Graded Conforming Delaunay Tetrahedralization with Bounded Radius-Edge Ratio*

Siu-Wing Cheng† Sheung-Hung Poon†

July 9, 2002

Abstract

We propose an algorithm to compute a conforming Delaunay mesh of a polyhedral domain in three dimensions. Arbitrarily small input angles are allowed. The output mesh is graded and has bounded radius-edge ratio everywhere.

1 Introduction

In finite element analysis, a domain needs to be partitioned into a cell complex for the purpose of numerical simulation and analysis [7]. A simplicial complex is a popular choice and it is also commonly known as a tetrahedral mesh. 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 faces of tetrahedra in the mesh. An important challenge in mesh generation is to construct a mesh with good quality. Our contribution is a simple Delaunay refinement algorithm that produces tetrahedra with provably good edge lengths and radius-edge ratio. Our algorithm is distinguished from previous ones [2, 5, 8, 11, 13, 14] by its ability to handle input angles less than $\pi/2$ and its theoretical guarantees.

Our input domain is a bounded volume in 3D whose boundary is specified by a *piecewise-linear complex* $\mathcal{P}$. 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. We make the simplifying assumption that each edge of $\mathcal{P}$ has two or more incident facets, e.g., polyhedron possibly with voids and holes. This assumption is not critical and it can be removed, with more work, without affecting our results.

Delaunay tetrahedralization is a popular tetrahedral mesh in theory and practice [6, 7]. For results using quadtree and octtree based methods, please refer to the papers by Bern et al. [1] and Mitchell and Vavasis [10]. Ruppert [12] proposed the Delaunay refinement algorithm to mesh a 2D polygonal domain. The mesh is *graded*, i.e., the shortest edge incident to every vertex v has length at least a constant factor of the local feature size at v. Every triangle has bounded aspect ratio. The size of the mesh is asymptotically optimal. Shewchuk [13] extended Delaunay refinement to 3D for polyhedral domains. A graded conforming Delaunay mesh is obtained but there are two differences. First, when some input angle is less than $\pi/2$, the algorithm may or may not terminate depending on the specific input instance. Second, for each tetrahedron τ, its *radius-edge ratio* (i.e., the ratio of the circumradius of τ to the shortest edge length of τ) is bounded by a constant. 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*. Nevertheless, bounded radius-edge ratio works well in some applications [9].

Recently, methods have been discovered to eliminate slivers when every input angle is at least $\pi/2$. Li and Teng [8] improved Delaunay refinement with a random point-placement strategy in line of Chew [4]. Cheng et al. [3] introduced *sliver exudation* to eliminate slivers from a Delaunay mesh of a periodic point set with bounded radius-edge ratio. Cheng and Dey [2] introduced *weighted Delaunay refinement* which extends sliver exudation to handle boundaries. Both algorithms by Li and Teng [8] and Cheng and Dey [2] produce a graded conforming Delaunay mesh with bounded aspect ratio and asymptotically optimal size.

Much less is known about handling polyhedral do-

*This work has been 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

mains with input angles less than $\pi/2$. Murphy et al. [11] showed the existence of a conforming Delaunay mesh, but their method produces tetrahedra of poor shape and unnecessarily many vertices. Cohen-Steiner et al. [5] proposed an improved method and they experimentally studied the effectiveness of their algorithm. Shewchuk [14] attacked the problem differently and generated a *constrained Delaunay tetrahedralization*. In the above results, gradedness is not guaranteed and the radius-edge ratio is not guaranteed to be bounded everywhere. It is sometimes unavoidable that the edge length and the shape of tetrahedra deteriorate near a small input angle. Thus, it is conceivable that there are lower bound on edge length and upper bound on radius-edge ratio that use constant factors depending on the input angle. Nevertheless, no such result is known till now.

For the purposes of this paper, we measure three types of angles as follows. 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 the supporting plane of F. The *angle between uv and F* is $\min\{\angle puv : p \in L \cap \mathrm{int}(F)\}$. 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 $\min_{u \in H_1 \cap H_2}\{\angle puq : p \in L_u \cap \mathrm{int}(F_1), q \in L_u \cap \mathrm{int}(F_2)\}$. Throughout this paper, ϕ denotes the smallest angle in the domain measured as described above. We assume that $\phi < \pi/2$ as the other case has been solved [2, 8].

We present an algorithm MESH that constructs a conforming unweighted Delaunay tetrahedralization given a polyhedral domain. The mesh is graded and has bounded radius-edge ratio everywhere (Theorem 1 in Section 10). Let $\mu \in (0, 1/7]$ and $\rho_0 > 16$ be two a priori chosen constants. Our algorithm encloses the input edges within a buffer zone whose size is proportional to local feature size. For every tetrahedron τ, if τ does not lie inside the buffer zone, its radius-edge ratio $\rho(\tau) \leq \rho_0$; otherwise, $\rho(\tau) \leq \rho_1$ where ρ_1 depends on μ and ϕ. The shortest edge incident to a vertex v has length at least a factor δ of the local feature size at v where δ depends on μ and ϕ.

The rest of the paper is organized as follows. Section 2 gives some basic definitions and an overview of our algorithm. Section 3 describes the augmentation of the input complex with the buffer zone before MESH processes it. Section 4 describes MESH. Sections 5–7 prove that the output mesh is conforming. Sections 8–10 prove the bounds on edge length and radius-edge ratio. In Section 11, we discuss some future work.

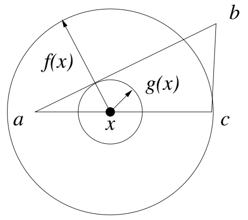

Figure 1: The large and small circles have radii $f(x)$ and $g(x)$ respectively.

2 Preliminaries and overview

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. For a point x, the *local gap size* $g(x)$ 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 points that we are interested in (Lemmas 16 and 17 in Section 10). In general, local gap sizes do not satisfy the Lipschitz property. However, the Lipschitz property holds under certain conditions and this sufficient for our purposes.

LEMMA 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 $g(y) + ||x - y||$. So B intersects two elements of $\mathcal{P}$, one of which does not contain y. Denote this element by E. Since $y \in e$ and $x \in \mathrm{int}(e)$, E does not contain x. So $\mathrm{radius}(B) \geq g(x)$. ◻

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 $S \cap S'$.

We enclose the edges of $\mathcal{P}$ with a buffer zone. We compute spheres centered at points on edges of $\mathcal{P}$. The buffer zone boundary is the outer boundary of the union of these spheres. $\mathcal{P}$ is then augmented with the buffer zone boundary to yield a new complex $\mathcal{Q}$. The idea is to apply Delaunay refinement to $\mathcal{Q}$ to mesh the space outside the buffer zone such that the tetrahedralization of the space inside the buffer zone is automatically induced. The spheres are judiciously chosen so that consecutive ones are orthogonal. The intuition is that the space outside the buffer zone will have non-acute angle, thus allowing the use of Delaunay refinement. There are still two difficulties to overcome. First, we need to guarantee that unnecessarily short edges are not forced when constructing the buffer zone. Second, we need a method to triangulate the spherical buffer zone boundary.

3 Augmenting $\mathcal{P}$

We describe the buffer zone and its merging with $\mathcal{P}$ to yield $\mathcal{Q}$. Several properties of the buffer zone and $\mathcal{Q}$ are described in Lemmas 2–5. It suffices to know the construction of the buffer zone and $\mathcal{Q}$, Lemma 2 and Lemma 4 to understand MESH (Section 4), prove boundary conformity (Sections 5–7) and prove termination of MESH (Sections 8 and 9). Lemma 3 is used with Lemma 2 to prove Lemma 5 which is then used in Section 10 to analyze the edge lengths and radius-edge ratio.

3.1 Protecting spheres

Let μ be some fixed constant chosen from $(0, \frac{1}{7}]$. For each edge e of $\mathcal{P}$, we create some spheres with centers lying on e. We call these *protecting spheres*. First, for each 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:

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

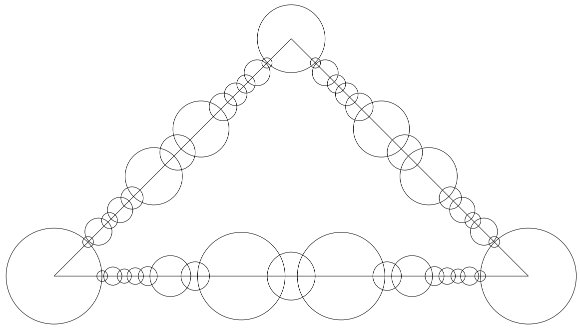

Figure 2: $\mu = 1/7$ and the base angle is $\pi/4$.

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

By construction, S_u and S_{u_v} are orthogonal and so are S_v and S_{v_u}. Third, 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 *consecutive* 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 consecutive protecting spheres are orthogonal.

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 2 shows the protecting spheres created for the sides of an isosceles triangle. The following lemma states that each protecting sphere S_x obtained has radius $\Theta(\mu \cdot g(x))$, the distance between two neighboring centers is lower bounded by their local gap sizes and the local gap sizes of two neighboring centers do not differ much. The proof of Lemma 2 can be found in Appendix 12.1.

LEMMA 2 *Let* $c_1 = 2\pi/(3\sqrt{3}) > 1$ *and* $c_2 = \min\{\sqrt{3}/2, \sin\phi\} < 1$. *There exist constants* $c_3 < c_2$ *and* $c_4 < 1$ *such that for each edge* uv *of* $\mathcal{P}$, *the following hold.*

(i) S_{u_v} *and* S_{v_u} *are orthogonal to* S_u *and* S_v *respectively. The two ratio* $\frac{\mathrm{radius}(S_{u_v})}{g(u_v)}$ *and* $\frac{\mathrm{radius}(S_{v_u})}{g(v_u)}$ *lie in* $[c_2\mu, c_1\mu]$.

(ii) $Split(u_v, v_u)$ *terminates and returns a sequence* $\mathcal{S}$ *of protecting spheres covering* $u_v v_u$. *Any two consecutive protecting spheres in* $\mathcal{S}$ *are orthogonal. For any* $S_z \in \mathcal{S} - \{S_{u_v}, S_{v_u}\}$, *the ratio* $\frac{\mathrm{radius}(S_z)}{g(z)}$ *lies in* $[c_3\mu, 3\mu]$.

(iii) Let x *and* y *be two neighboring centers of protecting spheres on* uv. *Then* $\|x - y\| > c_3\mu \cdot \max\{g(x), g(y)\}$ *and* $g(y) \geq c_4\mu \cdot g(x)$.

3.2 Buffer zone

Given a set $\mathcal{S}$ of spheres, we use $\mathrm{Bd}(\bigcup_{S\in\mathcal{S}} S)$ to denote the outer boundary of $\bigcup_{S\in\mathcal{S}} S$. Let $\mathcal{B} = \mathrm{Bd}(\bigcup S_x)$, 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. $\mathcal{B} \cap \bigcup_{S_x\in\mathcal{S}_{uv}} S_x$ 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 consecutive 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 $\mathcal{B} \cap S_z$. For each ring, we define its *width* as the distance between the parallel planes containing the two holes. Lemma 3 states that the width of each ring is lowered bounded by the local gap size and so is the radius of each hole on $\mathcal{B} \cap S_x$ for any protecting sphere S_x. Moreover, $\mathcal{B}$ encloses the edges of $\mathcal{P}$ without causing any unwanted self-intersection or intersection with $\mathcal{P}$. The proof of Lemma 3 can be found in Appendix 12.2.

LEMMA 3 *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 consecutive to* S_x. *The minimum distance between* $\mathcal{B} \cap S_x$ *and* $\mathcal{B} \cap S_y$ *is at least* $c_7\mu^3 \cdot g(x)$.

