HKUST Theoretical Computer Science Center Research Report HKUST-TCSC-2004-04

Quality Meshing for Polyhedra with Small Angles*

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

August 4, 2004

Abstract

We present an algorithm to compute a Delaunay mesh conforming to a polyhedron possibly with small input angles. The radius-edge ratio of most output tetrahedra are bounded by a constant, except possibly those that are provably close to small angles. Furthermore, the mesh is not unnecessarily dense in the sense that the edge lengths are at least a constant fraction of the local feature sizes at the edge endpoints. This algorithm is simple to implement as it eliminates most of the computation of local feature sizes and explicit protective zones. Our experimental results validate that few skinny tetrahedra remain and they lie close to small acute inputs.

1 Introduction

The need for meshing a polyhedral domain with well shaped tetrahedra occurs as an important problem in finite element methods. Mainly there are two approaches known for the problem, *octtree based refinement* [1, 12] and the *Delaunay refinement* [6, 8, 10, 11, 14, 15]. Among these two, often the latter is preferred for its directional independence and better quality meshing in general.

The Delaunay refinement techniques developed with the work of Chew [5] and Ruppert [14] who showed how to mesh a polygonal domain in 2D with quality triangles. Shewchuk extended the technique to polyhedral domains in three dimensions [15]. This extension has two shortcomings. First, it guarantees bounded radius-edge ratio but not bounded aspect ratio. The radius-edge ratio of a tetrahedron is the ratio of its circumradius to its smallest edge length. Bounded radius-edge ratio eliminates all kinds of poor tetrahedra but not *slivers*. Second, the Delaunay refinement works only with input domains where no angle is smaller than $\frac{\pi}{2}$. The *sliver exudation* method of Cheng et al. [3] addressed the sliver problem though for unbounded domains. Chew [6] and Li and Teng [10] proposed non-deterministic point placement to eliminate slivers for bounded domains. Cheng and Dey [2] combined the sliver exudation method with the Delaunay refinement to give a deterministic algorithm for meshing bounded domains with tetrahedra having bounded aspect ratio.

*Research is supported by RGC, Hong Kong, China, under grants HKUST6088/99E, HKUST6190/02E, Army Research Office, USA under grant DAAD19-02-1-0347 and NSF, USA under grant DMS-0310642.

†Department of Computer Science Science, HKUST, Hong Kong. Email: `scheng@cs.ust.hk`

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

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

However, these works did not address the problem of small angles, specifically input angles less than $\frac{\pi}{2}$. Shewchuk [16] addressed the problem of small angles with constrained Delaunay triangulation though without any guarantees about the shape quality of the tetrahedra. The question of handling small angles with Delaunay refinement was left open. Murphy, Mount and Gable [13] and Cohen-Steiner, de Verdière and Yvinec [7] presented algorithms that can triangulate any polyhedral domain with Delaunay tetrahedra though without any quality guarantee. Recently, Cheng and Poon [4] proposed a Delaunay meshing algorithm that can handle small angles with quality guarantees. The output mesh is graded, and the radius-edge ratio of all tetrahedra are bounded by a constant. The constant depends on the smallest input angle in the vicinity of input edges, but it is independent of the input domain elsewhere. Although this algorithm is a significant theoretical progress for dealing with small angles, its practical viability is doubtful. The main reason is that they construct a protecting region (a union of balls) around each input edge explicitly, which is complicated and time intensive. This construction involves computing a feature related quantity called the local gap size for a number of points on the edges, intersections among spheres and input facets, and intersections among spheres. Moreover, the subsequent refinement has to deal with curved edges and spherical surfaces.

In this paper we present a new algorithm for meshing polyhedra which has similar guarantees as that of [4], but it is much simpler to implement. Our algorithm produces a graded Delaunay mesh. The radius-edge ratio of most output tetrahedra are bounded by a constant. Any remaining tetrahedron of poor shape are provably close to some small input angle. In this new algorithm, we get rid of the explicit construction of any protecting region, and limit the computation of local feature sizes at the input vertices only. The elimination of these computations has enabled us to implement the algorithm. We present experimental results to demonstrate that our algorithm indeed handles polyhedra with small angles. The results validate our claim on the quality of the output tetrahedra and show that few skinny tetrahedra remain.

2 Preliminaries

Basic definitions. We need the following definitions most of which have been introduced in earlier works.

Skinny tetrahedra. Let R, ℓ be the circumradius and shortest edge length respectively of a tetrahedron τ. Let $\rho(\tau) = R/\ell$. We say τ is ρ_0-*skinny* or simply *skinny* if $\rho(\tau) > \rho_0$.

Input domain. We assume that the input domain is a polyhedron bounded by a 2-manifold which is the underlying space of a *piecewise linear complex*, or PLC in short. A PLC is a collection of closed *elements* which are vertices, edges and facets that form a complex in the following sense. The intersection of any two elments is either empty or is a collection of lower dimensional elements. We adopt the convention that an edge includes its endpoints, and a facet includes its vertices and edges. Also, we assume that the boundary of each facet is a collection of simple polygonal cycles.

Incident and adjacent elements. Let ∂F denote the boundary of an element $F \in \mathcal{P}$. We say that an element $F \in \mathcal{P}$ is *incident* to an element $G \in \mathcal{P}$ if either $F \subseteq \partial G$ or $G \subseteq \partial F$. Two elements of $\mathcal{P}$ are *adjacent* if either they are incident or they are non-incident but their intersection is non-empty.

Local feature size. The local feature size for $\mathcal{P}$ is a function $f : \mathbb{R}^3 \mapsto \mathbb{R}$ where $f(x)$ is the radius of the smallest ball centered at x intersecting two non-adjacent elements of $\mathcal{P}$.

Input angles. We define two types of input angles at each vertex u of $\mathcal{P}$. For any two incident edges of u, we measure the angle between them. We call such angles *edge-edge angles*. For any edge uv and a facet F incident to u such that uv is neither incident on F nor coplanar with F, the angle between uv and F is $\min\{\angle puv : p \in F, p \neq u\}$. We call such angles *edge-facet angles*. At an edge of $\mathcal{P}$, we measure the internal and external dihedral angles at the edge.

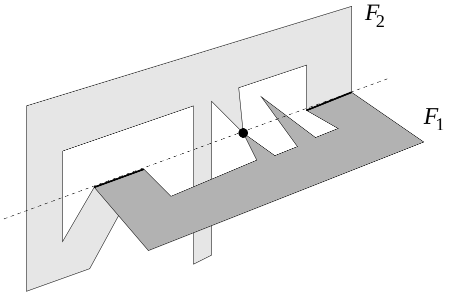

Figure 1: There are two edges and one isolated vertex in $F_1 \cap F_2$.

The remaining angles involve two adjacent non-coplanar facets F_1 and F_2. Let H_1 and H_2 be the support planes of F_1 and F_2, respectively. Notice that F_i may cross the line $H_1 \cap H_2$ and the intersection of F_1 and F_2 may be disconnected; see Figure 1. In general, $F_1 \cap F_2$ consists of a sequence of edges and isolated vertices along the line $H_1 \cap H_2$. The planes H_1 and H_2 divide $\mathbb{R}^3$ into four regions. We call a region a *wedge* of F_1 and F_2 if it contains points on both F_1 and F_2 other than those on the line $H_1 \cap H_2$. The dihedral angle subtended by the wedge at the line $H_1 \cap H_2$ is called its *wedge angle*.

Throughout this paper, we use α_m to denote the minimum input angle in $\mathcal{P}$. We will need the following lemma about the angles between a line segment and a plane that it is attached to.

Lemma 2.1 *Let ab be a line segment. Let H be a plane that contains b but not a. Let a' be the orthogonal projection of a onto H. Then for any point $x \in H$ other than b, $\cos\angle abx = \cos\angle aba' \cdot \cos\angle a'bx$.*

Proof. Let i, j, and k be three orthonormal vectors with origin at b such that i lies on $a'b$, j is normal to H and on the same side of H as ab, and k is on the same side of $a'b$ as bx. Then

$$\begin{aligned}
\cos\angle abx &= \frac{(a-b)}{\|a-b\|} \cdot \frac{(x-b)}{\|x-b\|} \\
&= (i\cos\angle aba' + j\sin\angle aba') \cdot (i\cos\angle a'bx + k\sin\angle a'bx) \\
&= \cos\angle aba' \cdot \cos\angle a'bx.
\end{aligned}$$

□

Lemma 2.1 implies that if we rotate a ray around b in H starting from the direction of $a' - b$, the angle between ab and the ray increases continuously until the ray is collinear with $a'b$. After that the angle decreases continuously until it returns to $a' - b$. See Figure 2.

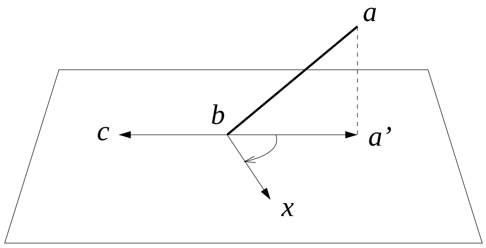

Figure 2: As the ray bx rotates a full circle in clockwise order starting from ba', $\angle abx$ strictly increases as x goes from a' to c, and then strictly decreases as it goes from c to a'.

Sharp features. We call an edge *sharp* if the internal or external dihedral angle at the edge is acute. Sharp edges can be identified by taking the dot product of the outward normals of their incident facets. If the dot product is negative, the edge is sharp.

We call a vertex u *sharp* if an edge-edge angle or an edge-facet angle at u is acute, or if u is an endpoint of a sharp edge. Two edges uv and uw form an acute angle if the dot product $(v-u)\cdot(w-u)$ is positive. Checking if an edge uv forms an acute angle with an incident facet F of u requires a bit more work. We first check that uv is not incident to F and uv is not coplanar with F. Then the angle between uv and F is acute if and only if $\angle vuw$ is acute for some vertex w of F.

We call a wedge *sharp* if its wedge angle is acute. Our algorithm only needs to identify sharp vertices and sharp edges. That is, our algorithm only works with edge-edge angles, edge-facet angles, and dihedral angles. The wedges, wedge angles, and the minimum input angle α_m are used in the analysis only.

3 Algorithm

The algorithm has three distinct phases, an initialization phase called Initialize, a protection phase called Protect and a refinement phase called Refine. In Initialize we protect the sharp vertices and compute the initial Delaunay triangulation. In Protect, we first refine edges and facets so that the the 3D Delaunay triangulation conforms to $\mathcal{P}$. Then, we determine certain protecting balls around sharp edges. Some points are disallowed to be inserted in the Refine phase in these protecting balls and the vertex balls around sharp vertices. Notice that there is no explicit construction of a protecting region. After Protect, Refine starts splitting skinny tetrahedra which may trigger more splitting of the edges and facets. The algorithm maintains a vertex set $\mathcal{V}$ which is updated along with its Delaunay triangulation $\mathrm{Del}\,\mathcal{V}$ as more vertices are generated.

At any generic step of the algorithm, the vertices on any edge of $\mathcal{P}$ divide it into *subsegments*. The subsegments inside the vertex balls are protected by the vertex balls, so we are not concerned about them. We call a subsegment *sharp* if it lies outside the vertex balls and on a sharp edge. It is *non-sharp* otherwise. A *circumball* of a subsegment is a ball with the subsegment endpoints on its boundary. The *diametric ball* of a subsegment is its smallest circumball. A point p *encroaches* a subsegment if p is not an endpoint of the subsegment and lies on or inside its diametric ball. This definition of subsegment encroachment is stricter than the usual one as it includes the case where p lies on the ball boundary. This helps to simplify proofs later.

When no subsegment is encroached by any vertex in $\mathcal{V}$, the facets are decomposed into *subfacets* defined as follows. For each facet F of $\mathcal{P}$, consider the 2D Delaunay triangulation of the vertices in $\mathcal{V} \cap F$. The Delaunay triangles in F are its subfacets. Thanks to the stricter definition of subsegment encroachment, the circumcenters of subfacets lie strictly in the interior of F, see Cheng and Dey [2]. A *circumball* of a subfacet h is a ball with the vertices of h on its boundary. The *diametric ball* of h is its smallest circumball. A point p *encroaches* h if p lies inside its diametric ball. Notice that we cannot use the stricter definition of encroachment as in the subsegment case because more than three vertices in $\mathcal{V} \cap F$ may lie on the same empty circle. In this case, the 2D Delaunay triangulation of $\mathcal{V} \cap F$ is degenerate.

There are several subroutines used by our algorithm. We describe them in Sections 3.1–3.5 and then give the algorithm in Section 3.6.

3.1 Sharp vertex protection

First, we protect the sharp vertices with empty protecting balls called *vertex balls*. Points are disallowed to be inserted inside these vertex balls at later stages. This will mean that certain skinny tetrahedra are not removed since their removal causes insertions of vertices inside these vertex balls. Because of this constraint we compute the feature sizes at the sharp vertices explicitly and use it to compute the vertex balls. This allows us to argue that the skinny tetrahedra that we left out lie near the sharp vertices and edges of the input.

For each sharp vertex u, we compute its distance from all elements of $\mathcal{P}$ which are not incident to u. This distance is the local feature size $f(u)$ at u. We put a ball $\widehat{u}$ with radius $f(u)/4$. (In the rest of this paper, we use $\widehat{x}$ to denote a ball centered at a point x. When x is a vertex of $\mathcal{P}$, $\widehat{x}$ refers to its vertex ball.) The points where the boundary of $\widehat{u}$ intersects edges of $\mathcal{P}$ are inserted into the vertex set $\mathcal{V}$. We protect a subset of $\widehat{u} \cap \mathcal{P}$ using the method of Cohen-Steiner, de Verdière and Yvinec [7]. At any generic step of the algorithm, $\mathcal{V}$ contains vertices on the arc where a facet F incident to u intersects the boundary of $\widehat{u}$. The segments connecting consecutive points on such an arc form *shield subsegments*. Let ab be a shield subsegment. If the angle of the sector aub on F is at least π or ab is encroached, it is split by the following method called SOS according to Cohen-Steiner et al. [7]. If the angle of the sector aub on F is at least π, we insert the midpoint x on the arc between a and b on the boundary of $\widehat{u} \cap F$. The subsegment ab is replaced with two shield subsegments ax and bx; see Figure 3(a). So this type of splitting happens at most once for u, in which case a and b lie on the boundary edges of F incident to u. If the angle of the sector aub is less than π and ab is encroached, we insert the midpoint x on the shorter arc between a and b on the boundary of $\widehat{u}$. The subsegment ab is replaced with two shield subsegments ax and bx; see Figure 3(b).

