Quality Meshing with Weighted Delaunay Refinement*

Siu-Wing Cheng† Tamal K. Dey‡

Abstract

Delaunay meshes with bounded circumradius to shortest edge length ratio have been proposed in the past for quality meshing. The only poor quality tetrahedra called *slivers* that can occur in such a mesh can be eliminated by the *sliver exudation* method. This method has been shown to work for periodic point sets, but not with boundaries. Recently a randomized point-placement strategy has been proposed to remove slivers while conforming to a given boundary. In this paper we present a deterministic algorithm for generating a weighted Delaunay mesh which respects the input boundary and has no poor quality tetrahedron including slivers. This success is achieved by combining the weight pumping method for sliver exudation and the Delaunay refinement method for boundary conformation. We show that an incremental weight pumping can be mixed seamlessly with vertex insertions in our weighted Delaunay refinement paradigm.

1 Introduction.

In finite element methods a 3D domain is often partitioned with tetrahedra. The quality of their shapes influences the quality of the finite element solution [18]. This motivated the decade long research on generating meshes with guaranteed aspect ratio called *quality meshes* [1, 3, 4, 6, 7, 15, 16, 17]. A considerable literature has built up on the subject, see the books [2, 9]. We review only a few of them in the context of the work in this paper.

Bern, Eppstein and Gilbert [3] pioneered a quadtree based triangulation approach for producing quality meshes with close to optimal size in two dimensions. Mitchell and Vavasis [15] extended this technique to triangulate polyhedra with guaranteed aspect ratio in higher dimensions. This line of work provided many crucial insights into the problem though the elements produced by this method have a biased alignment due to the axis parallel boxes used in quadtree/octree subdivisions. Delaunay based triangulations do not have this problem and they are widely used in mesh generation for their uniqueness and many other nice properties, see [2, 9]. As a result researchers also concentrated on computing meshes as a subcomplex of a Delaunay mesh with guaranteed quality. Chew proposed a simple circumcenter insertion method for the problem in 2D [4] which produces an uniform mesh of quality triangles. Ruppert [16], in a pioneering work called Delaunay refinement, showed how circumcenter insertion can be used to produce a quality graded mesh with optimal size.

Many of the concepts including the analysis of the *local feature size* introduced by Ruppert are the basis of the further developments in the area. Shewchuk made an important progress in extending the Delaunay refinement to 3D domains with boundaries [17]. This refinement eliminates tetrahedra that have poor ratio of their circumradius to the shortest edge length henceforth called *radius-edge* ratio. Consequently the resulting mesh satisfies the *ratio property*, i.e., all tetrahedra have radius-edge ratio below a threshold. Although most of the poor quality tetrahedra are removed by the ratio property, one type of bad tetrahedra called *slivers* are not eliminated, see [8]. Slivers are formed by four points placed almost uniformly around the equator of a sphere. Slivers have bounded radius-edge ratio, but they have negligible volume which make them poor quality. Cheng et al. [7] proposed the *sliver exudation* method to get rid of the slivers from a Delaunay mesh that already have the ratio property. They introduced the *pumping* technique that assigns a weight to each vertex which prohibits any sliver to be incident on them in the weighted Delaunay triangulation. Subsequently, Edelsbrunner et al. [10] developed a method to perturb the points so that the unweighted Delaunay tetrahedralization has the ratio property and is sliver-free.

Unfortunately, the algorithms of [7, 10] could not handle boundaries and was applied to periodic point sets. In a recent work Edelsbrunner and Guoy [11] experimented with the sliver exudation method which reveals that the technique is very effective to eliminate slivers though boundary handling remains a challenge. Recently, Li and Teng [12] showed for the first time that it is possible to construct a sliver-free Delaunay mesh, in the presence of boundaries, with a randomized point-placement strategy in line of Chew [6]. The analysis of this algorithm is a nice example of the confluence of the results developed over the years by Chew [6], Ruppert [16], Shewchuk [17] and Edelsbrunner et al. [10].

In this paper we present the first *deterministic* algorithm to construct a Delaunay mesh, with the ratio property and without any sliver, of a 3D domain with boundaries.

*This research has been supported by the Research Grant Council, Hong Kong, China, project no. HKUST8088/99E and NSF under grant CCR-9988216.

†Dept. of Computer Science, HKUST, Clear Water Bay, Hong Kong. e-mail: scheng@cs.ust.hk

‡Dept. of CIS, Ohio State U., Columbus, OH 43210, USA. e-mail: tamaldey@cis.ohio-state.edu

We combine the Delaunay refinement with the sliver exudation in a seamless manner that results in what we call the *weighted Delaunay refinement*. We show that this technique achieves graded meshing and also close to optimal size. Our work can be viewed as another milestone to the line of research initiated by Chew and Ruppert [4, 16], advanced by Dey et al. [8], Shewchuk [17], Cheng et al. [7], Edelsbrunner et al. [10] and Li and Teng [12].