3.3 The new complex $\mathcal{Q}$

We merge $\mathcal{B}$ with $\mathcal{P}$ to produce a new complex $\mathcal{Q}$. $\mathcal{B}$ 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 $x \in uv$ is divided by the facets of $\mathcal{P}$ incident to uv into curved rectangular patches; and for each vertex v of $\mathcal{P}$, $\mathcal{B} \cap S_v$ is divided by the facets of $\mathcal{P}$ incident to v into spherical patches. These 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 facets are the *curved edges* of $\mathcal{Q}$. The vertices of $\mathcal{Q}$ consists of the endpoints of linear and curved edges.

For any protecting sphere S_x and any curved facet E on $\mathcal{B} \cap S_x$, ∂E consists of curved edges that lie at the intersections between S_x and either facets of $\mathcal{P}$ or protecting spheres consecutive to S_x. Moreover, these two kinds of curved edges alternate in ∂E. How many edges can a facet F of $\mathcal{P}$, where $x \in \partial F$, contribute to ∂E? If x is not a vertex of $\mathcal{P}$, the answer is clearly at most one as E is rectangular. Suppose that x is a vertex of $\mathcal{P}$. Observe that x appears 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 at most one edge to ∂E. However, a hole on $\mathcal{B} \cap S_x$ may contribute several edges to ∂E when x is a vertex of $\mathcal{P}$.

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

LEMMA 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 adjacent to* F. *If* e *and* F *do not lie on the same*

plane or sphere, then the normal to F at $e \cap F$ is tangent to and coplanar with e.

Lemma 4 motivates the use of Delaunay refinement in the space outside $\mathcal{B}$. In essence, we compute a mesh that approximates $\mathcal{Q}$ and respects the input boundary. Due to Delaunay refinement (modified to handle curved elements), the edge lengths in the final mesh will be proportional to the local feature sizes with respect to $\mathcal{Q}$. For each point p, let $\widehat{f}(p)$ denote the local feature size at p with respect to $\mathcal{Q}$.[1] Lemma 5 states that if p lies on or outside $\mathcal{B}$, $\widehat{f}(p) = \Omega(g(p))$. This will allow us to relate the edge lengths in the final mesh to the local feature sizes with respect to $\mathcal{P}$ in Section 10. The proof of Lemma 5 can be found in Appendix 12.3.

LEMMA 5 *For any point p on or outside $\mathcal{B}$, $\widehat{f}(p) \geq \lambda\mu^8 \cdot g(p)$ for some constant $\lambda < 1$.*

4 Algorithm MESH

We introduce some 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 with endpoints p and q.

MESH approximates $\mathcal{Q}$ by a Delaunay subcomplex. We initialize a set $\mathcal{V}$ as the set of vertices of $\mathcal{Q}$. The initial complex is the Delaunay tetrahedralization, $\mathrm{Del}\,\mathcal{V}$, of $\mathcal{V}$. $\mathcal{V}$ induces several types of geometric objects that guide MESH to refine the mesh by inserting vertices into $\mathcal{V}$. We first define these objects.

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. Let $\widehat{pq}$ be a helper arc on e. The *circumcap* K of $\widehat{pq}$ is the smallest cap on S that contains $\widehat{pq}$. If the angular width of $\widehat{pq}$ is less than π, the *normal sphere* of $\widehat{pq}$ is the orthogonal sphere of S at ∂K and $\widehat{pq}$ is *encroached* by a point v if v lies inside its normal sphere. If the angular width of $\widehat{pq}$ is larger than $\pi/3$, $\widehat{pq}$ is *wide*.

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. If a convex polygon P with more than three vertices appears as a boundary facet of CH_x, then we triangulate P as follows. Let L be the supporting plane of P. The circumcap of P is the cap on S_x that is bounded by $L \cap S_x$ and separated from CH_x by L. First, for each helper arc $\widehat{pq}$ such that $p, q \in \partial P$ and $\widehat{pq}$ lies on the circumcap of P, we insert pq as a diagonal in P. Then we arbitrarily complete the triangulation of P. Afterwards, 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 3. Let H be the plane containing a helper triangle t. The *circumcap* of t is the cap K on S_x that is bounded by $H \cap S_x$ and separated from CH_x by H. If the angular diameter of K is less than π, the *normal sphere* of t is the orthogonal sphere of S_x at ∂K and t is *encroached* by a point v if v lies inside its normal sphere. If the angular diameter of K is larger than $\pi/3$, t is *wide*.

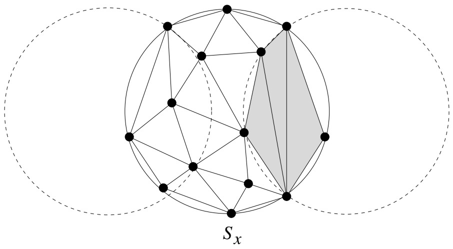

Figure 3: The figure shows S_x and two protecting spheres consecutive to S_x. Some boundary triangles of CH_x are shown. The non-shaded triangles are helper triangles. The shaded ones are not as the vertices of each shaded triangle lie on the boundary of the same hole on $\mathcal{B} \cap S_x$.

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 2D Delaunay triangulation of $\mathcal{V} \cap F$. Note that 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 curved boundary edges of flat facets. The *circumcap* of a subfacet τ is the disk bounded by the circumcircle of τ. The *normal sphere* of τ is the equatorial sphere of τ. If a point v lies inside the normal sphere of τ, τ is *encroached* by v.

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. When no rule is applicable, the subcomplex of $\mathrm{Del}\,\mathcal{V}$ covering the input domain is the final mesh. Recall that $\rho_0 > 16$ is an a priori chosen constant.

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

Rule 2: Pick a helper triangle t that is wide or encroached by a vertex in $\mathcal{V}$. Preference is given to

[1] $\widehat{f}(p)$ is the radius of the smallest ball centered at p that intersects two disjoint elements of $\mathcal{Q}$.

wide helper triangles. Let v be the center of the circumcap of t. If v does not encroach upon any helper arc, insert v. Otherwise, reject v and apply Rule 1 to split the helper arcs encroached by v.

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

Rule 4: Let v be the circumcenter of a tetrahedron τ such that $\rho(\tau) > \rho_0$ and no vertex of τ lies inside $\mathcal{B}$. If v does not encroach upon any helper arc, helper triangle or subfacet, insert v. Otherwise, reject v, apply Rule 1 to split the helper arcs encroached by v, and then apply Rules 2 and 3 to split the helper triangles and subfacets encroached by v.

This completes the description of MESH. The rest of the paper focuses on proving the guarantees offered by MESH. We will see that MESH never inserts a vertex inside $\mathcal{B}$, i.e., the vertices inside $\mathcal{B}$ are always the endpoints of linear edges of $\mathcal{Q}$.

5 Properties of orthogonality

This section presents three geometric results regarding orthogonal spheres. We introduce some notations. Given a sphere S and a point p outside S, $K(p,S)$ denotes the cap on S visible from p. 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. If S is a plane (infinite sphere), then $K^\perp$ is the equatorial sphere of K. For any point $q \in \partial K(p,S)$, pq is tangent to S, so p is the center of $K(p,S)^\perp$.

CLAIM 1 *Let S be a sphere. Let S' be a sphere such that $S \cap S'$ is an equator of S'. For any point z on the plane containing $S \cap S'$ and outside S', $K(z,S')^\perp$ is orthogonal to S.*

Proof. Let x and y be the centers of S and S' respectively. Recall that z is the center of $K(z,S')^\perp$. Let r be the radius of $K(z,S')^\perp$. Since xyz is a right-angled triangle, we have $||x-z||^2 = ||x-y||^2 + ||y-z||^2$. By Pythagoras theorem again, we have $||x-y||^2 = \mathrm{radius}(S)^2 - \mathrm{radius}(S')^2$. Since S' and $K(z,S')^\perp$ are orthogonal, we have $||y-z||^2 = \mathrm{radius}(S')^2 + r^2$. It follows that $||x-z||^2 = \mathrm{radius}(S)^2 + r^2$ and so $K(z,S')^\perp$ is orthogonal to S. ◻

CLAIM 2 *Let p and q be two non-diametral points on a sphere S centered at x. Let $\mathcal{N}$ be the set of spheres orthogonal to S that pass through p and q. There exists a unique circle C such that*

(i) C is coplanar with pqx, C passes through p and q, and $\mathcal{N}$ is the set of spheres that pass through C.

(ii) The locus of the centers of spheres in $\mathcal{N}$ is the line ℓ through the center of C perpendicular to the plane containing C.

Proof. Let H be the plane containing p, q and x. Take $N \in \mathcal{N}$. The two circles $H \cap S$ and $H \cap N$ are orthogonal in the sense that they intersect at right angle. It can be verified that there is a unique circle C on H that is orthogonal to $H \cap S$ and passes through p and q. Thus, $H \cap N = C$ and $\mathcal{N}$ is the set of spheres that pass through C. This proves (i) and (ii) is an easy corollary of (i). ◻

CLAIM 3 *Let S be a sphere. Let K_1 and K_2 be caps on S with angular diameter less than π. If $K_2 \subseteq K_1$, $K_1^\perp$ encloses $K_2^\perp$.*

Proof. Fix the center of K_2 and grow it to a cap K such that ∂K is tangent to ∂K_1. So $K \subseteq K_1$. Clearly, $K^\perp$ encloses $K_2^\perp$. If we treat the contact point between ∂K and ∂K_1 as a degenerate circle, $K^\perp$ and $K_1^\perp$ belong to the system of orthogonal spheres as described in Claim 2 in the limiting case. So $K_1^\perp$ encloses $K^\perp$ and hence $K_2^\perp$. ◻

6 Locations of centers

We study the locations of the circumcap centers when MESH inserts them. To this end, we need to associate helper arcs, helper triangles and subfacets with elements of Q. We first introduce some notations. Given a helper arc, helper triangle or subfacet σ, let K_σ denote the circumcap of σ. Hence, if the angular diameter of K_σ is less than π, $K_\sigma^\perp$ is the normal sphere of σ. We extend the definition for any circular arc β. The circumcap K_β is the smallest cap on the equatorial sphere of β that contains β. $K_\beta^\perp$ is defined as before if the angular diameter of K_β is less than π.

A helper arc *belongs to* the curved edge that contains it. A subfacet τ *belongs to* the flat facet F if F lies outside $\mathcal{B}$ and F contains the vertices of τ. Note that it is possible that $\tau \not\subseteq F$. This definition does not cover all subfacets but we are not concerned as we will see in Section 7 that MESH never deals with subfacets that are not covered. Clearly, a subfacet belongs to at most one flat facet. A helper triangle t *belongs to* a curved facet F if there exists a connected subset $\gamma \subseteq \mathrm{int}(K_t) \cap F$ such that $\mathrm{cl}(\gamma)$ contains the vertices of t. This definition is somewhat complicated due to the fact that t does not lie on $\mathcal{B}$ and the vertices of t may lie on the boundaries of two curved facets. The following result shows that t belongs to exactly one curved facet under the right conditions.

LEMMA 6 *Assume that* MESH *has not inserted any vertex inside $\mathcal{B}$. If there is no wide or encroached helper arc, each helper triangle belongs to exactly one curved facet.*

Proof. Let t be a helper triangle on CH_x. Assume to the contrary that t does not belong to any curved facet. Then some helper arc α must cross K_t and cut ∂K_t into two arcs A_1 and A_2 such that each $\mathrm{int}(A_i)$ contains a vertex of t. Note that α does not lie on $S_x \cap S_y$ for any protecting sphere S_y consecutive to S_x. Otherwise, A_1 or A_2 lies inside S_y which implies that a vertex of t lies inside $\mathcal{B}$, contradicting our assumption. It follows that α lies at the intersection of S_x and a facet of $\mathcal{P}$, i.e., S_x is the equatorial sphere of α.