When no shield subsegment corresponds to a sector at u with angle π or more, the shield subsegments around u create a set of *shield subfacets* incident to u. Figure 3(c) shows an example. It turns out that the diametric ball of a shield subfacet lies in the union of the vertex ball $\widehat{u}$ and the diametric ball of the corresponding shield subsegment. Since $\widehat{u}$ is kept empty throughout the algorithm, it is sufficient to keep the diametric balls of shield subsegments empty to ensure that shield subfacets appear in $\operatorname{Del} \mathcal{V}$.

In Initialize, we only insert the points where the incident edges of u intersects the boundary of $\widehat{u}$, and split shield subsegments that correspond to sectors with angles π or more. The encroachment of shield subsegments is handled in the next phase Protect. The following technical lemma will be useful.

Lemma 3.1 *Let u be a vertex. Let B be the ball centered at u with radius $f(u)/2$. Then*

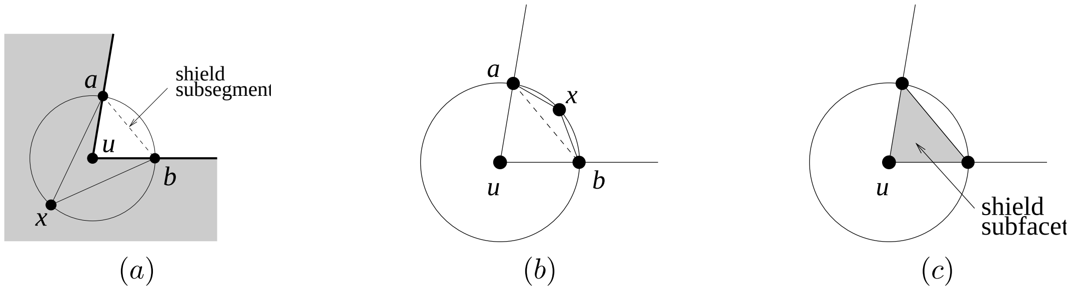

Figure 3: Shield subsegment, shield subfacet, and splitting.

(i) The diametric ball of any shield subsegment inside $\widehat{u}$ lies in B.

(ii) B is disjoint from any vertex ball other than $\widehat{u}$.

(iii) B is disjoint from any edge or facet not incident to u.

Proof. The diametric ball of any shield subsegment inside $\widehat{u}$ has radius at most $f(u)/4$. So any point in it is at distance $f(u)/2$ or less from u. This proves (i). For any vertex ball $\widehat{w} \neq \widehat{u}$, its radius is equal to $f(w)/4 \leq \|u - w\|/4$. Also, radius$(B) = f(u)/2 \leq \|u - w\|/2$. So the distance from B to $\widehat{w}$ is at least $\|u - w\|/4$. This proves (ii). The correctness of (iii) follows from the fact that the distance between u and an non-incident edge or facet is at least $f(u)$. ◻

3.2 Edge splitting

Edges are split in both the Protect and Refine phases using a subroutine SplitEdges, which recovers the edges of $\mathcal{P}$ as union of Delaunay edges. SplitEdges splits any subsegment (sharp, non-sharp or shield) that is encroached using SplitE until no such segment exists.

SplitE(e)

If e is a shield subsegment, split it with SOS else insert the midpoint of e in Del $\mathcal{V}$.

Notice that any point inserted by SplitE cannot encroach the vertex balls of sharp vertices. When SplitEdges terminates, each edge of $\mathcal{P}$ appears as a union of Delaunay edges.

3.3 Facet splitting

After we recover all the edges, we start splitting facets using a subroutine SplitFacets so that they appear in Del $\mathcal{V}$ as union of subfacets. Standard Delaunay refinement insists that no encroached subfacet exists. With small input angles, such a condition may never be satisfied for all subfacets. Instead, we check only that the subfacets appear in Del $\mathcal{V}$. We argue that such a condition can be satisfied for a polyhedron after sufficient but finite amount of splitting.

Let h be a subfacet that does not appear in Del $\mathcal{V}$. Certainly, h cannot be a shield subfacet since there is no encroached shield subsegment when the algorithm reaches the facet splitting step. We basically insert the circumcenter of h if $h \notin \mathrm{Del}\,\mathcal{V}$, but some exception handling is needed if the circumcenter encroaches a subsegment. The procedure SPLITF below gives the details.

SPLITF(h)

(i) Compute the circumcenter c of h;

(ii) If c does not encroach any subsegment, insert c into $\mathrm{Del}\,\mathcal{V}$. Otherwise, let F be the facet containing h.

  (a) pick a subsegment g encroached by c with preference for those in ∂F or on F (shield subsegment), and

  (b) call SPLITE(g).

3.4 Ball splitting

In the PROTECT phase, after all the edges and facets are recovered, further splittings are done using the subroutine SPLITBALLS to reduce the diametric balls of the sharp subsegments roughly to the order of local feature sizes. In order to avoid the computation of local feature sizes, this is achieved in a roundabout way. SPLITBALLS splits any subfacet or subsegment h that is encroached by the midpoint of a sharp subsegment s provided that h and s are contained in disjoint elements of $\mathcal{P}$. The intuition is that some of the new vertices inserted to split these subfacets or subsegments will encroach s and cause s to be split. We claim (Proof of Theorem 4.1) that, at the end of SPLITBALLS, all diametric balls of sharp subsegments become small.

At the end of the PROTECT phase, for each sharp subsegment, we double the radius of its diametric ball with the center fixed and call this a *protecting ball*. These protecting balls and the vertex balls at the sharp vertices constitute the entire set of *protecting balls* for the REFINE phase. Although sharp subsegments may be subdivided further in the REFINE phase, the protecting balls always refer to those computed at the end of the PROTECT phase.

3.5 Refinement of tetrahedra

As in usual Delaunay refinement we attempt to insert the circumcenters of the skinny tetrahedra in this phase. But, some modifications are needed to handle small angles. We disallow the insertion of the circumcenters of skinny tetrahedra inside any protecting ball. The reason is that once these points are allowed to be inserted, they can cause perpetual splittings of the subsegments or subfacets. The tetrahedra incident to the vertices and edges of $\mathcal{P}$ sustaining small enough angles can never satisfy the radius-edge ratio condition. So, at some stage of the algorithm, these tetrahedra should not be split. This means that some skinny tetrahedra may remain at the end. We prove that all such tetrahedra lie close to sharp vertices or edges.

3.6 QMESH

The following algorithm QMESH triangulates $\mathcal{P}$ with the claimed theoretical guarantees. QMESH uses the following subroutine ENCROACH to test whether a point c can be inserted and if not, return the appropriate point to be inserted.

ENCROACH(c)

- If c does not encroach any subsegment or subfacet, return c.
- If c encroaches some subsegment, return its midpoint.
- If c encroaches some subfacet h, return the circumcenter p of h if p does not encroach any subsegment. Otherwise, return the midpoint of a subsegment encroached by p.

QMESH($\mathcal{P}$)

INITIALIZE. Initialize $\mathcal{V}$ to be the set of vertices of $\mathcal{P}$. Compute the vertex balls. Insert the intersections between their boundaries and the edges of $\mathcal{P}$ into $\mathcal{V}$. If any shield subsegment forms a sector with angle π or more, split it with SOS. Compute Del $\mathcal{V}$.

PROTECT. Repeatedly apply a rule from the following list until no rule is applicable. Rule i is applied if it is applicable and no rule j with $j < i$ is applicable.

RULE 1(SPLITEDGES). If there is an encroached subsegment e, call SPLITE(e).

RULE 2(SPLITFACETS). If there is a subfacet h that does not appear in Del $\mathcal{V}$, call SPLITF(h).

RULE 3(SPLITBALLS). Let s be a sharp subsegment on an edge e. If the midpoint of s encroaches a subsegment or subfacet h, where h and e are contained in disjoint elements of $\mathcal{P}$, split h accordingly using SPLITE(h) or SPLITF(h).

At the end of PROTECT, we double the sizes of the diametric balls of sharp subsegments. These expanded balls and the vertex balls are the *protecting balls*. The sharp subsegments may be split further in the next phase, but the locations and sizes of these protecting balls do not change.

REFINE. Repeatedly apply a rule from the following list until no rule is applicable. Rule i is applied if it is applicable and no rule j with $j < i$ is applicable. The parameter $\rho_0 > 2/(1 - \tan(\pi/8))$ is a constant chosen a priori.

RULE 4 (SPLITEDGES). If there is an encroached subsegment e, call SPLITE(e).

RULE 5 (SPLITFACETS). If there is a subfacet h that does not appear in Del $\mathcal{V}$, call SPLITF(h).

RULE 6 (SPLITTET). Assume that there is a tetrahedron with radius-edge ratio exceeding ρ_0. Let z be its circumcenter. If z does not lie on or inside any protecting ball, then

- compute $p :=$ENCROACH(z), and
- insert p into $\mathcal{V}$.

4 Analysis

In this section, we prove that QMESH terminates with a graded Delaunay mesh. All tetrahedra have bounded radius-edge ratio, except possibly those near sharp vertices and edges (in terms of local feature sizes). In sections 4.1 and 4.2, we analyze some critical distances when subsegments and subfacets are considered for splitting. These results are instrumental to proving that rules 1, 2, 4, and 5 never insert any vertex too close to existing ones. Then, in sections 4.3.1, 4.3.2, and 4.3.3 we lower bound the distances among the vertices created in PROTECT and REFINE. These lower bounds allow us to prove the termination and the mesh quality in Section 4.3.4. Recall that $\alpha_m > 0$ denotes the smallest input angle in $\mathcal{P}$.

First of all, we show that the vertex balls of sharp vertices are kept empty.

Lemma 4.1 *Throughout* QMESH, *no vertex is inserted inside the vertex ball of any sharp vertex.*

Proof. Only rules 2 and 5 may possibly insert a vertex inside a vertex ball $\widehat{u}$. In this case, the vertex inserted is the circumcenter x of a subfacet h. We show that this is impossible by deriving a contradiction.

Let F be the facet containing h. Since no subsegment is encroached when rule 2 or 5 is applied, x lies in the interior of F. The facet F is incident to u; otherwise, F is disjoint from $\widehat{u}$ by Lemma 3.1(iii), contradicting the fact that $x \in F$ lies inside $\widehat{u}$.

Let x lie inside a sector on $\widehat{u} \cap F$ bounded by au and bu, where ab is a shield subsegment. If x lies inside the shield subfacet abu, a vertex of h would have encroached ab, contradicting the fact that no subsegment is encroached. If x does not lie inside abu, x would have been rejected for encroaching ab, contradicting the fact that x is inserted by rule 2 or 5. $\square$

4.1 Critical distances: subsegments

In this section, we study the effects of splitting an encroached subsegment in some situations. Our goal is to show that the distances from the inserted vertex x to other vertices of $\mathcal{P}$ are $\Omega(f(x))$. In Lemma 4.2, we first tackle the case of splitting a shield subsegment, and we study encroachment due to points on the edges of $\mathcal{P}$. We extend this result in Lemma 4.3 to all subsegments, shield and non-shield.

Lemma 4.2 *Let ab be a shield subsegment that lies on a facet F of $\mathcal{P}$. Assume that ab is encroached by either an endpoint of a subsegment or the midpoint of a sharp subsegment, which is different from ab. Let x be the vertex to be inserted to split ab. There is a constant $\lambda_1 > 0$ such that:*

(i) For any element E of $\mathcal{P}$ such that $E \neq F$ and $E \not\subseteq \partial F$, the distance between x and E is at least $\lambda_1 f(x)$.

(ii) For any vertex $v \in \mathcal{V}$, $\|v - x\| \geq \lambda_1 f(x)$.

Proof. We prove the lemma for the constant $\lambda_1 = \sin^2 \alpha_m \sin(\pi/4)/10$. Let u be the sharp vertex whose vertex ball contains ab. By the Lipschitz condition, $f(x) \leq f(u) + \|u - x\|$. Also, $\|u - x\| = f(u)/4$. This gives

$$f(u) \geq \frac{4}{5} f(x). \tag{4.1}$$

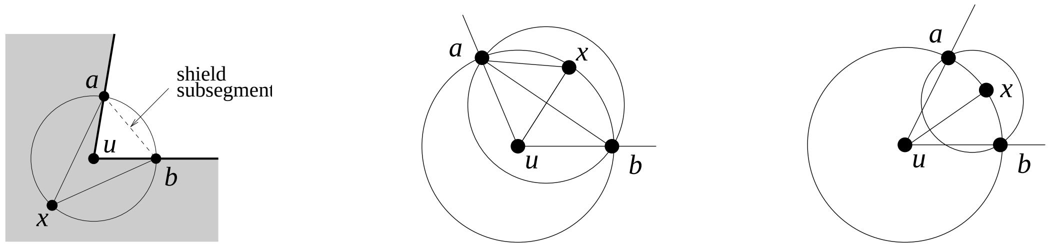

(a) (b) (c)

Figure 4: Splitting a shield subsegment.