2 Definitions.

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

Piecewise linear complex. The input we consider is a *piecewise linear complex* (PLC) with no input angle less than $\pi/2$ as in previous works [17, 12]. A collection P of vertices, edges and facets in $\mathbb{R}^3$ is called a PLC if (i) all elements on the boundary of an element in P also belong to P, and (ii) if any two elements intersect, then their intersection is a lower dimensional element in P.

Adjacent elements. We call two elements of P *adjacent* if either they are incident or they are non-incident but their closure intersect. For example, two segments sharing an endpoint are adjacent and two facets sharing a segment are adjacent. As another example, if a segment does not lie on the boundary of a facet but they share a vertex, then they are adjacent. A vertex of P is not adjacent to any other element that is not incident to it.

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

Quality measures. Volumes of tetrahedra in a normalized sense capture their quality. Poor quality tetrahedra have small normalized volume. The ratio of circumradius R to the shortest edge length ℓ of a tetrahedron gives a fair indication of its quality. An upper bound ρ_0 on the ratio $\frac{R}{\ell}$ eliminates all tetrahedra of poor quality except slivers. Of course, this demarcation of 'poor' vs. 'good' is dependent on the chosen bound ρ_0. In a sense we parameterize a tetrahedron τ with the ratio $\frac{R}{\ell}(\tau)$ and call it poor if this parameter is above a threshold ρ_0. Since the *radius-edge ratio* does not distinguish slivers as poor, we also parameterize τ by the ratio $\sigma(\tau) = \text{vol}(\tau)/\ell^3$. When $\frac{R}{\ell}(\tau) < \rho_0$ and when $\sigma(\tau)$ is less than a threshold σ_0, we call τ a sliver.

Weighted Delaunay triangulation. We use $\widehat{x}$ to denote a *weighted point* at x with weight X^2. The weighted point $\widehat{x}$ can be interpreted as a sphere centered at x with radius X. Notice that any point can be thought of as a weighted point with weight zero. The *power distance* $\pi(\widehat{x}, \widehat{y})$ between two weighted points $\widehat{x}$ and $\widehat{y}$ is given by

$$\pi(\widehat{x}, \widehat{y}) = \|x - y\|^2 - X^2 - Y^2.$$

If $\pi(\widehat{x}, \widehat{y}) = 0$, then we say that $\widehat{x}$ and $\widehat{y}$ are *orthogonal*. If $\pi(\widehat{x}, \widehat{y})$ is greater (smaller) than 0, then they are *further (closer) than orthogonal* from each other. The *bisector plane* of $\widehat{x}$ and $\widehat{y}$ is the locus of points at equal power distances from $\widehat{x}$ and $\widehat{y}$.

Let τ be a simplex of dimension one or more in $\mathbb{R}^3$, i.e., τ is any of an edge, a triangle or a tetrahedron, where each vertex of τ is weighted. The *smallest orthosphere* of τ is the smallest sphere, say $\widehat{x}$, so that $\widehat{x}$ is orthogonal to each weighted vertex of τ. The smallest orthosphere is the counterpart of the *diametric spheres* for simplices with unweighted vertices; see Figure 1. Notice that for a tetrahedron there is only a single sphere orthogonal to all of its four weighted vertices which is its smallest orthosphere. The center and radius of the smallest orthosphere of any simplex is called its *orthocenter* and *orthoradius* respectively.

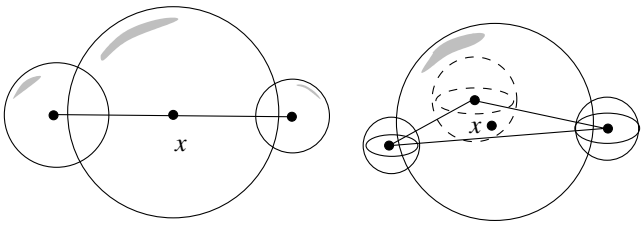

Figure 1: Smallest orthospheres of an edge and a triangle with orthocenter x.

For a weighted point set $\mathcal{V}$, we say a tetrahedron spanning four points of $\mathcal{V}$ is *weighted Delaunay* if its orthosphere is further than orthogonal away from any other weighted point in $\mathcal{V}$. A *weighted Delaunay triangulation* of $\mathcal{V}$ is the collection of all weighted Delaunay tetrahedra along with their triangles, edges and vertices.

3 Weighted Delaunay Refinement.

3.1 Overview. The Delaunay refinement algorithm as originally proposed by Ruppert [16] and later extended to 3D by Shewchuk [17] iteratively inserts circumcenters of poor quality tetrahedra that have radius-edge ratio above a threshold. Whenever a circumcenter x lies so close to an element F in the input PLC that some of its subsets cannot appear in the current Delaunay triangulation, then x is rejected and F is subdivided instead. It can be proved that a new vertex x is inserted at least $c \cdot f(x)$ distance away from all other vertices and input elements for some constant $c > 0$. This lower bound on distances guarantees the termination since only a finite number of vertices can be accommodated in a bounded domain with a lower bound on the interpoint distances.