If A_i is less than a semicircle for some i, then K_α contains A_i. But then α is encroached by the vertex of t in $\mathrm{int}(A_i)$, contradicting our assumption that no helper arc is encroached. Suppose that A_1 and A_2 are semicircles. Then α passes through the center of K_t. If any endpoint of α lies outside K_t, then K_α contains K_t and so some vertex of t encroaches upon α, contradiction. Otherwise, both endpoints of α lie on ∂K_t. Thus, the vertices of t and the endpoints of α are vertices of a boundary facet of CH_x. Recall that when triangulating the boundary of CH_x, we first connect the endpoints of α with a diagonal. But then this diagonal cuts t and so t would not exist, contradiction. This completes the proof that t belongs to at least one curved facet.

Lastly, t cannot belong to two curved facets, otherwise the definition would imply that the interior of two curved facets intersect. ◻

Clearly, for a helper arc α, the center of K_α lies on the curved edge that α belongs to. In fact, the center is the midpoint of α. The next two lemmas show that for a subfacet (resp. helper triangle) σ, the center of K_σ lies on the flat facet (resp. curved facet) that σ belongs to. With slightly more work in Section 7, these two lemmas will allow us to show that Rules 2 and 3 never insert a vertex inside $\mathcal{B}$.

LEMMA 7 *Assume that* MESH *has not inserted any vertex inside $\mathcal{B}$ and there is no wide or encroached helper arc. Let τ be a subfacet belonging to a flat facet F. The center of K_τ lies on F.*

Proof. Let v be the center of K_τ. Let H be the plane containing F. If $v \notin F$, K_τ intersects ∂F at an arc β such that β cuts K_τ into two parts, one contains v and the other contains $K_\tau \cap F$. This implies that $K_\tau \cap F$ lies inside $K_\beta^\perp$. Since the vertices of τ lie on $K_\tau \cap F$, some vertex of τ lies inside $K_\beta^\perp$. The emptiness of K_τ implies that β lies within a helper arc α. By Claim 3, $K_\alpha^\perp$ encloses $K_\beta^\perp$, so α is encroached by some vertex of τ, contradiction. ◻

LEMMA 8 *Assume that* MESH *has not inserted any vertex inside $\mathcal{B}$ and there is no wide or encroached helper arc. Let t be a helper triangle belonging to a curved facet F. The center of K_t lies on F.*

Proof. Suppose that $F \subseteq S_x$ for a protecting sphere S_x. Let v be the center of K_t. Assume to the contrary that v lies outside F.

Case 1: v lies outside $\mathcal{B} \cap S_x$. So v lies inside some protecting sphere S_y consecutive to S_x. Note that $v \in K(y, S_x)$. Let p be a vertex of t that does not lie on $S_x \cap S_y$ (such a vertex exists by the definition of helper triangle). Since p lies outside S_y, K_t intersects $S_x \cap S_y$ at an arc β. The emptiness of K_t implies that β lies within a helper arc α. Since α is not wide by assumption, the angular width of β is less than π, so $K_\beta^\perp$ is defined. Since $\beta = K_t \cap S_x \cap S_y$ and $v \in K(y, S_x)$, the angular diameter of K_t is also less than π, otherwise the angular width of β would be at least π. So the normal sphere of t is $K_t^\perp$. Let H be the plane through x and the endpoints of β. By Claim 1, $K_\beta^\perp$ is orthogonal to S_x. Since S_y and $K_t^\perp$ are also orthogonal to S_x, Claim 2 implies that S_y, $K_t^\perp$ and $K_\beta^\perp$ intersect at the circle $H \cap K_\beta^\perp$. It follows that the caps K_β, K_t and $K(y, S_x)$ contain β and their boundaries pass through the endpoints of β. See Figure 4. Using this and the fact that $v \in K(y, S_x)$, we get $K_t \subseteq K_\beta \cup K(y, S_x)$. This implies that the vertex p of t lies inside K_β as p lies outside S_y. By Claim 3, p lies inside $K_\alpha^\perp$, contradicting the assumption that α is not encroached.

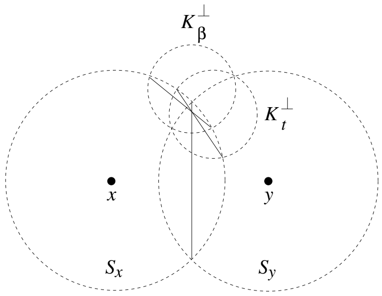

Figure 4: The three solid line segments delimit the caps K_β, K_t and $K(y, S_x)$ on S_x.

Case 2: v lies on a curved facet other than F on $\mathcal{B} \cap S_x$. We will reduce to case 1 by properly choosing a vertex p of t and a sphere to play the role of S_y in case 1. Let η be the shortest geodesic on S_x from v to a vertex q of t. Clearly, $\eta \subseteq K_t$. We claim that η does not cross $S_x \cap S_w$ for any protecting sphere S_w consecutive to S_x. Observe that $K(w, S_x)$ does not lie inside K_t, otherwise the emptiness of K_t would imply that $K_t \cap K(w, S_x)$ contains some wide helper arc. Thus, $K_t - \mathrm{int}(K(w, S_x))$ is star-shaped with respect to v and shortest geodesics on S_x originating from v. So our claim follows. Since $v \notin F$, our claim implies that η enters F from another curved facet at a curved edge e where e is also incident on a flat facet F' adjacent to F. Let L be the plane containing F'. Note that L passes through x. Consider the infinite sphere bounded by L with the halfspace containing v as its inside. We denote this infinite sphere by L^+. L^+ will play the role of S_y in case 1. We claim that t has a vertex outside L^+. Recall that the destination of η is a vertex q of t and η intersects L. If $q \notin L$, then $q \notin L^+$, otherwise η would be more than a semicircle and so $K_t \cap L$ is a complete circle. Since K_t is empty and $e \subseteq K_t \cap L$, this implies that e is a complete circle, contradiction. If $q \in L$, then $q \in e$. If t does not have any vertex outside L^+, then $K_t \subseteq L^+$. Using this and the fact that t belongs to F, we conclude that for any neighborhood $N(q)$ around q, $N(q) \cap F$ has points inside L^+. However, since η enters F at q, for a sufficiently small neighborhood $N(q)$ around q, $N(q) \cap F$ does not lie inside L^+, contradiction. This proves our claim that t has a vertex p outside L^+. To summarize, we have the same setting as in case 1 with S_y substituted by L^+: L^+ is orthogonal to S_x, $v \in L^+$, $p \notin L^+$, and $K_t \cap S_x \cap L$ is an arc within a helper arc α. Thus, the argument in case 1 shows that p encroaches upon α, contradiction. ∎

7 Boundary conformity

We are ready to prove that $\mathrm{Del}\,\mathcal{V}$ is conforming whenever no helper arc is wide or encroached and no subfacet is encroached. Thus, the output mesh is conforming when MESH terminates (termination will be proved in Sections 8 and 9). We start with a result characterizing the subcomplex of $\mathrm{Del}\,\mathcal{V}$ inside $\mathcal{B}$.

LEMMA 9 *Assume that* MESH *has not inserted any vertex inside $\mathcal{B}$. Let S_x and S_y be two consecutive protecting spheres. When there is no wide or encroached helper arc, (i)–(iii) hold. When there is no wide or encroached helper arc/triangle, (i)–(iv) hold.*

(i) For any flat facet F incident to x and any helper arc $\widehat{pq} \subseteq \mathcal{B} \cap S_x \cap F$, the equatorial sphere of pqx is empty.

(ii) For any helper arc endpoint $p \in S_x \cap S_y$, the equatorial sphere of pxy is empty.

(iii) For any helper arc $\widehat{pq} \subseteq S_x \cap S_y$, the circumsphere of $pqxy$ is empty.

(iv) For any helper triangle pqr on CH$_x$, the circumsphere of $pqrx$ is empty.

Proof. Consider (i). Let $\alpha = \widehat{pq}$. Let S be the equatorial sphere of pqx. Observe that the centers of S_x, S and $K_\alpha^\perp$ lie on a straight line. Since x lies on S but outside $K_\alpha^\perp$, the center of S lies between x and the center of $K_\alpha^\perp$. 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. Consider (ii). 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. Consider (iii). The circumsphere of $pqxy$ is the equatorial sphere of pxy which is empty by (ii). We can prove (iv) by considering the circumcap and normal sphere of pqr and employing the same arguments in proving (i). ∎

Next, we bootstrap from Lemma 9 to show that MESH never inserts any vertex inside $\mathcal{B}$.

LEMMA 10 MESH *never inserts any vertex inside $\mathcal{B}$.*

Proof. Assume to the contrary that MESH wants to insert a vertex v 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 Lemmas 6 and 8, v is not inserted by Rule 2. If v is inserted by Rule 3 to split an encroached subfacet τ, then τ does not lie inside $\mathcal{B}$ by Lemma 9. In fact, Lemma 9 further implies that τ belongs to a flat facet outside $\mathcal{B}$. But then v does not lie inside $\mathcal{B}$ by Lemma 7, contradiction. The remaining possibility is that Rule 4 wants to insert v inside some protecting sphere S_x. It follows that there is no wide or encroached helper arc/triangle. Let v be the circumcenter of the tetrahedron τ. By Rule 4, τ has no vertex inside $\mathcal{B}$. At least one vertex of τ is outside S_x as S_x cannot be the empty circumsphere of τ (x lies inside S_x). Let S be the circumsphere of τ. Let K be the cap on S_x that is bounded by $S_x \cap S$ and lies inside S. Since x does not lie inside S, the angular diameter of K is less than π, so $K^\perp$ is defined. If $K \cap K_t = \emptyset$ for all helper triangle t on CH_x, then $K \subseteq K(y, S_x)$ for some protecting sphere S_y consecutive to S_x. It follows that $\mathrm{Bd}(S_x \cup S_y)$ encloses S and hence τ, contradicting the fact that MESH has not inserted any vertex inside $\mathcal{B}$. Next, take a helper triangle t_0 on CH_x such that $K \cap K_{t_0} \neq \emptyset$. Starting from t_0, we visit a sequence of helper triangles $t_0, t_1, t_2, \cdots$ to derive a contradiction as follows.

Case 1: $K \subseteq K_{t_i}$. Clearly S lies inside $\mathrm{Bd}(S_x \cup K^\perp)$. So any vertex of τ outside S_x lies inside $K^\perp$. Since $K_{t_i}^\perp$ encloses $K^\perp$ by Claim 3, some vertex of τ encroaches upon t_i, contradiction.

Case 2: $K \not\subseteq K_{t_i}$. The vertices of t_i divide ∂K_{t_i} into three arcs and by emptiness of K, $\partial K \cap \partial K_{t_i}$ lie on one arc, say the one between vertices u and v of t_i.