We first lower bound $\|a - x\|$ in the case where the sector aub on F has angle less than π and ab is not encroached by u. See Figure 4(c). Let p be the point encroaching ab (subsegment endpoint or midpoint of a sharp subsegment). Let m be the midpoint of ab and let $\widehat{m}$ be the diametric ball of ab. Notice that the sector aub lies on the sector $\widehat{u} \cap F$. By Lemma 3.1(iii), $\widehat{m}$ cannot intersect the edges in ∂F other than the two edges incident to u, which bound $\widehat{u} \cap F$. By Lemma 3.1(i) and (iii), p must lie on some input edge e incident to u. Also, p lies outside $\widehat{u}$: if p is a subsegment endpoint, p does not lie inside $\widehat{u}$ by Lemma 4.1; otherwise, p is the midpoint of sharp subsegment and no sharp subsegment lies inside a vertex ball by definition. We claim that e is not a boundary edge of F. Suppose the contrary. Then p lies on e which bounds the sector $\widehat{u} \cap F$, and p lies outside $\widehat{u}$. But $\widehat{m}$ does not intersect the bounding edges of the sector $\widehat{u} \cap F$ outside $\widehat{u}$. So p cannot lie on or inside $\widehat{m}$, a contradiction. This proves our claim that e is not a boundary edge of F. Also, it cannot be coplanar with F. Otherwise, for $p \in e$ to encroach $\widehat{m}$, e would have to cut through the sector aub which is impossible. It follows that the angle between e and F is an input angle which is at least α_m. Since p lies on or inside $\widehat{m}$, $\|a - m\| \geq \|p - m\|$. Since p does not lie inside $\widehat{u}$, $\|p - u\| \geq f(u)/4$. Thus the distance between p and F is at least $f(u) \sin \alpha_m / 4$, which implies that $\|p - m\| \geq f(u) \sin \alpha_m / 4$. Hence

$$\|a - x\| \geq \|a - m\| \geq \frac{\sin \alpha_m}{4} f(u) \stackrel{(4.1)}{\geq} \frac{\sin \alpha_m}{5} f(x). \tag{4.2}$$

This finishes the proof of the lower bound of $\|a - x\|$ in the case of Figure 4(c).

We proceed to prove (i). Let E be an element of $\mathcal{P}$ such that $E \neq F$ and $E \not\subseteq \partial F$. There are two cases.

Case 1: E is disjoint from F. Then the distance between E and x is at least $f(x)$.

Case 2: E is adjacent to F. We claim that x is not further from au and bu than all the edges of F. First, by Lemma 3.1(iii), we only need to prove the claim for the bounding edges of the sector $\widehat{u} \cap F$. Since the sector aub lies on the sector $\widehat{u} \cap F$, the shortest line segment connecting x to the bounding edges of $\widehat{u} \cap F$ must link to u or intersect au or bu. This proves our claim.

We show that the distance D between x and au is at least $f(x) \sin \alpha_m \sin(\pi/4)/5$. Refer to Figure 4. In the case of Figure 4(a), the angle of the sector aub is π or more, so $D \geq \|u - x\| = f(u)/4$ which is at least $f(x)/5$ by (4.1). In the cases of Figures 4(b) and (c), the angle of the sector aub is less than π. So the nearest point on au to x lies in the interior of au. In the case of Figure 4(b), u lies on or inside $\widehat{m}$, so $\angle aux \geq \angle aub/2 \geq \pi/4$. Thus $D \geq \|u - x\| \sin(\pi/4) = f(u) \sin(\pi/4)/4$ which is at least $f(x) \sin(\pi/4)/5$ by (4.1).

In the case of Figure 4(c), u does not lie inside $\widehat{m}$, so $\angle xau \geq 3\pi/8 > \pi/4$. Thus $D \geq \|a - x\| \sin(\pi/4)$ which is at least $f(x) \sin \alpha_m \sin(\pi/4)/5$ by (4.2).

If E and F are coplanar, the distance between x and E is at least $D > \lambda_1 f(x)$. If E and F are not coplanar, the angle of any wedge induced by E and F is at least α_m. It follows that the distance between x and E is at least $D \sin \alpha_m > \lambda_1 f(x)$.

Next, we prove (ii). There are two cases.

Case 1: the sector aub has angle π or more, or ab is encroached by u. We have only executed INITIALIZE and possibly Rule 1 in PROTECT so far. Thus all vertices in $\mathcal{V}$ lie on input edges or the intersections between boundaries of vertex balls and facets. By Lemma 3.1(iii), the vertex nearest x is either u or a vertex on an incident edge e of u. In the first case, by (4.1), $\|u - x\| = f(u)/4 \geq f(x)/5 > \lambda_1 f(x)$. Consider the second case. If e is not a boundary edge of F, the distance between x and e is at least $\lambda_1 f(x)$ by (i). If e is a boundary edge of F, we have shown in proving (i) that the distance from x to ∂F is at least the distance from x to au, which is at least $f(x) \sin \alpha_m \sin(\pi/4)/5 > \lambda_1 f(x)$.

Case 2: ab is not encroached by u. Then $\angle aub \leq \pi/2$. Any vertex encroaching $\widehat{m}$ must lie on an element of $\mathcal{P}$; otherwise, such a vertex would have been rejected by QMESH. It follows that any vertex in $\mathcal{V} \backslash F$ are further away than the distance lower bound in (i). Consider the vertices in $\mathcal{V} \cap F$. We claim that $\widehat{m} \cap F$ does not contain any shield subsegment endpoint other than a and b. Clearly, all shield subsegments inside $\widehat{u} \cap F$ other than ab lie outside $\widehat{m}$. Since the shield subsegments outside $\widehat{u}$ lie inside other vertex balls, they are disjoint from $\widehat{m}$ by Lemma 3.1(i) and (ii). This proves our claim. Since $\widehat{m} \cap F$ does not contain any shield subsegment endpoint other than a and b, $\widehat{m}$ does not contain any vertex in the interior of F because such a vertex would have been rejected by QMESH. Thus the distance from x to any vertex in $\mathcal{V} \cap F$ is at least the distance from x to the boundary of $\widehat{m}$. This distance is equal to $\|a - x\|(\sin(\angle axu) - \cos(\angle axu)) \geq \|a - x\|(\sin(3\pi/8) - \cos(3\pi/8)) > \|a - x\|/2$. By (4.2), we conclude that the distance from x to any vertex in $\mathcal{V} \cap F$ is greater than $\lambda_1 f(x)$.

□

Next, we extend the intervertex distance lower bound in Lemma 4.2 to all subsegments, shield and non-shield.

Lemma 4.3 *Let x be the vertex that splits a subsegment ab. Let p be an endpoint of a subsegment or the midpoint of a sharp subsegment, which is different from ab. Suppose that p encroaches ab. Then for all $v \in \mathcal{V}$, $\|v - x\| = \Omega(f(x))$.*

Proof. If ab is a shield subsegment, Lemma 4.2(ii) implies that $\|v - x\| = \Omega(f(x))$ for all $v \in \mathcal{V}$. Suppose that ab is a non-shield subsegment. Let e be the edge containing ab. We first show that $\|p - x\| = \Omega(f(x))$. There are two cases.

Case 1: p lies on an edge e' of $\mathcal{P}$ (including its endpoints). Then $e \neq e'$ in order that p encroaches ab. If e and e' are disjoint, then $\|p - x\| \geq f(x)$. If e and e' are adjacent, they must make an acute angle in order that p encroaches ab. Let u be the sharp vertex $e \cap e'$. As x lies outside $\widehat{u}$, $\|u - x\| \geq f(u)/4$. The Lipschitz condition implies that

$f(x) \leq f(u) + \|u - x\| \leq 5\,\|u - x\|$. Since the angle between e and e' is at least α_m, $\|p - x\| \geq \|u - x\| \sin \alpha_m \geq f(x) \sin \alpha_m / 5$.

Case 2: p lies in the interior of a facet F of $\mathcal{P}$. So p was inserted on the boundary of a vertex ball $\widehat{w}$ to split a shield subsegment on F. In order that p encroaches ab, w cannot be an endpoint of e, which implies that both $f(x)$ and $f(w)$ are at most $\|w - x\|$. Thus $\|p - x\|$ is at least the distance from x to $\widehat{w}$ which is $\|w - x\| - f(w)/4 \geq 3\|w - x\|/4 \geq 3f(x)/4$.

Since p encroaches ab, the above analysis implies that the radius of the diametric ball of ab is $\Omega(f(x))$. Let v be any vertex in $\mathcal{V}$. If v does not lie inside the diametric ball of ab, then clearly $\|v - x\| = \Omega(f(x))$. Otherwise, v encroaches ab. It follows that v must lie on an element of $\mathcal{P}$; otherwise, v would be the circumcenter of some skinny tetrahedron, and QMESH would have rejected v for encroaching ab. Since v lies on an edge or facet, we can apply the same analysis in cases 1 and 2 to show that $\|v - x\| = \Omega(f(x))$. ◻

4.2 Critical distances: subfacets

Our goal is to show that when a non-Delaunay subfacet h splits, the distances from its circumcenter x to the existing vertices and the elements of $\mathcal{P}$ are $\Omega(f(x))$. The major challenge is to show that the circumradius of h is $\Omega(f(x))$. Due to our new strategy of splitting non-Delaunay subfacets, the analysis for the standard Delaunay refinement does not apply.

In this section, we focus on the case where the circumballs of h contain vertices on the elements of $\mathcal{P}$ only, and defer the full analysis until Section 4.3. We first look at some important special cases in Section 4.2.1 and lower bound the distance from x to any encroaching point. In Section 4.2.2, we apply these results to lower bound the circumradius of h. The lower bounds on the distances from x to the existing vertices and the elements of $\mathcal{P}$ then follow. The next lemma on the distance between x and sharp vertices will be useful.

Lemma 4.4 *Let x be the circumcenter of a subfacet on a facet F. For any sharp vertex w of F, $\|w - x\| \geq f(x)/9$.*

Proof. Let h be the subfacet with the circumcenter x. Let C be the circumcircle of h. We claim that for any sharp vertex w of F, $\|w - x\| \geq f(w)/8$. Otherwise, since w does not lie inside C, radius$(C) \leq \|w - x\| < f(w)/8$. Thus the vertices of h lie inside the vertex ball $\widehat{w}$, but this contradicts Lemma 4.1. By the claim, the Lipschitz condition implies that $f(x) \leq f(w) + \|w - x\| < 9\,\|w - x\|$. ◻

4.2.1 Splitting subfacets: special cases

The next three lemmas show that the distance between the circumcenter x of a subfacet from an encroaching point is $\Omega(f(x))$ under some special conditions. These special cases will arise in the analysis in Section 4.2.2 and Section 4.3.

Lemma 4.5 *Assume that no subsegment is encroached by a vertex. Let E and F be two non-coplanar adjacent facets. Let x be the circumcenter of a subfacet on F. Let q be the closest* *point to x on the intersection line of the support planes of E and F. Suppose that one of the the following holds:*

(i) qx intersects an edge in ∂F disjoint from E.

(ii) $q \in E$ and an edge in ∂F intersects qx at x or in its interior.

(iii) q is an isolated vertex in $E \cap F$.

Then for any point $u \in E$, $\|u - x\| = \Omega(f(x))$.

Proof. Consider (i) (Figure 5(a)). Let p be the intersection point of qx and the edge in ∂F. Since this edge is disjoint from E, $\max\{\|p - x\|, \|u - x\|\} \geq f(x)$. If $\|u - x\| \geq f(x)$, we are done. Otherwise, $\|p - x\| \geq f(x)$ and $\|q - x\| \geq \|p - x\| \geq f(x)$. If the wedge of E and F containing u and x is sharp, then $\|u - x\| \geq \|q - x\| \sin \alpha_m \geq f(x) \sin \alpha_m$. If the wedge is non-sharp, then $\angle uqx \geq \pi/2$ and so $\|u - x\| \geq \|q - x\| \geq f(x)$.

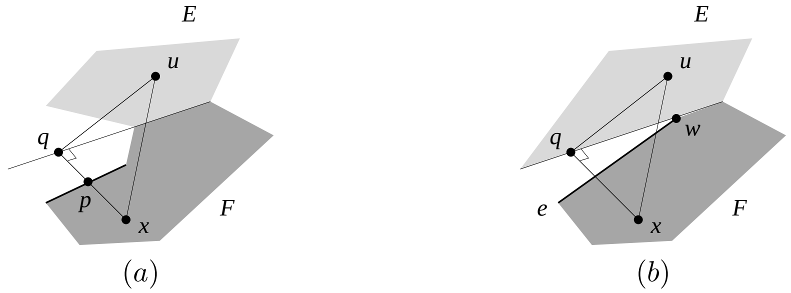

Figure 5:

Consider (ii) (Figure 5(b)). Let e be the edge in ∂F intersecting qx at x or in its interior. We assume that (i) does not apply, so e is incident to a vertex w in $E \cap F$. We show that $\|q - x\| = \Omega(f(x))$, which implies that $\|u - x\| = \Omega(f(x))$ as in case 1. The angle between qw and e is acute. As $q \in E$, this angle is at least an input angle at w. So w is sharp and this angle is at least α_m. Thus $\angle xwq \geq \alpha_m$. By Lemma 4.4, $\|w - x\| \geq f(x)/9$. It follows that $\|q - x\| \geq \|w - x\| \sin \alpha_m \geq f(x) \sin \alpha_m / 9$.

Consider (iii) (Figure 6). We assume that (ii) does not apply. So no edge of F intersects qx at x or in its interior. As x is the circumcenter of a subfacet, x lies in the interior of F. So qx does not lie on an edge of F. Let ℓ be the intersection line of the support planes of E and F. We first show that q is a sharp vertex. There are two cases.

Case 1: E does not lie strictly on one side of ℓ locally at q. So a halfline f on ℓ delimited by q intersects E locally at q. As $qx \subseteq F$ and x lies in the interior of F, the quadrant formed by qx and f on the plane of F must contain an edge e of F. See Figure 6(a). The edge e makes an acute angle with E. So q is sharp.

Case 2: E lies strictly on one side of ℓ locally at q. Let g be the halfline on the support plane of E that is delimited by q and perpendicular to ℓ. If the two edges of E incident to q lie in one of the quadrants formed by ℓ and g, then they make an acute angle at q. See Figure 6(b). So q is sharp. Suppose that the two edges of E incident to q lie in different quadrants formed by ℓ and g. There are two cases.

