c_3\mu \cdot g(v)$ by Lemma 8.1. In the second case, $\mathrm{length}(e) \geq c_3\mu \cdot g(v)$ by (7). In both cases, Lemma 8.5 implies that $\mathrm{length}(e) = \Omega(f(v))$. The remaining case is that v lies on or outside $\mathcal{B}$. Then Lemmas 7.7 and 8.9 imply that $\mathrm{length}(e) = \Omega(f(v))$.

Consider (ii). If τ does not lie inside $\mathcal{B}$, rule 4 guarantees that $\rho(\tau) \leq \rho_0$. Otherwise, there are two cases.

Case 1: There exists a protecting sphere S_x such that $\tau = pqrx$ for some helper triangle pqr on CH_x. Let e be the shortest edge of τ. Note that e is incident to one of p, q, and r, say p. By the Lipschitz condition and (7), we get $f(p) \geq f(x) - \mathrm{radius}(S_x) \geq f(x) - 3\mu \cdot g(x)$. As $g(x) \leq f(x)$, we have $f(p) \geq (1 - 3\mu) \cdot f(x)$. Then by (i), $\mathrm{length}(e) = \Omega(f(p)) = \Omega(f(x))$.

Since the angular diameter of the cap K_{pqr} is at most $\pi/3$, the circumradius of τ is less than $\mathrm{radius}(S_x)$, which is at most $3\mu \cdot g(x)$ by (7). As $g(x) \leq f(x)$, we conclude that $\rho(\tau) = O(1)$.

Case 2: There exists adjacent protecting spheres S_x and S_y such that $\tau = pqxy$ for some helper arc $\widehat{pq}$ on $S_x \cap S_y$. Assume that S_x is not smaller than S_y. Let e be the shortest edge of τ. Since S_x and S_y intersect at right angle, $e \neq xy$. So e is incident to p or q. We can show that $\mathrm{length}(e) = \Omega(f(x))$ as in case 1. The circumradius of τ is less than $\mathrm{radius}(S_x)$, which is at most $3\mu \cdot g(x)$ by (7). As $g(x) \leq f(x)$, we conclude that $\rho(\tau) = O(1)$.

$\square$

9 Discussion

Our approach is based on protecting the vertices and edges of the input domain with an explicit buffer zone. The buffer zone disallows the insertions of vertices near the sharp input angles. Thus it prevents the indefinite splitting of mesh elements that may happen when input elements meet at a sharp angle. This approach caters for the worst-case. It does not follow the same adaptive and on-demand strategy of conventional Delaunay refinement. A purely adaptive and on-demand method would be a better choice for implementation. There is a recent advance in this direction. By restricting the input to polyhedra, Cheng et al. [5] limits the explicit construction of buffer zones to protect input vertices only. The output mesh is graded, and the radius-edge ratio of almost all tetrahedra are bounded by a constant independent of the input domain. The skinny ones left are provably close to input vertices or edges where the input angles are acute. Although there is no radius-edge ratio bound in terms of the input angle for these skinny tetrahedra, the experimental results in [5] show that only a few skinny tetrahedra remain. In the presence of small angles, it remains an open problem how to construct a conforming Delaunay mesh with bounded aspect ratio. Can the weighted Delaunay refinement [3] or the randomized refinement [12] techniques be applied?

References

[1] M. Bern, D. Eppstein and J. Gilbert. Provably good mesh generation. *J. Comput. Syst. Sci.*, 48 (1994), 384–409.

[2] B. Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete & Computational Geometry*, 9 (1993), 145–158.

[3] S.-W. Cheng and T.K. Dey. Quality meshing with weighted Delaunay refinement. *SIAM J. Comput.*, 33 (2003), 69–93.

[4] S.-W. Cheng, T. K. Dey, H. Edelsbrunner, M. A. Facello and S.-H. Teng. Sliver exudation. *J. ACM*, 47 (2000), 883–904.

[5] S.-W. Cheng, T.K. Dey, E. Ramos, and T. Ray. Quality Meshing for Polyhedra with Small Angles. To appear in *Proc. 20th ACM Sympos. Comput. Geom.*, 2004.

[6] L. P. Chew. Guaranteed-quality Delaunay meshing in 3D. *Proc. 13th ACM Sympos. Comput. Geom.*, 1997, 391–393.

[7] D. Cohen-Steiner, E.C. de Verdiere and M. Yvinec. Conforming Delaunay Triangulations in 3D. *Proc. 18th ACM Sympos. Comput. Geom.*, 2002, 199–208.

[8] T. K. Dey and J. Pach. Extremal problems for geometric hypergraphs. *Discrete Comput. Geom.*, **19**, (1998), 473–484.

[9] H. Edelsbrunner. *Algorithms in Combinatorial Geometry.* Springer-Verlag, Berlin, 1987.

[10] H. Edelsbrunner. *Geometry and topology for mesh generation*. Cambridge University Press, 2001.

[11] P.L. George and H. Borouchaki. *Delaunay triangulation and meshing : application to finite elements*. Hermes, 1998.

[12] X.-Y. Li and S.-H. Teng. Generating well-shaped Delaunay meshes in 3D. *Proc. 12th ACM-SIAM Sympos. Discrete Alg.*, 2001, 28–37.

[13] G.L. Miller, D. Talmor, S.-H. Teng, and N. Walkington. On the radius-edge condition in the control volume method. *SIAM J. Numer. Analy.*, 36 (1999), 1690–1708.

[14] S. A. Mitchell and S. A. Vavasis. Quality mesh generation in higher dimensions. *SIAM J. Comput.*, 29 (2000), 1334–1370.

[15] M. Murphy, D. M. Mount, C. W. Gable. A point-placement strategy for conforming Delaunay tetrahedralization, *Int'l. J. Comput. Geom. Appl.*, 11 (2001), 669–682.

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

[17] J. R. Shewchuk. Tetrahedral mesh generation by Delaunay refinement. *Proc. 14th ACM Sympos. Comput. Geom.*, 1998, 86–95.

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