Case 2.1: There is a helper triangle t_{i+1} on CH_x that shares uv with t_i. If $K \subseteq K_{t_i} \cup K_{t_{i+1}}$, (refer to Claim 2) we move a point z from the center of $K_{t_i}^\perp$ towards the center of $K_{t_{i+1}}^\perp$ and stop as soon as $\partial K(z, S_x)$ is tangent to ∂K. Tangency implies that $K \subseteq K(z, S_x)$, so $\mathrm{Bd}(S_x \cup K(z, S_x)^\perp)$ encloses S. Since z lies between the centers of $K_{t_i}^\perp$ and $K_{t_{i+1}}^\perp$, Claim 2 implies that $\mathrm{Bd}(K_{t_i}^\perp \cup K_{t_{i+1}}^\perp)$ encloses $K(z, S_x)^\perp$. So $\mathrm{Bd}(S_x \cup K_{t_i}^\perp \cup K_{t_{i+1}}^\perp)$ encloses $S_x \cup K(z, S_x)^\perp$ and hence S. Hence, some vertex of τ encroaches upon t_i or t_{i+1}, contradiction. If $K \not\subseteq K_{t_i} \cup K_{t_{i+1}}$, we continue to visit t_{i+1}. We will never return to t_i as $K \cap K_{t_i} \subset K \cap K_{t_{i+1}}$.

Case 2.2: If t_i is the only helper triangle on CH_x incident to uv, u and v are the endpoints of a helper arc $\alpha \subseteq S_x \cap S_y$ for some protecting sphere S_y consecutive to S_x. Let z be the center of $K_\alpha^\perp$. By Claim 1, $K_\alpha^\perp$ is orthogonal to S_x, so $K_\alpha^\perp = K(z, S_x)^\perp$. If $K \subseteq K_{t_i} \cup K(z, S_x)$, we conclude as in case 2.1 that $\mathrm{Bd}(S_x \cup K_{t_i}^\perp \cup K_\alpha^\perp)$ encloses S. So some vertex of τ encroaches upon t_i or α, contradiction. If $K \not\subseteq K_{t_i} \cup K(z, S_x)$, then $K \subseteq K(z, S_x) \cup K(y, S_x)$ which implies that $\mathrm{Bd}(S_x \cup K_\alpha^\perp \cup S_y)$ encloses S. Since no vertex of τ lies inside S_x or S_y, some vertex of τ encroaches upon α, contradiction. ◻

Finally, we put together Lemmas 6–10 and summarize the main results of this section.

COROLLARY 1 MESH *never inserts a vertex inside* $\mathcal{B}$.

(i) Whenever no helper arc is wide or encroached, the following hold.

- *(a) Subfacets inside* $\mathcal{B}$ *are not encroached.*
- *(b) Each subfacet that does not lie inside* $\mathcal{B}$ *belongs to exactly one flat facet. Each helper triangle belongs to exactly one curved facet.*
- *(c) The center of the circumcap of a subfacet (resp. helper triangle)* σ *lies on the flat facet (resp. curved facet) that* σ *belongs to.*

(ii) Whenever no helper arc/triangle is wide or encroached, the following hold.

- *(a) For any two consecutive protecting spheres* S_x *and* S_y *and any helper arc* $\widehat{pq} \subseteq S_x \cap S_y$, $pqxy \in \mathrm{Del}\,\mathcal{V}$.
- *(b) For any protecting sphere* S_x *and any helper triangle* pqr *on* CH_x, $pqrx \in \mathrm{Del}\,\mathcal{V}$.

(iii) Whenever no helper arc is wide or encroached and no subfacet is encroached, $\mathrm{Del}\,\mathcal{V}$ *is conforming.*

8 Between adjacent elements

The termination of MESH hinges on the fact that we will not keep generating encroached helper arc, helper triangle or subfacet. In particular, if a new vertex inserted on one element encroaches upon something on an adjacent and non-incident element and if this happens indefinitely, then algorithm will not terminate. In this section, we show that this cannot happen. Lemmas 11, 12 and 13 analyze the cases for helper arc, helper triangle and subfacet respectively. Lemmas 11 and 12 are stated more generally for their usage in Section 9.

LEMMA 11 *Let* β *be an arc on a curved edge* e *such that the angular width of* β *is less than* π. *If* E *is an element of* $\mathcal{Q}$ *such that* E *is adjacent to* e *and* $e \not\subseteq \partial E$, *then* E *does not intersect the inside of* $K_\beta^\perp$.

Proof. Case 1: e lies at the intersection between a protecting sphere S_x and a facet of $\mathcal{P}$. Then E is a curved edge or curved facet lying on a protecting sphere S_y consecutive to S_x. Since e lies outside S_y and e meets $S_x \cap S_y$ at right angle, the cone of rays from x through $S_x \cap S_y$ and the cone of rays from x through ∂K_β do not cross. Observe that S_y and $K_\beta^\perp$ lie inside their corresponding cones. Thus, S_y does not intersect the inside of $K_\beta^\perp$ and neither does E.

Case 2: $e \subseteq S_x \cap S_y$ for two consecutive protecting spheres S_x and S_y. The endpoints of e lie on two facets F_1 and F_2 of $\mathcal{P}$. Note that $x, y \in F_1 \cap F_2$. Let H_i be the halfplane that is bounded by the supporting line of xy and contains the endpoint of e on F_i. For i = 1 or 2, since e meets H_i at right angle, either H_i avoids $K_\beta^\perp$ or H_i is tangent to $K_\beta^\perp$. Observe that either $E \subseteq H_i$ for some i or E is separated from e by H_1 and H_2. It follows that E does not intersect the inside of $K_\beta^\perp$. ◻

LEMMA 12 *Suppose that there is no wide or encroached helper arc. Let t be a helper triangle belonging to a curved facet F. Let $K \subseteq K_t$ be a cap with the same center as K_t and angular diameter less than π. Let E be an element of $\mathcal{Q}$ adjacent to F. For any vertex $v \in \mathcal{V} \cap E$, v does not lie inside $K^\perp$.*

Proof. Let $F \subseteq S_x$ for some protecting sphere S_x. Assume to the contrary that v lies inside $K^\perp$. Observe that $E \not\subseteq S_x$, otherwise the emptiness of K would be contradicted.

Case 1: E is a curved edge or curved facet lying on a protecting sphere S_y consecutive to S_x. In order that $K^\perp$ intersects E, K must cross $S_x \cap S_y$. Otherwise, the cone from x through ∂K and the cone from x through $S_x \cap S_y$ do not cross, implying that S_y does not intersect the inside of $K^\perp$, contradiction. By emptiness of K, $K \cap S_x \cap S_y$ is an arc β within a helper arc α. Let z_β and z be the centers of $K_\beta^\perp$ and $K^\perp$ respectively. By Claim 1, $K_\beta^\perp$ is orthogonal to S_x. Since S_y and $K^\perp$ are also orthogonal to S_x, Claim 2 implies that S_y, $K^\perp$ and $K_\beta^\perp$ intersect at the same circle and y, z_β and z are collinear. If y lies between z_β and z, the subset of S_y inside $K_\beta^\perp$ lies outside $K^\perp$. Since β lies on the subset of S_y inside $K_\beta^\perp$, β is outside $K^\perp$, contradicting the fact that $\beta \subseteq K$. If y does not lie between z_β and z, the subset of S_y inside $K^\perp$ is equal to the subset of S_y inside $K_\beta^\perp$. Since v lies on the subset of S_y inside $K^\perp$, v lies inside $K_\beta^\perp$ and hence $K_\alpha^\perp$ by Claim 3. This contradicts the assumption that α is not encroached.

Case 2: E is a flat facet or a curved boundary edge of a flat facet. Let H be the plane containing the corresponding flat facet. Note that H passes through x. Since v lies inside $K^\perp$, $K^\perp$ intersects $S_x \cap H$ at an arc β within a helper arc α. Since $K_\beta^\perp$ and $K^\perp$ are orthogonal to S_x, $H \cap K_\beta^\perp = H \cap K^\perp$ by Claim 2. Since v lies inside $H \cap K^\perp$, v lies inside $K_\beta^\perp$. But then v also lies inside $K_\alpha^\perp$ by Claim 3, contradiction. ◻

LEMMA 13 *Suppose that there is no wide or encroached helper arc. Let τ be a subfacet belonging to a flat facet F. Let E be an element of $\mathcal{Q}$ adjacent to F. For any vertex $v \in \mathcal{V} \cap E$, v does not lie inside $K_\tau^\perp$.*

Proof. The proof is similar to that of Lemma 12 by treating the supporting plane of F as an infinite sphere. ◻

9 Insertion radius

For each vertex 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/rejected by MESH, r_v is the minimum distance to a vertex in $\mathcal{V}$ at the time when v is inserted/rejected. In this section, we prove a lower bound on the insertion radii of vertices. Thus, MESH must terminate by a packing argument.

We first introduce some notations. Consider the time when MESH inserts/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 center of K_σ for a wide helper arc/triangle σ, the parent of v is undefined. Suppose that v is the center of K_σ where σ is a non-wide encroached helper arc/triangle or an encroached subfacet. If $\mathcal{V}$ has a vertex encroaching upon σ (i.e., lying inside $K_\sigma^\perp$), the the parent of v is the nearest encroaching vertex in $\mathcal{V}$. Otherwise, $K_\sigma^\perp$ is empty. What happens is that MESH rejected a vertex p because p encroached upon σ and this also prompted MESH to consider v. The parent of v is p in this case. If v is the circumcenter of a tetrahedron τ, the parent of v is the endpoint of the shortest edge of τ that appeared in $\mathcal{V}$ the latest. Finally, the parents of vertices of $\mathcal{Q}$ are undefined.

We will use induction. To this end, Lemma 14 relates the insertion radius of v to the insertion radius of its parent p and to $||p - v||$. The proof of Lemma 14 needs the following claim.

CLAIM 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 $K^\perp$. Since the angular diameter of K is at most $\pi/3$, $||v - z|| < \mathrm{radius}(K^\perp) \cdot \cos(\pi/3)$ which is at most $||q - z||/2$. By triangle inequality, $||q - v|| \geq ||q - z|| - ||v - z||$. It follows that

$$||q - v|| > ||q - z||/2.$$

Since p and v lie inside $K^\perp$, $||q - z|| \geq ||p - v||/2$. Thus, $||q - v|| > ||p - v||/4$. By triangle inequality, $||p - q|| \leq ||p - z|| + ||q - z||$ which is at most $2 \cdot ||q - z||$. Thus, $||q - v|| > ||p - q||/4$. ◻

LEMMA 14 *Let v be a vertex of $\mathcal{Q}$ or a vertex inserted/rejected by* MESH*. Let p be the parent of v.*

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

(ii) Otherwise, $r_v > ||p - v||/4$ and if $r_v \leq \widehat{f}(v)/4$, the following hold depending on the type of v:

Type 1: *p has type 2, 3 or 4 and $r_v > r_p/4$.*

Type 2 or 3: *p has type 4 and $r_v > r_p/4$.*

Type 4: $r_v > \rho_0 \cdot r_p$.

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

Case 1: v is the center of K_σ for a wide helper arc/triangle σ. The parent p is undefined in this case. If σ is a helper arc, let S be the equatorial sphere of σ. Note that S is either a protecting sphere or the equatorial sphere of the common hole between two consecutive protecting spheres. If σ is a helper triangle, let S be the protecting sphere that contains the vertices of σ. Let E be the element of Q that σ belongs to (i.e., E is a curved edge or curved facet depending on whether σ is a helper arc or helper triangle). Note that E lies on S. By Corollary 1, $v \in E$. Let $K \subseteq K_\sigma$ be the cap with center v and angular diameter $\pi/3$. Let B be the smallest ball centered at v that contains K. Let z be the center of S. Suppose that $\mathrm{int}(B)$ does not contain any vertex in $\mathcal{V}$. Then $r_v \geq \mathrm{radius}(B) = ||v - z|| \cdot 2\sin(\pi/12) \geq ||v - z||/2$. Observe that z lies on some linear edge of $\mathcal{Q}$ that stabs S. Since all linear edges are disjoint from $\mathcal{B}$, we have $||v - z|| \geq \widehat{f}(v)$ by definition. It follows that $r_v \geq \widehat{f}(v)/2$. Suppose that $\mathrm{int}(B)$ contains a vertex $w \in \mathcal{V}$. Observe that $\mathrm{Bd}(S \cup K^\perp)$ encloses B which implies that w lies inside $K^\perp$. If w is a vertex of $\mathcal{Q}$, then $w \notin E$ as vertices on E do not lie inside $K^\perp$, so $||v - w|| \geq \widehat{f}(v)$ by definition. Otherwise, w was inserted by MESH. We claim that w lies on an element E' of $\mathcal{Q}$ disjoint from E. If σ is a wide helper arc, then MESH has split helper arcs only so far, so w lies on some curved edge E'. By Lemma 11, for $K^\perp$ to enclose w, E' is disjoint from E. If σ is a wide helper triangle, then MESH has split helper arcs/triangles only so far, so w lies on some element E' of $\mathcal{Q}$. By Lemma 12, for $K^\perp$ to enclose w, E' is disjoint from E. This proves the claim. Our claim implies that $||v - w|| \geq \widehat{f}(v)$. It follows that $r_v = \min_{w \in B} ||v - w|| \geq \widehat{f}(v)$.