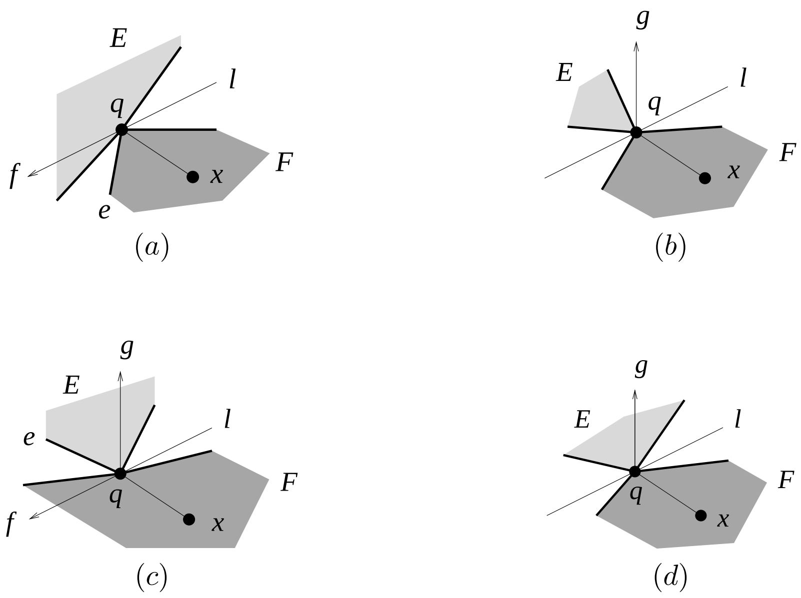

Figure 6:

Case 2.1: F does not lie strictly on one side of ℓ locally at q. A halfline f on ℓ delimited by q intersects F locally at q. The quadrant formed by f and g contains an edge e of E incident to q. See Figure 6(c). Thus e makes an acute angle with F and q is sharp.

Case 2.2: F lies strictly on one side of ℓ locally at q. Then the four edges of E and F incident to q form a convex cone with apex q. See Figure 10(d). The sum of the four face angles of the convex cone is less than 2π, which implies that at least one face angle is acute. So q is sharp.

This finishes the proof of q being sharp. By Lemma 4.4, $\|q - x\| \geq f(x)/9$. If the wedge of E and F containing u and x is sharp, then $\|u - x\| \geq \|q - x\| \sin \alpha_m \geq f(x) \sin \alpha_m/9$. If the wedge is non-sharp, then $\angle uqx \geq \pi/2$ and so $\|u - x\| \geq \|q - x\| \geq f(x)/9$. ◻

Lemma 4.6 *Assume that no subsegment is encroached by a vertex. Let E and F be two non-coplanar adjacent facets. Let x be the circumcenter of a subfacet h on F. Let u be a point on E such that u and x lie on a sharp wedge of E and F. Let q be the closest point to x on the intersection line of the support planes of E and F. Suppose that qu intersects an edge e in ∂E. Also, suppose that e is incident to a vertex w in $E \cap F$ such that $\|w - x\| = \Omega(f(x))$. Then $\|u - x\| = \Omega(f(x))$.*

Proof. Let x' be the orthogonal projection of x onto the support plane of E. Since u and x lie on a sharp wedge of E and F, x' lie on this sharp wedge too. There are two cases.

Suppose that the angle between e and qw is greater than or equal to $\angle x'wq$. See Figure 7(a). By Lemma 2.1, $\angle uwx$ is greater than the angle between e and wx. The angle between e and wx is at least an input angle which is at least α_m. So $\angle uwx \geq \alpha_m$, which implies that $\|u - x\| \geq \|w - x\| \sin \angle uwx = \Omega(f(x))$.

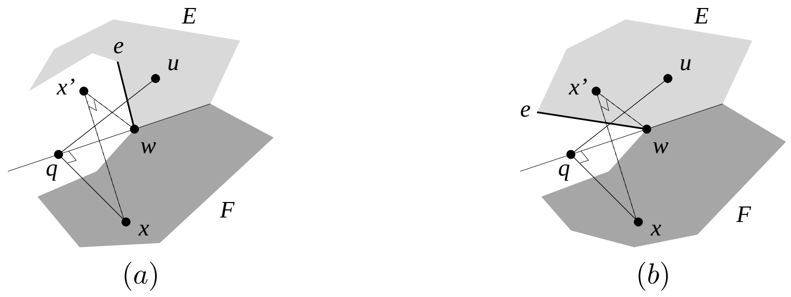

Figure 7:

Suppose that the angle between e and qw is less than $\angle x'wq$. See Figure 7(b). The angle between e and wx is at least α_m. By Lemma 2.1, $\angle qwx$ is greater than the angle between e and wx, which implies that $\angle qwx \geq \alpha_m$. Thus, $\|q - x\| = \|w - x\| \sin \angle qwx = \Omega(f(x))$. Then $\|u - x\| \geq \|x - x'\| \geq \|q - x\| \sin \alpha_m = \Omega(f(x))$. $\square$

Lemma 4.7 *Assume that no subsegment is encroached by a vertex. Let E and F be two non-coplanar adjacent facets. Let h be a subfacet on F with circumcenter x. Suppose that the diametric ball of h contains a point u on E such that u does not encroach any subsegment. Then the following hold:*

(i) F contains the segment joining x and the orthogonal projection of u onto the support plane of F.

(ii) If the wedge of E and F containing u and x is non-sharp, then $\|u - x\| = \Omega(f(x))$.

(iii) Suppose that the wedge of E and F containing u and x is sharp. Let q be the nearest point to x on the intersection line of the support planes of E and F. If $q \notin E \cap F$ or q lies on a non-sharp edge in $E \cap F$, then $\|u - x\| \geq \delta f(x)$ for some constant $\delta > 0$.

Proof. Let B and C be the diametric ball and circumcircle of h, respectively. Let ℓ be the intersection line of the support planes of E and F. Let u' be the orthogonal projection of u onto the support plane of F.

Consider (i). Note that u' lies inside C as u lies inside B. Assume to the contrary that (i) is false. So an edge e in ∂F separates x and u' inside C. By the emptiness of C, e contains a chord ab of C, which lies within a subsegment on e. See Figure 8(a). The point u lies outside the diametric ball of ab as u does not encroach any subsegment by assumption. The bisector of B and the diametric ball of ab is the plane H through ab perpendicular to F. Since ab separates u' and x, H bounds an open halfspace H^- such that $u' \in H^-$ and $x \notin H^-$. Since u does not lie inside the diametric ball of ab but inside B, $u \notin H^-$. But this contradicts the fact that u' is the orthogonal projection of u.

Consider (ii). Since the wedge of E and F containing u and x is non-sharp, u' and x do not lie on the same side of ℓ. No edge in $E \cap F$ intersects C. Otherwise, such an edge crosses C completely (due to the emptiness of C) and as this edge lies on ℓ, it must cross xu', contradicting (i). Since B contains $u \in E$, B must intersect an edge e in ∂E above F such that e is closer to x than u. See Figure 8(b). If e is disjoint from F, then $\|u - x\| \geq f(x)$ as e is closer to x

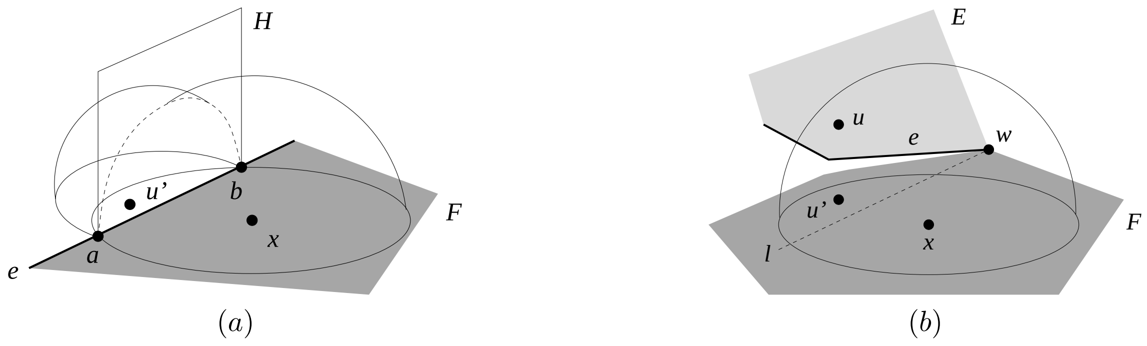

Figure 8:

than u. If e is incident to a vertex w of F, then w lies on ℓ and outside C. Observe that the angle between e and ℓ is acute. Since ℓ intersects $xu' \subseteq F$, this angle is at least an input angle at w. So w is sharp and $\|w - x\| \geq f(x)/9$ by Lemma 4.4. The angle between e and wx is at least α_m. The distance between x and e is at least $\|w - x\| \sin \alpha_m \geq f(x) \sin \alpha_m/9$, and so is $\|u - x\|$.

There are two cases in proving (iii).

Case 1: $qu \not\subseteq E$. Then an edge e in ∂E intersects qu at u or in its interior. Suppose that e is disjoint from F. Let $p = qu \cap e$. By considering the triangle qux in Figure 9(a), we have $\max\{\|u - x\|, \|q - x\|\} \geq \|p - x\| \geq f(x)$. If $\|u - x\| \geq f(x)$, we are done. Otherwise, $\|q - x\| \geq f(x)$ implies that $\|u - x\| \geq \|q - x\| \sin \alpha_m \geq f(x) \sin \alpha_m$.

Suppose that e is adjacent to F. So e is incident to a vertex w in $E \cap F$. We prove that w is sharp. Then by Lemma 4.4, $\|w - x\| \geq f(x)/9$ which allows us to apply Lemma 4.6 to show that $\|u - x\| = \Omega(f(x))$. Let H be the plane through w orthogonal to ℓ.

If e lies on H or e lies on the same side of H as x, let g be the halfline in the intersection of H and the support plane of F such that g is delimited by w and lies on the same side of ℓ as qx. See Figure 9(b). Lemma 2.1 implies that the angle between e and g is acute as it is less than the right angle between qw and g. As e lies on H or e lies on the same side of H as x, the quadrant formed by g and qw contains the orthogonal projection of e onto the support plane of F. Then Lemma 2.1 implies that the angle between e and wx is less than one of the angles that e makes with qw and g. So the angle between e and wx is acute. Since $x \in F$, the angle between e and wx is at least an input angle at w. So w is sharp.

If e and x lie on opposite sides of H, since qu intersects e, u and e lie on the same side of H. See Figure 9(c). The argument in the previous paragraph can be applied again to show that the angle between e and wu' is acute. As $u' \in F$ by (i), we conclude that w is sharp.

Case 2: $qu \subseteq E$. So $q \in E$. If $qx \not\subseteq F$, an edge of F intersects qx at x or in its interior. So $\|u - x\| = \Omega(f(x))$ by Lemma 4.5(ii). Suppose that $qx \subseteq F$. Then $q \in E \cap F$. The assumption of the lemma implies that q lies on a non-sharp edge, say e, in $E \cap F$. The two conditions $qu \subseteq E$ and $qx \subseteq F$ imply that a dihedral angle of e is equal to the angle of the wedge of E and F containing u and x. This is a contradiction as the wedge is sharp

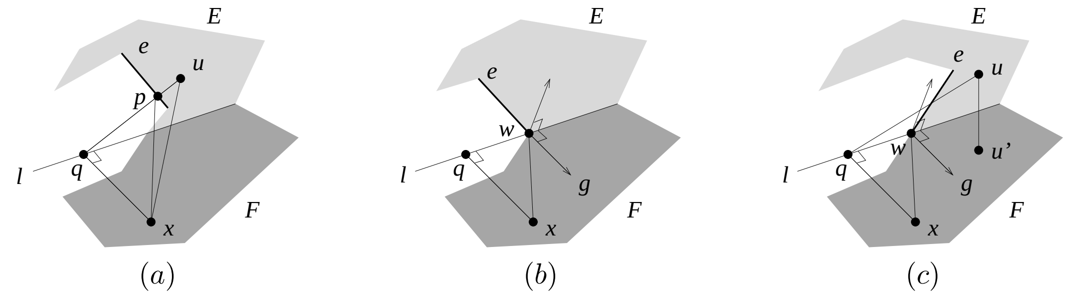

Figure 9:

and e is non-sharp.

$\square$

4.2.2 Splitting subfacets: general cases

We are to show that if a subfacet h is non-Delaunay and its circumballs contain vertices on the elements of $\mathcal{P}$ only, then its circumradius is $\Omega(f(x))$, where x is the circumcenter of h. The analysis consists of three steps. The first step is a technical result (Lemma 4.8) to be used later. Second, we prove in Lemma 4.9 that h has a non-empty circumball with radius $\Omega(f(x))$. Third, we prove in Lemma 4.10 that if h has a large non-empty circumball, its circumradius is $\Omega(f(x))$. This further implies that the distances from x to the elements of $\mathcal{P}$ are $\Omega(f(x))$.

Lemma 4.8 *Assume that no subsegment is encroached by a vertex. Let h be a subfacet on a facet F. Let p be a vertex on an element E of $\mathcal{P}$. If p lies inside some circumball of h, then $E \neq F$, $E \not\subseteq \partial F$, and E and F are non-coplanar.*

Proof. Assume to the contrary that $E = F$, or $E \subseteq \partial F$, or E is coplanar with F. Then p actually lies inside the circumcircle of h. If $E = F$ or $E \subseteq \partial F$, this contradicts the emptiness of the circumcircle of h. If E and F are coplanar, the circumcircle of h must cross a subsegment in ∂F in order to contain p. But then p encroaches this subsegment, a contradiction. $\square$

Lemma 4.9 *Assume that no subsegment is encroached by a vertex. Let x be the circumcenter of a subfacet h such that $h \notin \mathrm{Del}\,\mathcal{V}$. Let F be the facet of $\mathcal{P}$ containing h. Suppose that all vertices inside the diametric ball of h lie on the elements of $\mathcal{P}$. Then there is a circumball B of h that contains vertices $u, v \in \mathcal{V}$ such that:*

(i) *u lies on some element of $\mathcal{P}$.*

(ii) *v and the center of B lie in the same open halfspace bounded by the plane of F, and u lies in the other open halfspace.*

(iii) *If v lies on some element of $\mathcal{P}$, then $\mathrm{radius}(B) \geq \lambda_2 f(x) - D$ where D is the distance between x and the center of B, and $\lambda_2 > 0$ is some constant.*

Proof. Since $h \notin \operatorname{Del} \mathcal{V}$, the diametric ball of h is non-empty. Let H be the support plane of F. Let H^- be an open halfspace bounded by H that contains some vertices inside the diametric ball of h. Let H^+ be the opposite open halfspace.

By assumption, all vertices in H^- inside the diametric ball of h lie on some elements of $\mathcal{P}$. If we move the center of the diametric ball of h in the normal direction of H into H^+, we obtain a system of circumballs of h passing through C. Since $h \notin \operatorname{Del} \mathcal{V}$, one of these circumballs B contain vertices on both sides of H, say $u \in H^-$ and $v \in H^+$. It follows that (i) and (ii) are correct.

Let E_1 be a facet of $\mathcal{P}$ containing u (u may lie on its boundary). Suppose that v lies on an element of $\mathcal{P}$. Then v lies on some facet E_2 of $\mathcal{P}$ (including the boundary of E_2). By Lemma 4.8, neither E_1 nor E_2 is coplanar with F.

Let z be the center of B. If any two of E_1, E_2, and F are disjoint, then $\text{radius}(B) \geq f(z)$. The Lipschitz condition implies that $\text{radius}(B) \geq f(z) \geq f(x) - D$ and we are done. The remaining possibility is that E_1 and E_2 are identical or adjacent, and both are adjacent to F. In the rest of the proof, it suffices to lower bound $\|u - x\|$ or $\|v - x\|$ as $\text{radius}(B) \geq \max\{\|u - x\|, \|v - x\|\} - D$ by triangle inequality. For $i = 1$ and 2, let q_i be the closest point to x on the intersection line between the support planes of E_i and F. There are several cases.

If the wedge of E_1 and F containing u and x is non-sharp. then $\|u - x\| = \Omega(f(x))$ by Lemma 4.7(ii). Suppose that the wedge of E_1 and F containing u and x is sharp.

Case 1: $\|q_i - x\| \geq f(x)\sin(\alpha_m/2)/18$ for i=1 or 2. Without loss of generality assume that $\|q_1 - x\| \geq f(x)\sin(\alpha_m/2)/18$. Since the angle of the wedge of E_1 and F containing u and x is at least α_m, $\|u - x\| \geq \|q_1 - x\| \sin \alpha_m = \Omega(f(x))$.

Case 2: $\|q_i - x\| < f(x)\sin(\alpha_m/2)/18$ for $i = 1$ and 2. If $q_1 \notin E_1 \cap F$ or q_1 lies on a non-sharp edge in $E_1 \cap F$, then $\|u - x\| = \Omega(f(x))$ by Lemma 4.7(iii). If q_1 is an isolated vertex in $E_1 \cap F$, then $\|u - x\| = \Omega(f(x))$ by Lemma 4.5(iii). The remaining case is that q_1 lies on a sharp edge e_1 in $E_1 \cap F$.

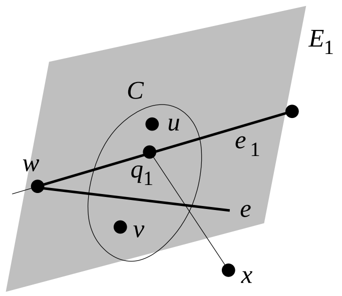

Figure 10:

Case 2.1: $E_1 = E_2$. In order that B contains u and v which are on opposite sides of F, B must intersect the support plane of E_1 in a circle C which contains u, v, and q_1. By the emptiness of the circumcircle of h, e_1 crosses C completely. Since $u, v \in E_1$, the segment uv intersects e_1 as well as another edge e in ∂E_1. If e_1 and e are disjoint, then $\text{radius}(B) \geq f(z)$. The Lipschitz condition implies that $\text{radius}(B) \geq f(z) \geq f(x) - D$ and we are done. Suppose that e_1 and e are adjacent. So e_1 and e share a vertex w in $E_1 \cap F$. Recall that e_1 is sharp and so its endpoint

w is also sharp by definition. See Figure 10. Since $\|q_1 - x\| < f(x)\sin(\alpha_m/2)/18$ by the assumption of case 2, Lemma 4.4 implies that $\|q_1 - x\| < \|w - x\|/2$. By triangle inequality, we get $\|w - q_1\| \geq \|w - x\| - \|q_1 - x\| \geq \|w - x\|/2$, which is at least $f(x)/18$ by Lemma 4.4. The angle between e_1 and e is at least α_m. So $\|q_1 - v\| \geq \|w - q_1\| \sin\alpha_m \geq f(x)\sin\alpha_m/18$. Since the wedge of E_1 and F containing u and x is sharp by assumption, $\angle xq_1v$ is non-acute. It follows that $\|v - x\| \geq \|q_1 - v\| = \Omega(f(x))$.

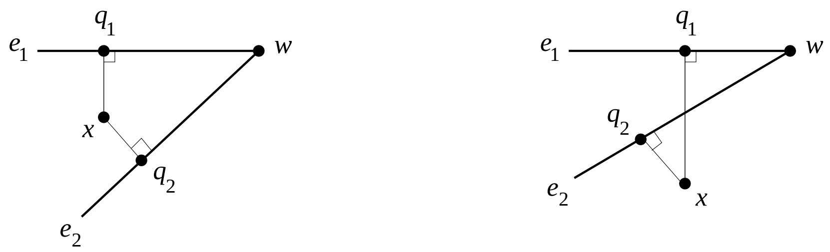

Figure 11: The possible configurations for case 2.2.1.

Case 2.2: $E_1 \neq E_2$. There are two cases depending on whether $q_2 \in E_2$.

Case 2.2.1: $q_2 \in E_2$. If $q_2 \notin E_2 \cap F$, an edge in ∂F must intersect q_2x at x or its interior. So Lemma 4.5(ii) implies that $\|v - x\| = \Omega(f(x))$. Suppose that $q_2 \in E_2 \cap F$. If q_2 is an isolated vertex in $E_2 \cap F$, Lemma 4.5(iii) implies that $\|v - x\| = \Omega(f(x))$. The remaining possibility is that q_2 lies on an edge e_2 in $E_2 \cap F$. Observe that $e_1 \neq e_2$; otherwise, e_1 would be incident to three distinct facets E_1, E_2, and F, violating the fact that $\mathcal{P}$ is a polyhedron. If e_1 and e_2 are disjoint, then $\max\{\|q_1 - x\|, \|q_2 - x\|\} \geq f(x)$, contradicting the assumption of case 2. If e_1 and e_2 share an endpoint, say w, then w is sharp as it is an endpoint of the sharp edge e_1. See Figure 11. The angle between e_1 and e_2 is at least α_m, and wx makes an angle at least $\alpha_m/2$ with e_1 or e_2. But then Lemma 4.4 implies that $\max\{\|q_1 - x\|, \|q_2 - x\|\} \geq \|w - x\| \sin(\alpha_m/2) \geq f(x)\sin(\alpha_m/2)/9$, contradicting the assumption of case 2.

Case 2.2.2: $q_2 \notin E_2$. Then an edge e in ∂E_2 intersects q_2v at v or its interior. If e is disjoint from F, let $p = q_2v \cap e$. By considering the triangle q_2vx, we have $\max\{\|v - x\|, \|q_2 - x\|\} \geq \|p - x\| \geq f(x)$, which implies that $\|v - x\| \geq f(x)$ by the assumption of case 2. Suppose that e is adjacent to F. So e is incident to a vertex w in $E_2 \cap F$. We claim that $\|w - x\| \geq f(x)/9$. If w is an endpoint of e_1, then w is sharp as e_1 is sharp. So $\|w - x\| \geq f(x)/9$ by Lemma 4.4. If w is not an endpoint of e_1, then $\max\{\|q_1 - x\|, \|w - x\|\} \geq f(x)$, which implies that $\|w - x\| \geq f(x)$ by the assumption of case 2. This proves our claim.

If the wedge of E_2 and F containing x and v is sharp, we have satisfied the conditions of Lemma 4.6, which implies that $\|v - x\| = \Omega(f(x))$. Assume that the wedge is non-sharp. Let G be the plane through w that is orthogonal to the intersection line of the support planes of E_2 and F. If x and v do not lie on the same side of G, then $\angle xwv \geq \pi/2$. See Figure 12(a). So $\|v - x\| \geq \|w - x\| \geq f(x)/9$. If x and v lie on the same side of G, e and x also lie on the same side of G as q_2v intersects e. See Figure 12(b). By Lemma 2.1, $\angle xwv$ is at least the angle between e and wx, which is at least α_m. So $\|x - v\| \geq \|w - x\| \sin\alpha_m \geq f(x)\sin\alpha_m/9$.

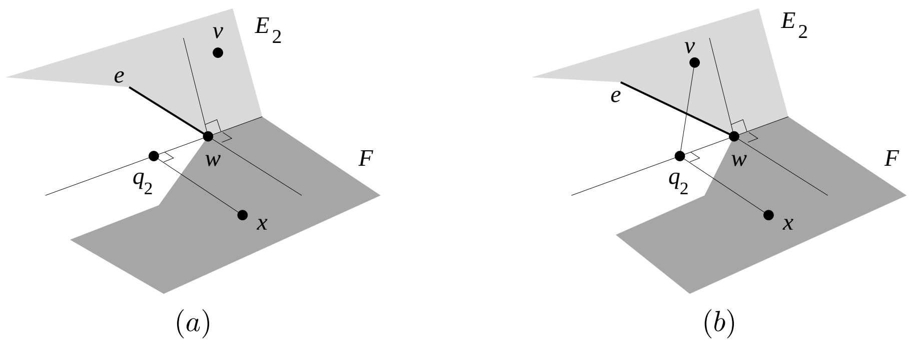

Figure 12:

□

Based on the large circumball guaranteed by the previous lemma, we go on to show that the circumradius of the non-Delaunay subfacet is also large. This also yields a lower bound on the distances from the circumcenter to the elements of $\mathcal{P}$.

Lemma 4.10 *Assume that no subsegment is encroached by a vertex. Let x be the circumcenter of a subfacet h on a facet F of $\mathcal{P}$. Suppose that there is a circumball B of h centered at z and B contains a vertex u such that:*

- *u lies on some element of $\mathcal{P}$,*
- *u and z lie in the different open halfspaces bounded by the plane of F, and*
- $\text{radius}(B) \geq cf(x) - \|x - z\|$ *for some constant c.*

Let $k_1(c) = \min\{\frac{c \tan \alpha_m}{2+\tan \alpha_m}, \delta\}$ and let $k_2(c) = \frac{\sin \alpha_m}{\sqrt{2}} k_1(c)$, where δ is the constant in Lemma 4.7(iii). Then

(i) The circumradius of h is at least $k_1(c)f(x)$.

(ii) Let E be an element of $\mathcal{P}$ such that $E \neq F$ and $E \not\subseteq \partial F$. If x does not encroach upon any subsegment in ∂F, the distance between x and E is at least $k_2(c)f(x)$.

Proof. Let C be the circumcircle of h. Let E_1 be the element of $\mathcal{P}$ containing u. By Lemma 4.8, E_1 is not coplanar with F. Without loss of generality, assume that F is horizontal and u is above F. Let H be the support plane of F. Let $D = \|x - z\|$. Let u' be the orthogonal projection of u onto H. Note that u' lies inside C as u lies inside B, and u and z are in the different halfspaces bounded by H. We first prove (i).

Case 1: E_1 is an edge. If E_1 and F are disjoint, then $\|u - x\| \geq f(x)$. Since u and z lie on opposite sides of F, $\text{radius}(C) \geq \|u - x\| \geq f(x)$ and we are done. If E_1 and F are adjacent, $E_1 \cap F$ is a vertex w. Let L be the plane through u, u', and w. Refer to Figure 13 which shows the cross-section on L.

Since w lies outside C, w lies outside $B \cap L$. We have

$$\tan \phi \leq \frac{r}{D} \leq \frac{\text{radius}(C)}{D}.$$

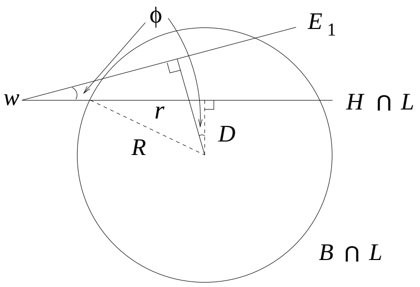

Figure 13: Cross-section on L.

By triangle inequality, $\text{radius}(C) \geq \text{radius}(B) - D$. Since $\text{radius}(B) \geq cf(x) - D$ by assumption, we have

$$\begin{aligned} \text{radius}(C) &\geq cf(x) - 2D \\ &\geq cf(x) - \frac{2\,\text{radius}(C)}{\tan\phi}, \end{aligned}$$

which implies that $\text{radius}(C) \geq \frac{c\tan\phi}{2+\tan\phi} f(x)$. By Lemma 4.7(i), $u' \in F$, so $\phi \geq \alpha_m$ as it is at least an input angle at w. It follows that $\text{radius}(C) \geq \frac{c\tan\alpha_m}{2+\tan\alpha_m} f(x) \geq k_1(c)f(x)$.

Case 2: E_1 is a facet. If E_1 and F are disjoint, then $\|u - x\| \geq f(x)$. Since u and z lie on opposite sides of F, $\text{radius}(C) \geq \|u - x\| \geq f(x)$ and we are done. Assume that E_1 and F are adjacent. The angle of the wedge of E_1 and F containing u and x is at least α_m. Let ℓ be the intersection line of the support planes of E_1 and F. Let q' be the nearest point to u' on ℓ. There are two cases depending on whether q' is inside or outside C.