Case 2: v is the midpoint of a non-wide encroached helper arc α. Note that v has type 1. Let e be the curved edge that α belongs to. Let q be the vertex in $\mathcal{V}$ such that $r_v = ||q - v||$. Recall that p is the parent of v. We first relate r_v to $||p - v||$. If q lies inside $K_\alpha^\perp$, then $p = q$ by definition of parent; otherwise, $||q - v|| > ||p - v||/4$ by Claim 4. Hence, $r_v = ||q - v|| > ||p - v||/4$. Next, we relate r_v to $\widehat{f}(v)$ and r_p. If p is a vertex of $\mathcal{Q}$, then $p \notin e$ as vertices on e do not lie inside $K_\alpha^\perp$, so $||p - v|| \geq \widehat{f}(v)$. If p lies on an element E of $\mathcal{Q}$ such that $e \not\subseteq \partial E$, Lemma 11 implies that e and E are disjoint and so $||p - v|| \geq \widehat{f}(v)$. Since $r_v > ||p - v||/4$, we get $r_v > \widehat{f}(v)/4$ for the above two cases. The remaining case is that p has type 4 or p lies on a curved/flat facet whose boundary contains e. Note that p has type 2, 3 or 4. What happens is that MESH attempted to insert p but since p encroached upon α, MESH rejected p and inserts v to split α now. In this case, q does not lie inside $K_\alpha^\perp$, otherwise the parent of v would be q instead. Since $q \in \mathcal{V}$ when p was rejected, $r_p \leq ||p - q||$. By Claim 4, $||q - v|| > ||p - q||/4$. It follows that $r_v > r_p/4$.

Case 3: v is the center of K_σ where σ is a non-wide encroached helper triangle or an encroached subfacet. Note that v has type 2 or 3. Let F be the curved facet or flat facet that σ belongs to, whichever is appropriate. Let q be the vertex in $\mathcal{V}$ such that $r_v = ||q - v||$. We first relate r_v to $||p - v||$. If q lies inside $K_\sigma^\perp$, then $p = q$ by definition of parent; otherwise, $||q - v|| > ||p - v||/4$ by Claim 4. Hence, $r_v = ||q - v|| > ||p - v||/4$. Next, we relate r_v to $\widehat{f}(v)$ and r_p. Suppose that p is a vertex of $\mathcal{Q}$ or p has type 1, 2 or 3. Vertices of type 1 are always inserted. If p has type 2 or 3, although p encroached upon σ, p was inserted as v has type 2 or 3. It follows that p is a vertex in $\mathcal{V} \cap E$ for some element E of $\mathcal{Q}$. We invoke

Lemma 12 if σ is a helper triangle or Lemma 13 if σ is a subfacet. The implication is that E is disjoint from F. Since $v \in F$ by Corollary 1, $||p - v|| \geq \widehat{f}(v)$. Since $r_v > ||p - v||/4$, we get $r_v > \widehat{f}(v)/4$. The remaining case is that p has type 4. By Rule 4, p was rejected for encroaching upon σ. In this case, q does not lie inside $K_\alpha^\perp$, otherwise the parent of v would be q instead. Since $q \in \mathcal{V}$ when p was rejected, $r_p \leq ||p - q||$. By Claim 4, $||q - v|| > ||p - q||/4$. It follows that $r_v > r_p/4$.

Case 4: v is the circumcenter of a tetrahedron τ. By definition, p is an endpoint of the shortest edge of τ. Let q be the other endpoint of this edge. If p is a vertex of $\mathcal{Q}$, by the definition of parent, q is also a vertex of $\mathcal{Q}$. This implies that $r_v = ||p - v|| = ||q - v|| = \widehat{f}(v)$. If p is not a vertex of Q, since $\rho(\tau) > \rho_0$, $r_v = ||p - v|| > \rho_0 \cdot ||p - q|| \geq \rho_0 r_p$. ◻

We prove one more claim and then derive the lower bounds for insertion radii in Lemma 15.

CLAIM 5 *Let v be a vertex of $\mathcal{Q}$ or inserted/rejected by* MESH. *Let p be the parent of v. If $r_v > c \cdot r_p$, then $\widehat{f}(v) < \widehat{f}(p) \cdot r_v/(c \cdot r_p) + 4r_v$.*

Proof. Since p is defined, $r_v > ||p - v||/4$ by Lemma 14. Using the Lipschitz property, we get $\widehat{f}(v) \leq \widehat{f}(p) + ||p - v|| < \widehat{f}(p) \cdot r_v/(c \cdot r_p) + 4r_v$. ◻

LEMMA 15 *Let v be a vertex of $\mathcal{Q}$ or inserted/rejected by* MESH. *If v is a vertex of $\mathcal{Q}$, then $r_v \geq \widehat{f}(v)$. Otherwise, there are four constants $C_1 > C_2 = C_3 > C_4 > 4$ such that if v has type i, then $r_v > \widehat{f}(v)/C_i$.*

Proof. We prove the lemma by induction using the constants $C_1 = 84\rho_0/(\rho_0 - 16)$, $C_2 = C_3 = (20\rho_0 + 16)/(\rho_0 - 16)$ and $C_4 = (4\rho_0 + 20)/(\rho_0 - 16)$. Before MESH starts, $r_v \geq \widehat{f}(v)$ 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 14 implies that the parent p of v is defined.

If v has type 1, by Lemma 14, p has type 2, 3 or 4 and $r_v > r_p/4$. By induction assumption, $\widehat{f}(p) < C_2 r_p$. By Claim 5, $\widehat{f}(v) < 4C_2 r_v + 4r_v = C_1 r_v$.

If v has type 2 or 3, by Lemma 14, p has type 4 and $r_v > r_p/4$. By induction assumption, $\widehat{f}(p) < C_4 r_p$. By Claim 5, $\widehat{f}(v) < 4C_4 r_v + 4r_v = C_2 r_v$.

If v has type 4, then $r_v > \rho_0 r_p$ by Lemma 14. By induction assumption, $\widehat{f}(p) < C_1 r_p$ regardless of whether p is a vertex of $\mathcal{Q}$ or p was inserted/rejected. By Claim 5, $\widehat{f}(v) < C_1 r_v/\rho_0 + 4r_v = C_4 r_v$. ◻

We are ready to prove that MESH terminates by a packing argument.

COROLLARY 2 MESH *terminates and for each output vertex v, its shortest incident edge has length at least $\widehat{f}(v)/(1 + C_1)$.*

Proof. Let vw be the shortest 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 15. If v appeared in $\mathcal{V}$ before w, then $||v - w|| \geq r_w \geq \widehat{f}(w)/C_1$ by Lemma 15. 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 $\mathcal{Q}$. Since $\widehat{f}_{\min} > 0$ and the input domain has bounded volume, there is a finite number of output vertices. It follows that MESH terminates. ◻

10 Mesh quality

In this section, we relate the edge lengths to local feature size with respect to $\mathcal{P}$, bound the radius-edge ratio and summarize the guarantees offered by MESH. We first prove in Lemmas 16 and 17 that $g(p) = \Omega(f(p))$ for each output vertex p.

LEMMA 16 *Let uv be an edge of $\mathcal{P}$. Let q be a point on uv. There exists a constant $k_1 < 1$ such that*

(i) If $||q-u|| \geq (\mu/2) \cdot f(u)$ and $||q-v|| \geq (\mu/2) \cdot f(v)$, then $g(q) \geq k_1 \mu \cdot f(q)$.

(ii) For any point p on or outside $\mathcal{B}$, $g(p) + ||p - q|| \geq k_1 \mu \cdot f(p)$.

Proof. We prove the lemma for the constant $k_1 = \sin\phi/4$. Consider (i). Let B be the ball centered at q with radius $g(q)$. If B intersects two disjoint elements of $\mathcal{P}$, $g(q) = f(q)$. Otherwise, we can assume that B touches u or the interior of an edge/facet of $\mathcal{P}$ incident to u. So $g(q) \geq ||q - u|| \cdot \sin\phi$. By the Lipschitz condition, $f(q) \leq f(u) + ||q - u||$. Since $||q - u|| \geq (\mu/2) \cdot f(u)$, we get $f(q) \leq ((2 + \mu)/\mu) \cdot ||q - u||$. So $f(q) \leq ((2 + \mu)/(\mu \sin\phi)) \cdot g(q) < g(q)/(k_1\mu)$. Consider (ii). Suppose that $||q - u|| < (\mu/2) \cdot f(u)$. Using the Lipschitz condition and the fact that

$||p-u|| \geq \mu \cdot f(u)$, we get $f(p) \leq f(u) + ||p-u|| \leq ((1+\mu)/\mu) \cdot ||p-u||$. Since $||q-u|| < (\mu/2) \cdot f(u)$, $||q-u|| < ||p-u||/2$. Using triangle inequality, we get $||p-q|| \geq ||p-u|| - ||q-u|| > ||p-u||/2$. Thus, $f(p) < ((2+2\mu)/\mu) \cdot ||p-q|| < ||p-q||/(k_1\mu)$. We get the same result for the case where $||q-v|| < (\mu/2) \cdot f(v)$. If $||q-u|| > (\mu/2) \cdot f(u)$ and $||q-v|| > (\mu/2) \cdot f(v)$, then using $f(p) \leq f(q) + ||p-q||$ and (i), we get $f(p) \leq g(q)/(k_1\mu) + ||p-q|| < (g(q) + ||p-q||)/(k_1\mu)$. ◻

LEMMA 17 *For each vertex p in the final mesh, $g(p) \geq k_2\mu \cdot f(p)$ for some constant $k_2 < k_1$.*

Proof. We prove the lemma for the constant $k_2 = \min\{k_1/2, k_1 \sin(\phi/2)/(1+\sin(\phi/2))\}$. If p is a vertex of $\mathcal{P}$, then $g(p) = f(p)$. Otherwise, if p is a linear edge endpoint, then for each endpoint v of the edge of $\mathcal{P}$ that contains p, $||p-v|| > (\mu/2) \cdot f(v)$. By Lemma 16(i), $g(p) \geq k_1\mu \cdot f(p)$. The remaining case is 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 $\mathcal{P}$, $g(p) = f(p)$. Suppose not. If B intersects an edge uv, then for any point $q \in B \cap uv$, $||p-q|| \leq g(p)$. Using Lemma 16(ii), we get $g(p) \geq (k_1\mu/2) \cdot f(p)$. Otherwise, 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 $H_1 \cap H_2$ nearest to p. Since pr makes an angle at least $\phi/2$ with H_1 or H_2, we have $||p-r|| \cdot \sin(\phi/2) \leq g(p)$. The orthogonal projections of pr onto H_1 and H_2 must intersect ∂F_1 or ∂F_2 at some point q. Observe that $||p-q|| \leq ||p-r||$, so $||p-q|| \leq g(p)/\sin(\phi/2)$. Using Lemma 16(ii), we get $k_1\mu \cdot f(p) \leq g(p) + ||p-q|| \leq g(p) \cdot (1+\sin(\phi/2))/\sin(\phi/2)$. ◻

We are ready to prove the main results of this paper.