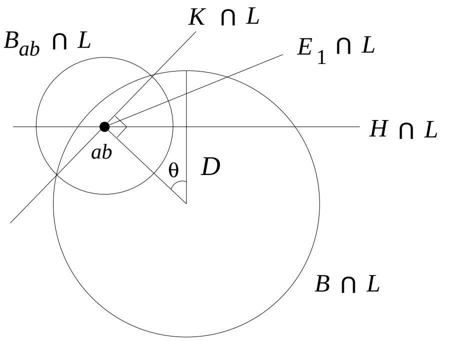

Figure 14: Cross-section on L.

Case 2.1: q' lies outside C. Let L be the plane through q', u, and u'. Then we can apply the analysis in case 1 to the cross-section on L to conclude that $\text{radius}(C) \geq k_1(c)f(x)$.

Case 2.2: q' lies inside C. There are two cases.

Case 2.2.1: $E_1 \cap F$ contains an edge e that stabs C. The emptiness of C implies that e contains a chord ab of C, and ab is part of a subsegment. It also follows that $q' \in ab$. Let B_{ab} be the diametric ball of ab. Let K be the plane that passes through $\partial B_{ab} \cap \partial B$. Let L be the plane through q', u, and u'. Consider the cross-section on L. Refer to Figure 14. Observe that ab lies on K. Since no subsegment is encroached, u lies outside B_{ab}. As u lies inside B by assumption, u must lie on the same side of K as z does. This fact with the condition that $u \in E_1$ implies that the angle of the wedge of E_1 and F containing u and x is at most the acute angle between K and H, which is the angle θ in Figure 14. We have $\tan\theta \leq \frac{\text{radius}(C)}{D}$. The triangle inequality implies that $\text{radius}(C) \geq \text{radius}(B) - D$, which is at least $cf(x) - 2D$ by assumption. Thus the same analysis in case 1 yields $\text{radius}(C) \geq k_1(c)f(x)$.

Case 2.2.2: No edge in $E_1 \cap F$ stabs C. Then the closest point to x on ℓ lies outside $E_1 \cap F$. Lemma 4.7(ii) and (iii) imply that $\|u - x\| \geq \delta f(x)$. Since u and z lie on opposite sides of F, $\text{radius}(C) \geq \|u - x\| \geq \delta f(x)$.

Consider (ii). If E and F are disjoint, the distance between x and E is at least $f(x)$ and we are done. Suppose that E and F are adjacent. First, we claim that the distance between x and ∂F is at least $\text{radius}(C)/\sqrt{2}$. If C does not intersect ∂F, the claim is trivially true. Otherwise, take any subsegment $s \in \partial F$ that intersects C (so s crosses C completely). Since x does not encroach s by assumption, the distance between x and s is at least $\text{radius}(C)/\sqrt{2}$. By our claim, the distance between x and $E \cap F$ is at least $\text{radius}(C)/\sqrt{2}$, which is at least $k_1(c)f(x)/\sqrt{2}$ by (i). The angle of any wedge of E and F is at least α_m. Hence, the distance between x and E is at least $\frac{\sin\alpha_m}{\sqrt{2}}k_1(c)f(x) = k_2(c)f(x)$. ◻

4.3 Termination and mesh quality

In this section, we prove that the inter-vertex distances are above certain thresholds as QMESH inserts vertices. This allows to prove the termination of QMESH by a packing argument. We then show that any remaining skinny tetrahedron is close to some sharp vertex or edge.

The lower bounds on inter-vertex distances in the INITIALIZE and PROTECT phases will be directly related to local feature sizes, while the REFINE phase will be handled using an inductive argument. The inductive argument makes use of a parent-child relation defined as follows.

Let x be a vertex inserted or rejected by QMESH. The *parent* of x is an input vertex or a vertex inserted or rejected by QMESH. If x is a vertex of $\mathcal{P}$ or a vertex inserted during INITIALIZE, its parent is undefined. Otherwise, the parent p is defined as follows.

- If x splits a subsegment (shield or non-shield) or a subfacet, p is the encroaching vertex nearest to x, among the vertices in the current $\mathcal{V}$ as well as any vertex that QMESH tries to insert but is going to reject (for encroaching the subsegment or subfacet split by x).
- If x is the circumcenter of a skinny tetrahedron τ, p is one of the endpoints of the shortest edge of τ. Between the two endpoints of the shortest edge, p is the one that appears in $\mathcal{V}$ later.

It is still possible that x has no parent. This happens when x is inserted to split a subsegment in rule 3. When a subfacet is split by x in rule 3, x may or may not have a parent. If the diametric ball of the subfacet is empty, x does not have any parent. But, if this ball is not empty, x has a parent, namely the vertex in $\mathcal{V}$ nearest to x. In the case where x is the circumcenter of a subfacet h, if the insertion of x is triggered by the circumcenter of a skinny tetrahedron in rule 6, the parent of x is not necessarily the triggering point. This is because the diametric ball of h may already contain some vertices in $\mathcal{V}$.

The *insertion radius* r_x of x is the distance from x to its nearest neighbor in the current $\mathcal{V}$ when x is inserted or rejected. So it follows from definition that if x is an input vertex, then $r_x \geq f(x)$.

In Sections 4.3.1–4.3.3, we lowerbound the insertion radii in each phase, which leads to the proof in Section 4.3.4 that QMESH terminates. We end this preamble with a technical result concerning insertion radius.

Lemma 4.11 *Let x be a vertex inserted or rejected by* QMESH*. Assume that its parent p is defined. If p is a vertex in $\mathcal{V}$, then $r_x = \|p - x\|$. Otherwise, the following hold:*

(i) If x splits a non-shield subsegment, subfacet, or tetrahedron, then r_x is equal to the radius of the diametric ball of the subsegment or subfacet, and $r_x \geq \|p - x\|$.

(ii) If x splits a shield subsegment, then r_x is at least $(1 - \tan(\pi/8))$ times the radius of the diametric ball of the subsegment, and $r_x \geq \frac{1-\tan(\pi/8)}{1+\tan(\pi/8)} \|p - x\|$.

Proof. Let σ be the subsegment, subfacet, or tetrahedron split by x. Let B be the diametric ball of σ. If p is a vertex in $\mathcal{V}$, it follows from definition that $r_x = \|p - x\|$. Otherwise, we know that B does not strictly contain any vertex.

If σ is a non-shield subsegment or subfacet, then x is the center of B and QMESH rejects p for lying inside B. So $r_x = \mathrm{radius}(B) \geq \|p - x\|$. If σ is a tetrahedron, then x is the center of B and p lies on the boundary of B. So $r_x = \mathrm{radius}(B) = \|p - x\|$. Suppose that σ is a shield subsegment. The difference from the previous cases is that x is not the center of B. Let y be the center of B. Let $\widehat{v}$ be the vertex ball that contains σ. The emptiness of B implies that r_x is at least the distance from x to the boundary of B, which is

$$\mathrm{radius}(B) - \|x - y\|. \tag{4.3}$$

Since the parent p is defined, x is not inserted in the INITIALIZE phase. So σ corresponds to a sector at v with angle less than π. Moreover, v cannot lie inside B as B is empty. Thus σ subtends at v an angle $\theta \leq \pi/2$. This implies that

$$\|x - y\| = \mathrm{radius}(B)\tan(\theta/4) \leq \mathrm{radius}(B)\tan(\pi/8). \tag{4.4}$$

Substituting (4.4) into (4.3) yields $r_x \geq (1 - \tan(\pi/8))\,\mathrm{radius}(B)$. Also, (4.4) implies that $\|p - x\| \leq \mathrm{radius}(B) + \|x - y\| \leq (1 + \tan(\pi/8))\,\mathrm{radius}(B)$. It follows that $r_x \geq (1 - \tan(\pi/8))\,\mathrm{radius}(B) \geq \frac{1-\tan(\pi/8)}{1+\tan(\pi/8)} \|p - x\|$. ◻

4.3.1 Initialization

Lemma 4.12 *During* INITIALIZE*, for each vertex $x \in \mathcal{V}$, $r_x \geq \frac{\sin \alpha_m}{5} f(x)$.*

Proof. If x is an input vertex, it follows from definition that $r_x \geq f(x)$. Suppose that x is a vertex inserted by QMESH. During INITIALIZE, all vertices inserted lie on the boundaries of vertex balls. So x lies on $\widehat{v}$ for some sharp vertex v. The Lipschitz condition implies that $f(x) \leq f(v) + \|v - x\| \leq 5f(v)/4$. By Lemma 3.1(ii) and (iii), the nearest neighbor of x is contained in $\widehat{v}$.

Suppose that x is the intersection between an incident edge of v and the boundary of $\widehat{v}$. Then the vertices inserted so far are intersections between edges and boundaries of vertex balls. So either the nearest neighbor of x is v or it lies on another incident edge of v. In the first case, $r_x = f(v)/4 \geq f(x)/5$. In the second case, $r_x \geq f(v) \sin \alpha_m/4 \geq f(x) \sin \alpha_m/5$.

The remaining possibility is that x is inserted to split a shield subsegment ab because ab corresponds to a sector with angle π or more. It is clear that v is the nearest neighbor of x in this case. So $r_x = f(v)/4 \geq f(x)/5$. $\square$

4.3.2 Protection

We first lower bound the insertion radius of a vertex that splits a subsegment.

Lemma 4.13 *Let x be the vertex inserted to split a subsegment during* PROTECT. *Then $r_x = \Omega(f(x))$.*

Proof. Let ab be the subsegment split by x. Let p be the point that encroaches ab and triggers the insertion of x. Notice that during PROTECT, all vertices in $\mathcal{V}$ lie on the elements of $\mathcal{P}$.

If p is a subsegment endpoint or the midpoint of a sharp subsegment, Lemma 4.3 implies that $r_x = \Omega(f(x))$. The remaining possibility is that p is the circumcenter of a subfacet h. QMESH attempts to insert p, but rejects p eventually for encroaching ab. It follows that p is the parent of x and no subsegment is encroached before the attempted insertion of p. Let F be the facet containing h. Let r be the circumradius of h.

Case 1: $h \in \operatorname{Del} \mathcal{V}$. In this case, QMESH tries to insert p in rule 3. Rule 3 applies because the diametric ball of h intersects a sharp edge disjoint from F. So $r \geq f(p)$.

Case 1.1: Neither a nor b lies inside the diametric ball of h. By Lemma 4.11, whether ab is shield or not, r_x is asymptotically lower bounded by the radius of the diametric ball of ab. So $r_x = \Omega(\|a - b\|)$. Since p lies on or inside the diametric ball of ab, we have $\|a - b\|/2 \geq r/\sqrt{2}$. So $r_x = \Omega(r) = \Omega(f(p))$. Also, by Lemma 4.11, $r_x = \Omega(\|p - x\|)$ whether ab is shield or not. Thus the Lipschitz condition implies that $f(x) \leq f(p) + \|p - x\| = O(r_x)$.

Case 1.2: a or b lies inside the diametric ball of h. The emptiness of the circumcircle of h implies that ab does not lie on F (including the boundary of F). Since the diametric ball of h strictly contains a or b, which lies on some element of $\mathcal{P}$, by shifting the ball center slightly, we obtain a circumball B of h where radius$(B) \geq r \geq f(p)$, and B satisfies the conditions of Lemma 4.10. By the definition of SPLITF, p does not encroach any subsegment in ∂F in order to cause the splitting of ab. Thus Lemma 4.10(ii) implies that $\|p - x\| \geq k_2(1) f(p)$. The Lipschitz condition implies that $f(x) \leq f(p) + \|p - x\| \leq \frac{1+k_2(1)}{k_2(1)} \|p - x\|$. By Lemma 4.11, $\|p - x\| = O(r_x)$ whether ab is shield or not. It follows that $f(x) = O(r_x)$.

Case 2: $h \notin \operatorname{Del} \mathcal{V}$. By Lemma 4.9, the conditions of Lemma 4.10 are satisfied.

Case 2.1: ab lies on F (including the boundary of F). The emptiness of the circumcircle of h implies that neither a nor b lies inside the diametric ball of h. By Lemmas 4.9 and 4.10(i), $r \geq k_1(\lambda_2)f(p)$. Then we can use the analysis in case 1.1 to obtain $r_x = \Omega(f(x))$.

Case 2.2: ab does not lie on F (including the boundary of F). By the definition of SPLITF, p does not encroach any subsegment in ∂F in order to cause the splitting of ab. So $\|p - x\| \geq k_2(\lambda_2)f(p)$ by Lemmas 4.9 and 4.10(ii). The Lipschitz condition implies that $f(x) \leq f(p) + \|p - x\| \leq \frac{1+k_2(\lambda_2)}{k_2(\lambda_2)}\|p - x\|$. By Lemma 4.11, $\|p - x\| = O(r_x)$ whether ab is shield or not. It follows that $f(x) = O(r_x)$.

$\square$

We are ready to give the full analysis of the lower bounds on inter-vertex distances at the end of the PROTECT phase. We also show that each point on a sharp subsegment lies inside a protecting ball that is not too small.

Lemma 4.14 *At the end of* PROTECT, *there are constants* $\lambda_3, \lambda_4 > 0$ *such that*

(i) For any vertex $x \in \mathcal{V}$, $r_x \geq \lambda_3 f(x)$.

(ii) For any point z *on a sharp subsegment, a ball centered at* z *with radius* $\lambda_4 f(z)$ *lies inside some protecting ball.*

Proof. We first prove (i). By Lemma 4.12, for any vertex x inherited from the INITIALIZE phase, $r_x \geq (\sin \alpha_m/5)f(x)$. Consider the insertion of a vertex x during PROTECT. If x splits a subsegment, then $r_x = \Omega(f(x))$ by Lemma 4.13. Suppose that x is the circumcenter of a subfacet h on a facet F. Thus no subsegment is encroached before the insertion of x, and so x lies in the interior of F. The vertex x does not encroach any subsegment; otherwise, x would be rejected.

If $h \notin \mathrm{Del}\,\mathcal{V}$, then $r_x \geq k_2(\lambda_2)f(x)$ by Lemmas 4.9 and 4.10. Suppose that $h \in \mathrm{Del}\,\mathcal{V}$. So x is inserted by rule 3. Rule 3 applies because the diametric ball of h intersects a sharp edge disjoint from F. It follows that $r \geq f(x)$, where r is the circumradius of h. There are several cases to consider.

Case 1: x has no parent. Then the diametric ball of h is empty and so $r_x \geq r \geq f(x)$.

Case 2: x has a parent p. So p is a vertex in $\mathcal{V}$ inside the diametric ball of h. Thus $r_x = \|p-x\|$. Recall that all vertices in $\mathcal{V}$ in the PROTECT phase lie on the elements of $\mathcal{P}$. So p lies on some facet E (including the boundary of E). By Lemma 4.8, E is not coplanar with F.

If E is disjoint from F, then $r_x = \|p - x\| \geq f(x)$. Suppose that E and F are adjacent. Let q be the nearest point to x on the intersection line of the support planes of E and F.

Case 2.1: $q \notin E \cap F$. If the wedge of E and F containing p and x is non-sharp, then $r_x = \|p - x\| = \Omega(f(x))$ by Lemma 4.7(ii). Otherwise, $r_x = \|p - x\| = \Omega(f(x))$ by Lemma 4.7(iii).

Case 2.2: $q \in E \cap F$. We claim that $\|q - x\| \geq r/\sqrt{2}$. Let C denote the circumcircle of h. If q does not lie inside C, then $\|q - x\| \geq r$. Otherwise, the emptiness of C implies that q lies on a subsegment s in $E \cap F$, and s crosses C completely. In order that x does not encroach s, the distance from x to s must be at least $r/\sqrt{2}$. This proves our claim. Since the angle of any wedge of E and F is at least α_m, the claim implies that $r_x = \|p - x\| \geq \|q - x\| \sin \alpha_m \geq \frac{\sin \alpha_m}{\sqrt{2}} \cdot r \geq \frac{\sin \alpha_m}{\sqrt{2}} f(x)$.

This finishes the proof of (i).

Consider (ii). Let m be the midpoint of a sharp subsegment ab. For any point $z \in am$, the ball centered at z with radius $\|a - m\|$ lies inside the protecting ball of ab. The Lipschitz condition implies that

$$f(z) \leq f(a) + \|a - z\| \leq f(a) + \|a - m\| \tag{4.5}$$

Next, we relate $f(a)$ and $\|a - m\|$ to $f(z)$. We claim that $\|a - b\| \geq \frac{\lambda_3}{1+\lambda_3} f(a)$. If a was not inserted before b, then $\|a - b\| \geq r_a$ which is at least $\lambda_3 f(a)$ by (i). If a was inserted before b, we have $\|a - b\| \geq \lambda_3 f(b)$ similarly. Then the Lipschitz condition implies that $f(a) \leq f(b) + \|a - b\| \leq \frac{1+\lambda_3}{\lambda_3} \|a - b\|$. Our claim implies that $\|a - m\| \geq \frac{\lambda_3}{2+2\lambda_3} f(a)$. Substituting this into (4.5), we get $f(z) \leq \frac{2+3\lambda_3}{\lambda_3} \|a - m\|$. Since the ball centered at z with radius $\|a - m\|$ lies inside the protecting ball of ab, decreasing its radius to $\frac{\lambda_3}{2+3\lambda_3} f(z)$ keeps it inside. $\square$

4.3.3 Refinement

The main difference of the Refine phase from the Protect phase is that skinny tetrahedra are split in the Refine phase. As in the case of standard Delaunay refinement analysis, we are to analyze the effect of splitting skinny tetrahedra on insertion radii using an inductive argument. However, there are issues that have to be addressed first. Due to the splitting of skinny tetrahedra, some vertices may not lie on the elements of $\mathcal{P}$. This means that Lemma 4.9(iii) does not apply. So we cannot borrow the proof of Lemma 4.14(i) for the Protect phase to conclude that the insertion radii of circumcenters of subfacets are also large in the Refine phase. Furthermore, unlike in the Protect phase, we need to look at circumcenters rejected by Qmesh as well for the subsequent inductive argument to work. We first tackle these issues in the next two lemmas. Notice that when Qmesh inserts or rejects a vertex in the Refine phase, its parent is always defined.

Lemma 4.15 *Let h be a subfacet on a facet F. Let C and x be the circumcircle and circumcenter of h, respectively. Suppose that x projects orthogonally to a point y on a sharp edge in ∂F such that $\|x - y\| \leq \lambda_4 f(x)/(3 + \lambda_4)$, where λ_4 is the constant in Lemma 4.14(ii). Then the following hold:*

(i) $f(y) \geq 3f(x)/(3 + \lambda_4)$.

(ii) $\lambda_4 f(y) \geq 3 \|x - y\|$.

(iii) *If* $h \notin \mathrm{Del}\,\mathcal{V}$, *then* $\mathrm{radius}(C) = \Omega(f(x))$.

Proof. Substituting the inequality $\|x-y\| \le \lambda_4 f(x)/(3+\lambda_4)$ into $f(y) \ge f(x) - \|x-y\|$ proves (i) and (ii). Consider (iii). First, observe that h is encroached by a vertex of $\mathcal{V}$ since otherwise h would have appeared in $\mathrm{Del}\,\mathcal{V}$. All encroaching vertices of h must lie on elements of $\mathcal{P}$ as vertices not lying on $\mathcal{P}$ (circumcenters of tetrahdera) and encroaching a subfacet are rejected in rule 6. Now we have all conditions of Lemma 4.9 in place. So, there is a circumball B of h such that B contains two vertices and one of them, say v, lies on the same side of h as the center of B. If v also lies on some element of $\mathcal{P}$, by Lemmas 4.9(iii) and 4.10(i), $\text{radius}(C) \ge k_1(\lambda_2)f(x)$ and we are done. Suppose that v does not lie on any element of $\mathcal{P}$. So v was inserted as the circumcenter of a skinny tetrahedron. By Lemma 4.14(ii), a protecting ball contains the ball centered at y with radius $\lambda_4 f(y)$. This implies that v could not be inserted inside it, i.e., $\|v-y\| \ge \lambda_4 f(y)$. Thus

$$\begin{aligned}
\|v-x\| &\ge \|v-y\| - \|x-y\| \\
&\overset{\text{(ii)}}{\ge} 2\lambda_4 f(y)/3 \\
&\overset{\text{(i)}}{\ge} 2\lambda_4 f(x)/(3+\lambda_4).
\end{aligned}$$

By triangle inequality, $\text{radius}(B)$ is at least $\|v-x\|$ minus the distance between x and the center of B. So Lemma 4.10(i) applies and yields $\text{radius}(C) \ge k_1(\frac{2\lambda_4}{3+\lambda_4})f(x)$. ∎

Lemma 4.16 *Assume that no subsegment is encroached by a vertex. Let x be the circumcenter of a subfacet inserted or rejected during* Refine*. Let p be the parent of x. If p is a vertex in $\mathcal{V}$ lying on some element of $\mathcal{P}$, then $r_x \ge \lambda_5 f(x)$ for some constant $\lambda_5 > 0$.*

Proof. Let h be the subfacet of which x is the circumcenter. Let F be the facet of $\mathcal{P}$ that contains h. Let F' be a facet containing p (including the boundary of F'). By Lemma 4.8, F and F' are non-coplanar.

If F and F' are disjoint, then $r_x = \|p-x\| \ge f(x)$. Suppose that F and F' are adjacent. If p and x lie on a non-sharp wedge of F and F', then $r_x = \|p-x\| = \Omega(f(x))$ by Lemma 4.7(ii). From now on, we assume that p and x lie on a sharp wedge of F and F'. Let y be the closest point to x on the intersection line of the support planes of F and F'.

Case 1: y lies on an edge in $F \cap F'$ (including the edge endpoints). If the edge is non-sharp, then $\|p-x\| = \Omega(f(x))$ by Lemma 4.7(iii). Suppose that the edge is sharp. Let W be the wedge of F and F' containing p and x. Recall that W is sharp by assumption. Since the angle of W is at least α_m, we have

$$\|p-x\| \ge \|x-y\| \sin \alpha_m. \tag{4.6}$$

Also, since $\angle pyx$ is at least the angle of W by Lemma 2.1, we have

$$\|p-x\| \ge \|p-y\| \sin \angle pyx \ge \|p-y\| \sin \alpha_m. \tag{4.7}$$

There are two cases depending on the length of $\|x-y\|$. Let C be the circumcircle of h. Let λ_4 be the constant in Lemma 4.14(ii).

Case 1.1: $\|x-y\| \geq \text{radius}(C)/2$. If $\|x-y\| \geq \lambda_4 f(x)/(3+\lambda_4)$, then by (4.6), $r_x = \|p-x\| = \Omega(f(x))$ and we are done.

Suppose that $\|x-y\| < \lambda_4 f(x)/(3+\lambda_4)$. Since $\text{radius}(C) \leq 2\,\|x-y\|$ by assumption, Lemma 4.15(ii) implies that $\lambda_4 f(y) \geq \|x-y\| + \text{radius}(C)$. Combining this with Lemma 4.14(ii), we conclude that the diametric ball of h lies inside some protecting ball. QMESH never inserts the circumcenter of a tetrahedron inside any protecting ball. It follows that QMESH attempts to insert x in rule 5 because $h \notin \text{Del}\,\mathcal{V}$. Hence, Lemma 4.15(iii) implies that $\text{radius}(C) = \Omega(f(x))$. So $\|x-y\| \geq \text{radius}(C)/2 = \Omega(f(x))$. Substituting this into (4.6) yields $r_x = \|p-x\| = \Omega(f(x))$.

Case 1.2: $\|x-y\| < \text{radius}(C)/2$. Recall that it is the assumption of case 1 that y lies on an edge in $F \cap F'$. By the emptiness of C, a subsegment on this edge crosses C completely. Clip this subsegment within C to obtain a chord s. Notice that y is the midpoint of s. Since $\|x-y\| < \text{radius}(C)/2$, we have the following two inequalities:

$$\text{length}(s) \;\geq\; \sqrt{3}\,\text{radius}(C), \tag{4.8}$$

$$\text{length}(s) \;\geq\; \|x-y\| + \text{radius}(C). \tag{4.9}$$

Since p does not encroach the subsegment containing s, p lies outside the diametric ball of s. It follows that

$$\|p-y\| \;\geq\; \text{length}(s)/2 \;\overset{(4.8)}{\geq}\; \sqrt{3}\,\text{radius}(C)/2. \tag{4.10}$$

If $\text{radius}(C) \geq 2\lambda_4 f(x)/(3+\lambda_4)$, then (4.10) and (4.7) imply that $r_x = \|p-x\| = \Omega(f(x))$ and we are done.

Suppose that $\text{radius}(C) < 2\lambda_4 f(x)/(3+\lambda_4)$. It follows that $\|x-y\| < \lambda_4 f(x)/(3+\lambda_4)$. The inequality (4.9) implies that if we double the size of the diametric ball of s, it contains the diametric ball of h. Therefore, the diametric ball of h lies inside some protecting ball. As in case 1.1, we conclude that $h \notin \text{Del}\,\mathcal{V}$ and so Lemma 4.15(iii) implies that $\text{radius}(C) = \Omega(f(x))$. Then (4.10) and (4.7) imply that $r_x = \|p-x\| = \Omega(f(x))$.

Case 2: y is an isolated vertex in $F \cap F'$. Then $r_x = \|p-x\| = \Omega(f(x))$ by Lemma 4.5(iii).

Case 3: $y \notin F \cap F'$. If the wedge of F and F' containing p and x is non-sharp, we apply Lemma 4.7(ii); otherwise, we apply Lemma 4.7(iii). In both cases, we conclude that $r_x = \|p-x\| = \Omega(f(x))$.

□

We are ready to apply induction to show that in the REFINE phase, the insertion radius of any vertex inserted or rejected is lower bounded by its local feature size.

Lemma 4.17 *If x exists in $\mathcal{V}$ before the invocation of* REFINE *or x is inserted or rejected by* QMESH *during* REFINE*, then $r_x = \Omega(f(x))$.*

Proof. If x exists in $\mathcal{V}$ before the invocation of REFINE, then $r_x \geq \lambda_3 f(x)$ by Lemma 4.14(i). Assume that x is inserted or rejected during REFINE. If x is inserted or rejected in rule i,

$4 \le i \le 6$, we say that x has type i. We prove a stronger statement by induction: if x has type i, then $r_x \ge f(x)/C_i$ for some constants $C_4 > C_5 > C_6 > 1$. In the induction steps, we develop several inequalities involving C_4, C_5, and C_6 that need to hold for the proof. We set the values of C_4, C_5, and C_6 at the end to satisfy all the inequalities.

Let p be the parent of x. Suppose that $r_x \ge c \cdot r_p$ for some constant c. By Lemma 4.11, if x has type 5 or 6, then $r_x \ge \|p - x\|$ and so

$$\frac{f(x)}{r_x} \le \frac{f(p)}{r_x} + \frac{\|p-x\|}{r_x} \le \frac{f(p)}{r_x} + 1 \le \frac{f(p)}{c \cdot r_p} + 1.$$

Otherwise, x has type 4 and $r_x \ge \frac{1-\tan(\pi/8)}{1+\tan(\pi/8)} \|p - x\|$. So

$$\frac{f(x)}{r_x} \le \frac{f(p)}{r_x} + \frac{\|p-x\|}{r_x} \le \frac{f(p)}{r_x} + \frac{1+\tan(\pi/8)}{1-\tan(\pi/8)} \le \frac{f(p)}{c \cdot r_p} + \frac{1+\tan(\pi/8)}{1-\tan(\pi/8)}.$$

Case 1: x has type 4. If p has type 4, p must be inserted and it is an endpoint of a subsegment (shield or non-shield). By Lemma 4.3, $r_x \ge \lambda_6 f(x)$ for some constant $\lambda_6 > 0$.