THEOREM 1 MESH *terminates and produces a Delaunay mesh $\mathcal{M}$ conforming to $\mathcal{P}$. There exists two constants δ and ρ_1 depending on μ and ϕ such that*

(i) For each vertex v of $\mathcal{M}$, the length of the shortest edge incident to v is at least $\delta \cdot f(v)$.

(ii) Let τ be a tetrahedron in $\mathcal{M}$. If τ does not have a vertex inside $\mathcal{B}$, then $\rho(\tau) \leq \rho_0$; otherwise, $\rho(\tau) \leq \rho_1$.

Proof. The termination of MESH has been proved in Corollary 2. Since MESH terminates, Corollary 1 implies that $\mathcal{M}$ is conforming.

We prove (i) for the constant $\delta = \min\{k_2\lambda\mu^9/(1+C_1), k_2c_3\mu^2\}$. Let v be a vertex of $\mathcal{M}$. Consider the case where v lies on or outside $\mathcal{B}$. Lemmas 5 and 17 imply that $\widehat{f}(v) \geq \lambda k_2\mu^9 \cdot f(v)$. By Corollary 2, the shortest edge incident to v has length at least $\widehat{f}(v)/(1+C_1)$ which is at least $(\lambda k_2\mu^9/(1+C_1)) \cdot f(v)$. Consider the case where v lies inside $\mathcal{B}$. Then v is a linear edge endpoint. By Lemma 2(iii), the shortest edge incident to v has length at least $c_3\mu \cdot g(v)$. By Lemma 17, $c_3\mu \cdot g(v) \geq k_2c_3\mu^2 \cdot f(v)$.

We prove (ii) for the constant $\rho_1 = 3\mu/(\delta(1-3\mu))$. If τ does not have a vertex inside $\mathcal{B}$, Rule 4 guarantees that $\rho(\tau) \leq \rho_0$. Otherwise, Corollary 1 implies that there are two possibilities.

Case 1: There exists a protecting sphere S_x such that $\tau = pqrx$ for some helper triangle pqr on CH_x. Since the angular diameter of the cap K_{pqr} is at most $\pi/3$, the circumradius of τ is less than $\mathrm{radius}(S_x) \leq 3\mu \cdot g(x)$. Assume that p is an endpoint of the shortest edge of τ. By (i), the shortest edge length of τ is at least $\delta \cdot f(p)$. Using the Lipschitz condition, we get $f(p) \geq f(x) - ||p-x|| \geq f(x) - 3\mu \cdot g(x) \geq (1-3\mu) \cdot f(x)$. Thus, the shortest edge length of τ is at least $\delta(1-3\mu) \cdot f(x)$. It follows that $\rho(\tau) < 3\mu/(\delta(1-3\mu))$.

Case 2: There exists consecutive protecting spheres S_x and S_y such that $\tau = pqxy$ for some helper arc $\widehat{pq}$ on $S_x \cap S_y$. The circumradius of τ is less than $\mathrm{radius}(S_x) \leq 3\mu \cdot g(x)$. Since x lies outside S_y and y lies outside S_x, xy is longer than some edge of τ (e.g., px or py). Thus, the shortest edge of τ is incident to p or q. Since p and q lie on S_x, case 1 shows that the shortest edge length of τ is at least $\delta(1-3\mu) \cdot f(x)$. So $\rho(\tau) < 3\mu/(\delta(1-3\mu))$. ◻

11 Conclusion

The constants may be improvable using a more refined analysis. We also plan an experimental study of the algorithm. We will look into the possibility of incorporating weighted Delaunay refinement [2] into our algorithm to eliminate slivers and guarantee bounded aspect ratio in the presence of small angles.

References

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

[2] S.-W. Cheng and T.K. Dey. Quality meshing with weighted Delaunay refinement. *Proc. 13th ACM-SIAM Sympos. Discrete Alg.*, 2002, 137–146.

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

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

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

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

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

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

[9] 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.

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

[11] M. Murphy, D. M. Mount, C. W. Gable. A point-placement strategy for conforming Delaunay tetrahedralization, *Proc. 11th ACM-SIAM Sympos. Discrete Alg.*, 2000, 67–74.

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

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

[14] J.R. Shewchuk. Mesh generation for domains with small angles. *Proc. 16th ACM Sympos. Comput. Geom.*, 2000, 1–10.

12 Appendix

12.1 Proof of Lemma 2

We first prove Lemma 2(i) as a separate claim.

CLAIM 6 *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 $\mathrm{radius}(S_{u_v})/g(u_v)$ and $\mathrm{radius}(S_{v_u})/g(v_u)$ 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 $2\mu < \cos(\mu\theta^u_{uv})$, $d < g(u)$ which implies that $u \in E$. 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, let ψ be the angle between uv and E. Since $\pi/2 > \psi \geq \phi^u_{uv}$ and $g(u_v) = \|u-u_v\| \cdot \sin\psi$, our claim is true. Let $R = \mathrm{radius}(S_{u_v})/g(u_v)$. It follows 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, $\sin\theta^u_{uv} = \min\{\sin(\pi/3), \sin\phi^u_{uv}\} \geq \min\{\sqrt{3}/2, \sin\phi\}$. If $\phi^u_{uv} \leq \pi/3$, then $\mu\theta^u_{uv}/\sin\phi^u_{uv} = \mu\phi^u_{uv}/\sin\phi^u_{uv}$ which is maximized when $\phi^u_{uv} = \pi/3$. If $\phi^u_{uv} > \pi/3$, then $\mu\theta^u_{uv}/\sin\phi^u_{uv} < \pi\mu/(3\sin(\pi/3)) = 2\pi\mu/(3\sqrt{3})$. ◻

Next, we show that when *Split*(x, y) is called, there is always a gap between S_x and S_y.

CLAIM 7 *Let $k = 1.099$. Whenever $Split(x, y)$ is called, the spheres centered at x and y with radii $k \cdot \mathrm{radius}(S_x)$ and $k \cdot \mathrm{radius}(S_y)$ do not intersect.*

Proof. Given a sphere S, let $\overline{S}$ denote the sphere with the same center as S and radius $k \cdot \mathrm{radius}(S)$. Let uv be an edge of $\mathcal{P}$. We first show that $\overline{S_{u_v}} \cap \overline{S_{v_u}} = \emptyset$. Since $\theta^u_{uv} \leq \pi/3$ and $g(u) \leq \|u - v\|$, it follows from definition that $\|u - u_v\| < 2\mu \cdot g(u) \leq 2\mu \cdot \|u - v\|$ and $\mathrm{radius}(S_{u_v}) \leq (\mu\pi/3) \cdot \|u-u_v\| < (2\pi\mu^2/3) \cdot \|u-v\|$. So $\|u-u_v\| + \mathrm{radius}(\underline{S_{u_v}}) < (2\mu + 2\pi\mu^2/3) \cdot \|u-v\| < \|u-v\|/2$, implying 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 $\overline{S_x} \cap \overline{S_y} = \emptyset$. Observe that z lies outside S_x and S_y. Since $\pi(z, S_x) = Z^2$ and line 3 of *Split* is satisfied,

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

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

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

Let E be an element of $\mathcal{P}$ such that $z \notin E$ and E touches the ball centered at z with radius $g(z)$. Let d be the minimum distance between x and E. By triangle inequality, (1) and (2), we get $d \leq ||x-z|| + g(z) < ((1+3\mu)/(3\mu)) \cdot ||x-z|| < k(1+3\mu)/(\mu(3-k)) \cdot \mathrm{radius}(S_x)$. If $x = u_v$, then $\mathrm{radius}(S_x) \leq c_1\mu \cdot g(x)$ by Claim 6, otherwise $\mathrm{radius}(S_x) = \mu \cdot g(x)$. Since $c_1 > 1$, $d < (c_1 k(1+3\mu)/(3-k)) \cdot g(x)$. By our choices of k, c_1 and μ, one can verify that $c_1 k(1+3\mu)/(3-k) < 1$. However, since $x, z \in \mathrm{int}(uv)$ and $z \notin E$, we have $x \notin E$ which implies $d \geq g(x)$, contradiction. ◻

The gap between S_x and S_y in Claim 7 implies that when we create a protecting sphere S_z between S_x and S_y, z cannot be too close to x and y and S_z cannot be too small. This is the main idea behind the proof of Lemma 2. The details are given below.

Proof of Lemma 2:

We prove the lemma for the constants $c_3 = c_2(k-1)/(1+c_2k)$ and $c_4 = (1/c_1) \cdot \min\{\sqrt{3}/2, \sin\phi\}$, where k is the constant in Claim 7. Clearly, (i) is equivalent to Claim 6.

Consider (ii). If $Split(u_v, v_u)$ does not terminate, Claim 7 implies that infinitely many non-intersecting protecting spheres are created in line 4 of *Split*. 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 consecutive protecting spheres created are orthogonal and hence overlapping. Thus, the spheres in $\mathcal{S}$ cover $u_v v_u$. Take a sphere $S_z \in \mathcal{S} - \{S_{u_v}, S_{v_u}\}$. By lines 3 and 4, $\mathrm{radius}(S_z)/g(z) \leq 3\mu$. 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. Claim 7 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 $||x-z|| \leq Z + \mathrm{radius}(S_x) \leq kZ/(k-1)$. Using this and Lemma 1, we get

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

If $x = u_v$, then $\mathrm{radius}(S_x) \geq c_2\mu \cdot g(x)$ by Claim 6, otherwise $\mathrm{radius}(S_x) = \mu \cdot g(x)$. So (3) yields $Z \geq c_2\mu(k-1) \cdot g(x)$. Substituting this into (4), we get $g(z) \leq Z(1+c_2\mu k)/(c_2\mu(k-1))$ which is less than $Z/(c_3\mu)$. Hence, $Z \geq c_3\mu \cdot g(z)$.

Consider (iii). 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 (i) and (ii). Suppose that $x = u$ or v. Then $y = u_v$ or v_u respectively. It follows from definition that $\mathrm{radius}(S_y) = \mu \tan(\mu\theta^u_{uv}) \cdot g(x)$. Note that $\tan(\mu\theta^u_{uv}) \geq \mu \sin\theta^u_{uv} \geq \mu \cdot \min\{\sqrt{3}/2, \sin\phi\} = c_1c_4\mu$. So $\mathrm{radius}(S_y) \geq c_1c_4\mu^2 \cdot g(x)$. Using this and the fact that $\mathrm{radius}(S_y) \leq c_1\mu \cdot g(y)$ by Claim 6, we get $g(y) \geq c_4\mu \cdot g(x)$. Suppose that $x \in \mathrm{int}(uv)$. Since S_x intersects S_y, $||x-y|| \leq \mathrm{radius}(S_x) + \mathrm{radius}(S_y)$ which is at most $3\mu(g(x)+g(y))$. Using this and Lemma 1, we get $g(y) \geq g(x) - ||x-y|| \geq (1-3\mu) \cdot g(x) - 3\mu \cdot g(y)$, so $g(y) \geq ((1-3\mu)/(1+3\mu)) \cdot g(x)$. Observe that $(1-3\mu)/(1+3\mu) > \mu > c_4\mu$. ◻

12.2 Proof of Lemma 3

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). Consider a hole on $\mathcal{B} \cap S_x$ bounded by $S_x \cap S_z$ for some protecting sphere S_z consecutive to S_x. By Lemma 2(ii), $\min\{\mathrm{radius}(S_x), \mathrm{radius}(S_z)\} \geq c_3\mu \cdot \min\{g(x), g(z)\}$. By Lemma 2(iii), $g(z) \geq c_4\mu \cdot g(x)$ which implies that $\min\{\mathrm{radius}(S_x), \mathrm{radius}(S_z)\} \geq c_3c_4\mu^2 \cdot g(x)$. Since S_x intersects S_z at right angle, the radius of the hole is at least $\min\{\mathrm{radius}(S_x), \mathrm{radius}(S_z)\}/\sqrt{2} \geq (c_3c_4\mu^2/\sqrt{2}) \cdot g(x) = c_5\mu^2 \cdot g(x)$.

Consider(ii). Let $\mathcal{B} \cap S_z$ be a ring adjacent to $\mathcal{B} \cap S_x$. We have $||x-z|| \leq \mathrm{radius}(S_x) + \mathrm{radius}(S_z)$ which is at most $3\mu \cdot g(x) + 3\mu \cdot g(z)$ by Lemma 2(ii). By Lemma 2(iii), $g(x) \geq c_4\mu \cdot g(z)$. It follows that

$$||x-z|| \leq ((3+3\mu c_4)/c_4) \cdot g(x). \tag{5}$$

Let d be the distance between x and the bisector plane of S_x and S_z. By orthogonality, $d = \mathrm{radius}(S_x)^2/||x-z||$. Since $\mathrm{radius}(S_x) \geq c_3\mu \cdot g(x)$ by Lemma 2(ii), $d \geq (c_3\mu \cdot g(x))^2/||x-z||$. By (5), we get $d \geq (c_3^2c_4\mu^2/(3+3\mu c_4)) \cdot g(x)$ which is larger than $c_6\mu^2 \cdot g(x)$.

(iii) follows from the facts that the distance between x and E is at least $g(x)$ and $\mathrm{radius}(S_x) \leq 3\mu \cdot g(x)$.

Consider (iv). Let d be the minimum distance between $\mathcal{B} \cap S_x$ and $\mathcal{B} \cap S_y$. If x and y do not lie on the same edge of $\mathcal{P}$, then $d \geq ||x-y|| - \mathrm{radius}(S_x) - \mathrm{radius}(S_y) \geq ||x-y|| - 3\mu(g(x)+g(y))$. Since $||x-y|| \geq \max\{g(x), g(y)\}$, we get $d \geq (1-6\mu) \cdot ||x-y|| \geq (1-6\mu) \cdot g(x)$. Observe that $1-6\mu \geq \mu > c_7\mu^3$. If x and y lie on the same edge of $\mathcal{P}$, then $\mathcal{B} \cap S_x$ and $\mathcal{B} \cap S_y$ are separated by a ring $\mathcal{B} \cap S_z$ adjacent to $\mathcal{B} \cap S_x$. By (ii), the width of $\mathcal{B} \cap S_z$ is at least $c_6\mu^2 \cdot g(z)$. By Lemma 2(iii), $g(z) \geq c_4\mu \cdot g(x)$. Therefore, $d \geq c_4c_6\mu^3 \cdot g(x) = c_7\mu^3 \cdot g(x)$. ◻

12.3 Proof of Lemma 5

We give an overview of our proof strategy. Let p be a point on or outside $\mathcal{B}$. 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 analyze the distance d between E and E'. The difficult cases are when E and E' lie on the same protecting sphere or two consecutive protecting spheres, or when E lies on a protecting sphere and E' is an adjacent flat facet. We proceed in three steps. First, we consider some geodesic γ from E to E' on $\mathcal{B}$ and show that $\mathrm{length}(\gamma) = O(d)$. Second, we argue that $\mathrm{length}(\gamma)$ is at least the distance between two disjoint boundary elements of some curved facet. Third, we show that this distance is at least a constant factor of the local gap size. These three steps are described in the Claims 8–10. Afterwards, we give the proof of Lemma 5.

CLAIM 8 *Let F be a curved facet on $\mathcal{B} \cap S_x$ for some protecting sphere S_x. The minimum distance between two disjoint boundary elements of F is at least $c_8\mu^4 \cdot g(x)$ for some constant $c_8 < c_7$.*

Proof. We prove the claim for $c_8 = \min\{c_5 \sin\phi, c_4c_7\}$. Let d be the minimum distance between two disjoint boundary elements of F. Since F has at least four boundary edges, d is achieved by the minimum distance between two disjoint boundary edges (including their endpoints), say e and e'.

Case 1: e and e' lie on some facets E and E' of $\mathcal{P}$ respectively. Note that $x \in E \cap E'$. If $E \cap E' = \{x\}$, x is a vertex of $\mathcal{P}$. Since the angle between E and E' is at least ϕ, we get $d \geq 2\sin(\phi/2) \cdot \mathrm{radius}(S_x) = 2\mu\sin(\phi/2) \cdot g(x) \geq \mu \sin\phi \cdot g(x)$. If $\{x\} \subset E \cap E'$, $E \cap E'$ is an edge of $\mathcal{P}$. Note that this edge passes through hole(s) on $\mathcal{B} \cap S_x$. So $d \geq 2r\sin(\phi/2)$ where r is the minimum radius of the hole(s). By Lemma 3(i), $r \geq c_5\mu^2 \cdot g(x)$, so $d \geq c_5\mu^2 \sin\phi \cdot g(x)$.

Case 2: e lies on a facet E of $\mathcal{P}$ and e' lies on the boundary of a hole on $\mathcal{B} \cap S_x$. This case can happen only when x is a vertex of $\mathcal{P}$. (Otherwise, $\mathcal{B} \cap S_x$ is a ring. Since all curved facets on a ring are rectangular, case 2 is impossible.) We have $e' \subseteq S_x \cap S_z$ for a protecting sphere S_z consecutive to S_x. If $z \notin E$, by Lemma 3(iii) and Lemma 2(iii), we get $d \geq (1-3\mu) \cdot g(z) \geq c_4\mu(1-3\mu) \cdot g(x)$. If $z \in E$, then $xz \subseteq \partial E$ which implies that E intersects $S_x \cap S_z$. Since $S_x \cap E$ is connected, it contains only one edge in ∂F and that edge is e. Observe that the adjacent edges of e in ∂F lie at the intersections between S_x and protecting spheres consecutive to S_x. It follows that one endpoint of e lies on $S_x \cap S_z$. Since e and e' are disjoint, they are separated by a curved edge on $S_x \cap S_z$ whose endpoints lie on two different facets of $\mathcal{P}$ incident to x. By case 1, we get $d \geq \min\{\mu\sin\phi, c_5\mu^2\sin\phi\} \cdot g(x)$.

Case 3: e and e' lie on boundaries of holes on $\mathcal{B} \cap S_x$. If e and e' lie on the same hole boundary $S_x \cap S_y$ for a protecting sphere S_y consecutive to S_x, then e and e' are separated by a curved edge on $S_x \cap S_y$ whose endpoints lie on two different facets of $\mathcal{P}$ incident to x. By case 1, we get $d \geq \min\{\mu\sin\phi, c_5\mu^2\sin\phi\} \cdot g(x)$. If e and e' lie on the boundaries of two holes $S_x \cap S_y$ and $S_x \cap S_z$ for two protecting spheres S_y and S_z consecutive to S_x, by Lemma 3(iv) and Lemma 2(iii), we get $d \geq c_7\mu^3 \cdot g(y) \geq c_4c_7\mu^4 \cdot g(x)$.

Finally, observe that $c_8\mu^4$ is at most the minimum of $\mu\sin\phi$, $c_5\mu^2\sin\phi$, $c_4\mu(1-3\mu)$ and $c_4c_7\mu^4$. ◻

CLAIM 9 *Let p and q be two points on two orthogonal spheres S and S'. Let η be the shortest geodesic between p and q on $\mathrm{Bd}(S \cup S')$. Then $||p-q|| \geq \mathrm{length}(\eta)/(5\pi)$.*

Proof. Let x and y be the centers of S and S' respectively. Let H be the plane through q, x and y. Let C_x and C_y be the circles $H \cap S$ and $H \cap S'$ respectively.

Case 1: $p \in H$. Consider the case where p and q lie on the same side of xy. Let r be the intersection point of C_x and C_y on the same side of xy as p and q. The length of η is at most the minimum tour length from p to r on C_x and from r to q on C_y which is at most $(||p-r|| + ||q-r||)\pi/2$. Since C_x and C_y intersect at right angle by orthogonality, $\angle prq$ in triangle pqr is at least $\pi/2$. So $\sqrt{2} \cdot ||p-q|| \geq ||p-r|| + ||q-r||$. This implies that $||p-q|| \geq (\sqrt{2}/\pi) \cdot \mathrm{length}(\eta)$. Consider the case where p and q lie on opposite sides of xy. Let r (resp. s) be the intersection point of C_x and C_y on the same side of xy as p (resp. q). Let q' be the point on C_y hit by a ray from q perpendicular to xy. Since p and q' lie on the same side of xy, the previous argument shows that $\sqrt{2} \cdot ||p-q'|| \geq ||p-r|| + ||q'-r||$. Since $||q'-r|| = ||q-s||$ and $||p-q|| \geq ||p-q'||$, we get

$$\sqrt{2} \cdot ||p-q|| \geq ||p-r|| + ||q-s||. \qquad (6)$$

Next, we compare $||p-q||$ with $||r-s||$. Without loss of generality, assume that triangle pqr contains the midpoint of rs. If $\angle prq$ in pqr is at least $\pi/2$, pq is the longest side of pqr and so $||p-q|| \geq ||r-s||/2$. If $\angle prq$ in pqr is less than $\pi/2$, then $\angle pxq$ in triangle pqx is at least $\pi/2$ and so $||p-q|| \geq ||p-x||$. Since px is a radial of C_x, we get $||r-s||/2 \leq ||p-x|| \leq ||p-q||$. In all,

$$||p-q|| \geq ||r-s||/2. \tag{7}$$

The length of η is at most the minimum tour length from p to r on C_x, from r to s on $S \cap S'$ and from s to q on C_y. Thus, $\text{length}(\eta) \leq (||p-r||+||r-s||+||q-s||)\pi/2$. By (6) and (7), we get $\text{length}(\eta) \leq (1+1/\sqrt{2})\pi \cdot ||p-q||$.

Case 2: $p \notin H$. Let p' be the point on C_x closest to p. Let d be the distance from p to H. Note that $d \leq ||p-q||$. The length of η is at most the minimum tour length from p to p' on S and from p' to q on $\text{Bd}(S \cup S')$. The tour length from p to p' is at most $\pi d/2 \leq (\pi/2) \cdot ||p-q||$. By case 1, the tour length from p' to q is at most $(1+1/\sqrt{2})\pi \cdot ||p'-q||$. Using triangle inequality, we get $||p'-q|| \leq ||p-p'|| + ||p-q|| \leq \sqrt{2}d + ||p-q|| \leq (\sqrt{2}+1) \cdot ||p-q||$. Hence, $\text{length}(\eta) \leq (\pi/2) \cdot ||p-q|| + ((\sqrt{2}+1)^2\pi/\sqrt{2}) \cdot ||p-q|| < 5\pi \cdot ||p-q||$. ◻

CLAIM 10 *Let E be an element of $\mathcal{Q}$ on $\mathcal{B} \cap S_x$ for some protecting sphere S_x. Let E' be an element of $\mathcal{Q}$ disjoint from E such that either $E' \subseteq \mathcal{B}$ or E' is a flat facet. The minimum distance between E and E' is at least $c_9\mu^8 \cdot g(x)$ for some constant $c_9 < c_8$.*

Proof. We prove the lemma for $c_9 = c_4c_8^2/(225\sqrt{2}\pi^4)$. Let d be the minimum distance between E and E'.