Suppose that p has type 5 or 6. Since p encroaches the subsegment split by x, the radius of the diametric ball of this subsegment is at least $r_p/\sqrt{2}$. By Lemma 4.11, $r_x \ge \frac{1-\tan(\pi/8)}{\sqrt{2}} \cdot r_p$. As $C_5 > C_6$, by induction assumption, $r_p \ge f(p)/C_5$ regardless of the type of p. This gives $\frac{f(x)}{r_x} \le \frac{\sqrt{2}}{1-\tan(\pi/8)} \cdot \frac{f(p)}{r_p} + \frac{1+\tan(\pi/8)}{1-\tan(\pi/8)} \le \frac{\sqrt{2}}{1-\tan(\pi/8)} \cdot C_5 + \frac{1+\tan(\pi/8)}{1-\tan(\pi/8)}$.

So the inequalities $\beta_1 C_5 + \beta_2 \le C_4$ and $C_4 \ge 1/\lambda_6$ need to hold, where $\beta_1 = \frac{\sqrt{2}}{1-\tan(\pi/8)}$ and $\beta_2 = \frac{1+\tan(\pi/8)}{1-\tan(\pi/8)}$.

Case 2: x has type 5. If p is a vertex lying on some element of $\mathcal{P}$, Lemma 4.16 implies that $r_x \ge \lambda_5 f(x)$. Otherwise, p must have type 6. Let h be the subfacet split by x. Let F be the facet containing h.

We claim that p is not a vertex in $\mathcal{V}$. Assume to the contrary that p is a vertex in $\mathcal{V}$. This implies that p exists before h is created. The diametric ball of h is contained in the union of the diametric balls of subfacets on F at any time in Refine before the creation of h. It follows that p encroached some subfacet when it was inserted. But then Qmesh should have rejected p, a contradiction.

By the claim, we can infer that Qmesh tries to insert p to split a skinny tetrahedron, but then it rejects p for encroaching h. And the diametric ball B of h is empty. Since p lies inside B, radius$(B) \ge r_p/\sqrt{2}$. Thus $r_x =$ radius$(B) \ge r_p/\sqrt{2}$. Since $r_p \ge f(p)/C_6$ by induction assumption, we have $\frac{f(x)}{r_x} \le \frac{\sqrt{2}f(p)}{r_p} + 1 \le \sqrt{2}C_6 + 1$.

So the inequalities $\sqrt{2}C_6 + 1 \le C_5$ and $C_5 \ge 1/\lambda_5$ need to hold.

Case 3: x has type 6. Let τ be the skinny tetrahedron of which x is the circumcenter. Observe that r_p is at most the shortest edge length of τ. This is because the other endpoint of this edge was already in $\mathcal{V}$ when p was inserted; refer to the definition of p. This implies that $r_x = \|p - x\| \ge \rho_0 r_p$.

If p exists as a vertex in $\mathcal{V}$ before the invocation of Refine, then $r_p \ge \lambda_3 f(p)$ by Lemma 4.14(i). So $r_x \ge \rho_0 \lambda_3 f(p)$. The Lipschitz condition implies that $f(x) \le f(p) + \|p - x\| \le \frac{\rho_0\lambda_3+1}{\rho_0\lambda_3} \cdot r_x$.

Suppose that p is inserted into $\mathcal{V}$ during REFINE. As $C_4 > C_5 > C_6$, by induction assumption, $r_p \geq f(p)/C_4$ regardless of the type of p. This gives $\frac{f(x)}{r_x} \leq \frac{f(p)}{\rho_0 \cdot r_p} + 1 \leq \frac{C_4}{\rho_0} + 1$. So the inequalities $C_4/\rho_0 + 1 \leq C_6$ and $C_6 \geq \frac{\rho_0 \lambda_3 + 1}{\rho_0 \lambda_3}$ need to hold.

We set $C_6 = \max\{\frac{\rho_0+\beta_1+\beta_2}{\rho_0 - \sqrt{2}\beta_1}, \frac{1}{\lambda_5}, \frac{1}{\lambda_6}, \frac{\rho_0\lambda_3+1}{\rho_0\lambda_3}\}$, $C_5 = \sqrt{2}C_6 + 1$, and $C_4 = \sqrt{2}C_5 + 1$. Notice that $C_i > 0$ for $i = 4, 5, 6$ since $\rho_0 > \sqrt{2}\beta_1 = \frac{2}{1-\tan(\pi/8)}$ as chosen by QMESH. □

4.3.4 Finale

Theorem 4.1 QMESH *terminates and returns a Delaunay mesh such that for any two mesh vertices a and b, $\|a - b\| = \Omega(\max\{f(a), f(b)\})$. The vertices of any skinny tetrahedron are at distance $O(f(x))$ from some sharp vertex x or some point x on a sharp edge.*

Proof. For any two vertices $a, b \in \mathcal{V}$ such that a appears in $\mathcal{V}$ not earlier than b, $\|a-b\| \geq \mu f(a)$ for some constant $\mu > 0$ (Lemma 4.17). So $f(b) \leq f(a) + \|a - b\| \leq \frac{1+\mu}{\mu}\|a - b\|$. In all, $\|a - b\| = \Omega(\max\{f(a), f(b)\})$. A packing argument shows that QMESH terminates.

Let τ be a skinny tetrahedron. Let R be the circumradius of τ. Note that the circumcenter q lies inside some protecting ball B. If B is the vertex ball at a sharp v, then $f(v)/4 \geq \|q-v\| \geq R$. So the vertices of τ are at distance $f(v)/4 + R \leq f(v)/2$ from v. If B is not a vertex ball, B is twice as large as the diametric ball B_{ab} of some sharp subsegment ab at the end of PROTECT. Let x be the midpoint of ab. Let X denote $\|a - b\|/2$. It suffices to show that $X \leq \sqrt{2}f(x)$.

Assume to the contrary that $X > \sqrt{2}f(x)$. Then B_{ab} intersects an edge or a facet of $\mathcal{P}$ disjoint from the edge containing ab. Let E be such an element closest to x. Since ab is not encroached at the end of PROTECT, no vertex of E lies inside B_{ab}. It follows that x projects orthogonally onto a point y in the interior of E. Let h be the subsegment or subfacet on E that contains y. Let w be the vertex of h closest to y. Since $\|w - x\| \geq X > \sqrt{2}f(x)$ and $\|x - y\| = f(x)$,

$$2\|x - y\|^2 < \|w - x\|^2. \tag{4.11}$$

Let $\widehat{z}$ be the diametric ball of h. Since w is the closest vertex of h to y, $\angle zyw \geq \pi/2$. This implies that $\|y - z\|^2 \leq \text{radius}(\widehat{z})^2 - \|w - y\|^2$. So

$$\begin{aligned}
\|x - z\|^2 &= \|x - y\|^2 + \|y - z\|^2 \\
&\leq \text{radius}(\widehat{z})^2 + \|x - y\|^2 - \|w - y\|^2 \\
&= \text{radius}(\widehat{z})^2 + 2\,\|x - y\|^2 - \|w - x\|^2,
\end{aligned}$$

which is less than $\text{radius}(\widehat{z})^2$ by (4.11). But then rule 3 (SPLITBALLS) should split h as x lies inside $\widehat{z}$, a contradiction. □

5 Experimental results

We experimented with a preliminary implementation of QMesh. We protected sharp vertices with a vertex ball as demanded by the algorithm. We inserted all points where the edges intersect the vertex balls, but did not protect the shield subsegments with SOS. Apparently this does not cause any trouble for the datasets tested so far. It may very well be true that a more refined analysis can get rid of the SOS from the algorithm.

Figure 15 shows the results. The tables below show that most of the tetrahedra have good radius-edge ratio. We took $\rho_0 = 2.2$ for the experiments. Tetrahedra with radius-edge ratio greater than this threshold are shown in the second column of Figure 15. All of them lie near some small angle region. Third column of the figure shows the triangulation of the input PLC and some highlights to exhibit how the algorithm copes with small input angles.

Model	# input points	# sharp elements
Anchor	28	27
Rail	48	24
Wiper	72	58
SimpleBox	32	30
meshtest	60	35

Model	# points inserted	# tetrahedra with R/ℓ		
		0.6-1.4	1.4-2.2	> 2.2
Anchor	1260	3156	1254	35
Rail	263	610	125	2
Wiper	1121	2074	785	6
SimpleBox	1334	3181	1131	40
meshtest	300	798	286	0

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. *SIAM J. Comput.*, 33 (2003), 69–93.

[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] S.-W. Cheng and S.-H. Poon. Graded conforming tetrahedralization with bounded radius-edge ratio. *Proc. 14th Annu. ACM-SIAM Sympos. Discrete Algorithms* (2003), 295–304.

[5] L. P. Chew. Guaranteed-quality triangular meshes. Report TR-98-983, Comput. Sci. Dept., Cornell Univ., Ithaca, New York, (1989).

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

[7] D. Cohen-Steiner, E. C. de Verdière and M. Yvinec. Conforming Delaunay triangulations in 3D. *Proc. Annu. Sympos. Comput. Geom.* (2002), 199–208.

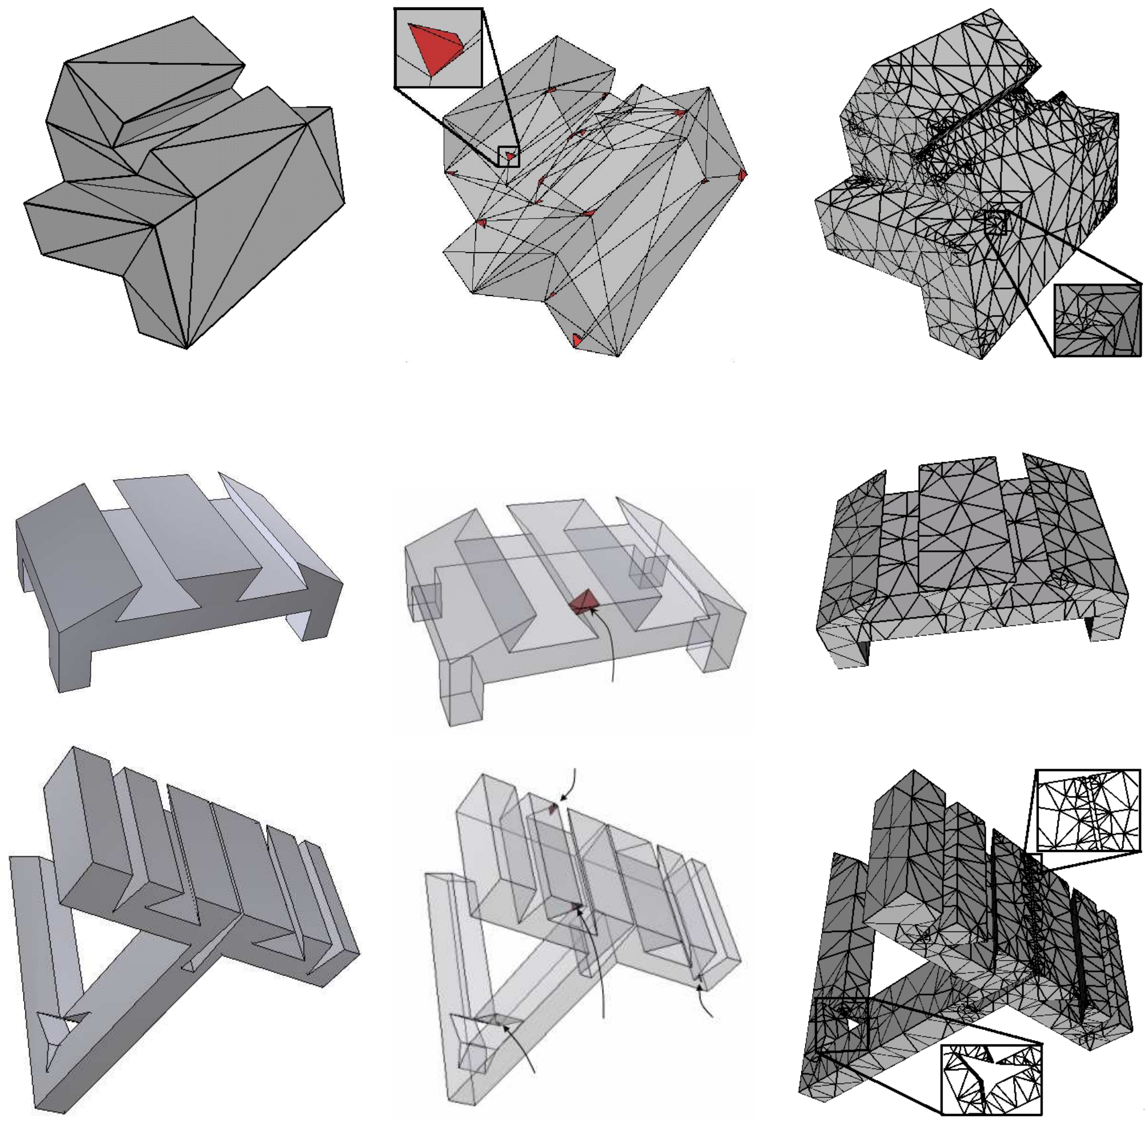

Figure 15: First column shows input PLCs, Anchor, Rail, Wiper from top to bottom. In the second column skinny tetrahedra that remain after Qmesh terminates are pointed out. Third column shows the triangulation of the PLC.

[8] T. K. Dey, C. Bajaj and K. Sugihara. On good triangulations in three dimensions. *Internat. J. Comput. Geom.*, 2 (1992), 75–95.

[9] H. Edelsbrunner. *Geometry and Topology for Mesh Generation.* Cambridge Univ. Press, England, 2001.

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

[11] G. L. Miller, D. Talmor, S.-H. Teng and N. Walkington. A Delaunay based numerical method for three dimensions: generation, formulation, and partition. *Proc. 27th Annu. ACM Sympos. Theory Comput.* (1995), 683–692.

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

[13] M. Murphy, D. M. Mount and C. W. Gable. A point placement strategy for conforming Delaunay tetrahedralization. *Intl. J. Comput. Geom. Appl.*, 11 (2001), 669–682.

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

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

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