Case 1: $E' \subseteq \mathcal{B}$. Let $E' \subseteq S_y$ for some protecting sphere S_y (y may be x). If $S_x \neq S_y$ and S_x and S_y are not consecutive, then by Lemma 3(iv), $d \geq c_7\mu^3 \cdot g(x)$ which is larger than $c_9\mu^8 \cdot g(x)$. Otherwise, $S_x = S_y$ or S_x and S_y are orthogonal. Let η be the shortest geodesic between E and E' on $\text{Bd}(S_x \cup S_y)$. For each hole on $\mathcal{B} \cap (S_x \cup S_y)$ crossed by η, we reroute around the hole boundary using the shorter arc. This yields a curve γ between E and E' on $\mathcal{B} \cap (S_x \cup S_y)$. If $S_x = S_y$, clearly $d \geq (2/\pi) \cdot \text{length}(\eta)$, otherwise $d \geq \text{length}(\eta)/(5\pi)$ by Claim 9. Observe that $\text{length}(\gamma) \leq (\pi/2) \cdot \text{length}(\eta)$. So we get

$$d \geq (2/(5\pi^2)) \cdot \text{length}(\gamma). \tag{8}$$

Case 1.1: γ intersects two disjoint boundary elements of some curved facet F on $\mathcal{B} \cap (S_x \cup S_y)$. By Claim 8, $\text{length}(\gamma) \geq c_8\mu^4 \cdot \min\{g(x), g(y)\}$. Since $g(y) \geq c_4\mu \cdot g(x)$ by Lemma 2(iii), $\text{length}(\gamma) \geq c_4c_8\mu^5 \cdot g(x)$. Substituting into (8), we get $d \geq (2c_4c_8\mu^5/(5\pi^2)) \cdot g(x) > c_9\mu^8 \cdot g(x)$.

Caes 1.2: every pair of curved edges that γ intersects consecutively are adjacent. Let e and e' be any such adjacent pair of curved edges. Let $p = \gamma \cap e$ and $q = \gamma \cap e'$. We extend γ by taking a detour on e from p to the closest endpoint of e and back to p. We do the same on e'. This yields a longer curve ψ. (ψ is self-intersecting but this is not a problem.) ψ passes through more than one vertex of $\mathcal{Q}$ on $\mathcal{B}$ since E and E' are disjoint. It follows that ψ passes through two vertices of some curved facet on $\mathcal{B} \cap (S_x \cup S_y)$. Case 1.1 shows that

$$\text{length}(\psi) \geq (2c_4c_8\mu^5/(5\pi^2)) \cdot g(x). \tag{9}$$

It remains to bound $\text{length}(\psi)$. Assume without loss of generality that e and e' bound a curved facet on S_x. Let C and C' be the supporting circles of e and e' respectively. Since e and e' meet at right angle (Lemma 4), $C \cap C'$ consists of two diametral points of C or C', say C'. Let B be the ball centered at p with radius $(c_8\mu^4/3) \cdot g(x)$. If $q \notin B$, then $||p-q|| > (c_8\mu^4/3) \cdot g(x)$. The detour on e or e' has length at most $2\pi \cdot \text{radius}(S_x) \leq 6\pi\mu \cdot g(x)$ which is at most

$$(18\pi/(c_8\mu^3)) \cdot ||p-q||. \tag{10}$$

If $q \in B$, we show in the following that the detour on e or e' has length at most $\sqrt{2}\pi \cdot ||p-q||$ which is smaller than (10). Let u be the common endpoint of e and e'. Let v and v' be the other endpoints of e and e' respectively. Let w be the point on C' diametrally opposite to u. Note that $C \cap C' = \{u, w\}$. Let H be the plane containing C'. Since the center of $B \cap H$ lies on the line containing uw, $B \cap H$ contains u or w. By Claim 8, $||p-v'||$ and $||q-v||$ are at least $c_8\mu^4 \cdot g(x)$ which implies that $v, v' \notin B$. We claim that $u \in B$. Otherwise, $w \in B$ which implies that the two arcs $B \cap C$ and $B \cap C'$ cross at w. Since e and e' cannot meet at w, we have $v \in B$ or $v' \in B$, contradiction. By our claim that $u \in B$, we get $||p-u|| \leq (c_8\mu^4/3) \cdot g(x)$ and $||q-u|| \leq (2c_8\mu^4/3) \cdot g(x)$. Observe that $||p-u|| < ||p-v'||$ and $||q-u|| < ||q-v||$. So both detours on e and e' pass through u. Since $\text{radius}(S_x) \geq c_3\mu \cdot g(x)$ by Lemma 2(ii), x is further from u than p and q. Thus, $\angle puq > \pi/4$ and so $\max\{||p-u||, ||q-u||\} \leq \sqrt{2} \cdot ||p-q||$. It follows that the detour on e or e' has length at most $\sqrt{2}\pi \cdot ||p-q||$.

By (10), we conclude that $\text{length}(\psi) \leq (36\pi/(c_8\mu^3)) \cdot \text{length}(\gamma)$. Substituting into (8) and (9), we get $d \geq (c_4c_8^2\mu^8/(225\pi^4)) \cdot g(x) > c_9\mu^8 \cdot g(x)$.

Case 2: E' is a flat facet. If $S_x \cap E' = \emptyset$, x is disjoint from the facet of $\mathcal{P}$ that contains E', so Lemma 3(iii) implies that $d \geq (1-3\mu) \cdot g(x) > c_9\mu^8 \cdot g(x)$.

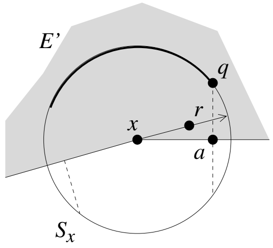

Figure 5: The shaded region is E'. The two dashed line segments delimit the two holes on $\mathcal{B} \cap S_x$ passed through by the boundary edges of E' incident to x. The bold arc is the curved edge $\mathcal{B} \cap S_x \cap E'$. Note that r cannot lie outside the right-angled triangle aqx. Otherwise, the point p, which is above r, would lie outside $\mathcal{B} \cap S_x$.

Otherwise, let $p \in E$ and $r \in E'$ be the points such that $||p - r|| = d$. If r lies on a curved boundary edge e of E', then $E \cap e = \emptyset$ as $E \cap E' = \emptyset$. So we can let $E' = e$ and apply case 1 to finish the analysis. If r lies on a linear boundary edge e of E', then E' lies inside $\mathcal{B}$ and so E' and e are incident to x. Then $||p - r||$ is at least the radius of the hole on $\mathcal{B} \cap S_x$ that e passes through. By Lemma 3(i), $||p - r|| \geq c_5\mu^2 \cdot g(x) > c_9\mu^8 \cdot g(x)$. It remains to consider $r \in \mathrm{int}(E')$. Observe that r is the orthogonal projection of p onto E' which implies that E' lies inside $\mathcal{B}$ and E' is incident to x. Since the subset of E' inside S_x is a cone with apex x (the angle of the cone may be greater than π), the ray from x through r reaches a point $q \in S_x \cap E'$. If $q \in \mathcal{B} \cap S_x \cap E'$, we keep it. Otherwise, q lies on a hole on $\mathcal{B} \cap S_x$ and we move q along $S_x \cap E'$ until q reaches $\mathcal{B} \cap S_x \cap E'$. Figure 5 shows the situation. Observe that in either case $\angle pqr \geq \pi/4$. It follows that $||p - q|| \leq \sqrt{2} \cdot ||p - r||$. Since $E \cap E' = \emptyset$, E is disjoint from the curved edge $\mathcal{B} \cap S_x \cap E'$ that contains q. By Case 1, $||p - q|| \geq (c_4 c_8^2 \mu^8/(225\pi^4)) \cdot g(x)$. Hence, $||p - r|| \geq (c_4 c_8^2 \mu^8/(225\sqrt{2}\pi^4)) \cdot g(x) = c_9\mu^8 \cdot g(x)$. ◻

Proof of Lemma 5

We show that $\widehat{f}(p) \geq c_{10}\mu^8 \cdot g(p)$ where $c_{10} = c_9/(2(1+3\mu) + c_9\mu^8)$. The lemma thus follows by setting $\lambda = c_9/9$ which is smaller than c_{10}. Recall that p lies on or outside $\mathcal{B}$. Let B be the ball centered at p with radius $\widehat{f}(p)$. If $B \cap \mathcal{B} = \emptyset$, then B intersects two flat facets of $\mathcal{Q}$ outside $\mathcal{B}$. Since p lies on or outside $\mathcal{B}$, at most one facet of $\mathcal{P}$ contains p. It follows that $\mathrm{radius}(B) \geq g(p)$. Consider the case where $B \cap \mathcal{B} \neq \emptyset$. Assume to the contrary that $\mathrm{radius}(B) < c_{10}\mu^8 \cdot g(p)$. We need two facts.

FACT 1 *Let p and q be two points. If p does not lie on any edge of $\mathcal{P}$, then $g(p) \leq g(q) + ||p - q||$.*

Proof. Let A be the ball centered at p with radius $g(q) + ||p - q||$. So A intersects the two elements of $\mathcal{P}$ that defines $g(q)$. Since p does not lie on any edge of $\mathcal{P}$ (including edge endpoints), at most one facet of $\mathcal{P}$ contains p. Thus, at most one of the elements of $\mathcal{P}$ that intersect A contains p, so $g(p) \leq \mathrm{radius}(A)$. ◻

FACT 2 *If B intersects a protecting sphere S_x, then $\mathrm{radius}(B) < (c_9\mu^8/2) \cdot g(x)$.*

Proof. By Fact 1, $g(p) \leq g(x) + ||p - x||$. Since B intersects S_x, we get $g(p) \leq g(x) + \mathrm{radius}(S_x) + \mathrm{radius}(B) < (1 + 3\mu) \cdot g(x) + c_{10}\mu^8 \cdot g(p)$. Thus, $g(p) < ((1 + 3\mu)/(1 - c_{10}\mu^8)) \cdot g(x)$ which implies that $\mathrm{radius}(B) < (c_{10}(1 + 3\mu)\mu^8/(1 - c_{10}\mu^8)) \cdot g(x)$. One can verify that $c_{10}(1+3\mu)/(1-c_{10}\mu^8) = c_9/2$. ◻

Take two disjoint elements E and E' of $\mathcal{Q}$ intersected by B. For any protecting sphere S_x intersected by B, by Lemma 3(i), the distances between p and the linear edges incident to x are at least $c_5\mu^2 \cdot g(x) > (c_9\mu^8/2) \cdot g(x)$. So neither E nor E' is a linear edge or an endpoint of a linear edge.

If both E and E' are flat facets, since they are disjoint, they lie on different facets of $\mathcal{P}$. Since at most one facet of $\mathcal{P}$ can contain p, we have $g(p) \leq \mathrm{radius}(B)$, contradicting the assumption that $\mathrm{radius}(B) < c_{10}\mu^8 \cdot g(p)$. Without loss of generality, it remains to consider $E \subseteq \mathcal{B} \cap S_x$ for some protecting sphere S_x. By Claim 10, the minimum distance between E and E' is at least $c_9\mu^8 \cdot g(x)$ which is larger than $2 \cdot \mathrm{radius}(B)$ by Fact 2. Thus, B cannot intersect both E and E', contradiction. ◻