Delaunay refinement achieves bounded radius-edge ratio but fails to remove slivers. This motivated the sliver exu-

dation method of Cheng et al. [7] which eliminates slivers from a Delaunay mesh that already have bounded radius-edge ratio. The algorithm assigns a weight in the interval $[0, \omega_0^2 N(v)^2]$ to each vertex v where $N(v)$ is the distance to its nearest neighbor and $\omega_0 \in (0, 1/2)$ is a constant. It is shown in [7] that there exists a weight in the mentioned interval for each vertex v so that if v is assigned that weight all slivers incident to v are removed from the weighted Delaunay triangulation of the vertex set. This is stated precisely in the following sliver theorem.

THEOREM 3.1. (SLIVER THEOREM [7].) *Given a periodic point set $\mathcal{V}$ and a Delaunay triangulation of $\mathcal{V}$ with radius-edge ratio $< \rho_0$, there exists $\sigma_0 > 0$ and a weight assignment in $[0, \omega_0^2 N(v)^2]$ for each vertex v in $\mathcal{V}$ such that $\sigma(\tau) \geq \sigma_0$ for each tetrahedron in the weighted Delaunay triangulation of $\mathcal{V}$.*

The required weights can be assigned in a deterministic manner as detailed in [7]. For a bounded domain, the weight assignment may challenge the boundary. We solve this problem by combining the Delaunay refinement with pumping while redefining the encroachment with respect to weighted points. At any generic step we maintain a weighted Delaunay triangulation of the existing vertices $\mathcal{V}$ and refine whenever a boundary element is encroached upon by a weighted vertex due to pumping, or an unweighted point (the orthocenter of a tetrahedron that has poor radius-edge ratio).

3.2 Weight assignment. One difficulty we face with the above strategy is that we do not know the nearest neighbor distance $N(v)$ for a vertex v beforehand as the vertex set changes over time. Therefore, any weight in the interval $[0, \omega_0^2 N(v)^2]$ at the current time may be inappropriate later when $N(v)$ decreases due to new insertions. As a remedy we assign a weight in the interval $[0, \omega_0^2 f(v)^2]$ and show that $N(v)$ remains lower bounded by a constant fraction of $f(v)$ throughout the algorithm. Once a weight is chosen, it will not change.

For pumping to work, we need to establish that following two properties hold throughout the algorithm. We will prove them in section 4.

> VERTEX GAP PROPERTY: For each vertex u, U is at most $\omega_1 < 1/2$ times the Euclidean distance of u to the nearest neighbor in the current vertex set $\mathcal{V}$.
>
> BOUNDARY GAP PROPERTY: For each vertex u, U is at most $\omega_2 < 1$ times the Euclidean distance of u to the nearest segment or facet not containing u for some constant $\omega_2 \leq 2\omega_1$.

3.3 Radius and skinny tetrahedron. One essential observation is that the circumradius of a simplex cannot be much larger than its orthoradius under the vertex gap property. This fact has been proved in [7]. We also provide the proof here for completeness.

LEMMA 3.1. *Let τ be a simplex of two, three, or four vertices. If the vertex gap property holds, then the orthoradius of τ is at least a fraction $\sqrt{1 - 4\omega_1^2}$ of its circumradius.*

Proof. Let $\widehat{x}$ be the orthosphere of τ. Let R denote the circumradius of τ. Let u be the vertex of τ that maximizes $\|u - x\|$. By orthogonality, $X^2 = \|u - x\|^2 - U^2$. By our choice of u, $\|u - x\|$ is at least the circumradius of τ. So $X^2 \geq R^2 - U^2$. The length of each edge of τ is at most $2R$. By the vertex gap property, $U^2 \leq 4\omega_1^2 R^2$. Thus, $X^2 \geq R^2 - 4\omega_1^2 R^2$. □

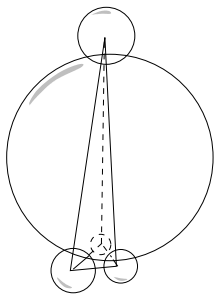

Figure 2: A skinny tetrahedron with a large orthosphere and a short edge.

We mentioned that large radius-edge ratio identifies poor quality tetrahedra other than slivers. Since we will deal with weighted points, we replace circumradius by orthoradius in the above ratio to identify the poor quality tetrahedra. This slackens our quality measure only by a constant factor due to Lemma 3.1. We call a tetrahedron *skinny* if its orthoradius is more than ρ_0 times its shortest edge length, where $\rho_0 > 0$ is a chosen constant. See Figure 2 for an illustration.

3.4 Subsegments and subfacets. The vertices in $\mathcal{V}$ on a segment of P subdivide it into *subsegments*. Let ab be a subsegment. A vertex p *encroaches* upon ab if $\widehat{p}$ is closer than orthogonal from the smallest orthosphere of ab. Note that $\widehat{p}$ may be just the point p. Suppose that there is no encroached subsegment. Consider the weighted vertices in $\mathcal{V}$ on a facet of P in isolation. The 2D weighted Delaunay triangulation of these vertices conforms to the boundary of the facet. Each triangle on the facet is called a *subfacet*. A vertex p *encroaches* upon such a subfacet abc if $\widehat{p}$ is closer than orthogonal from the smallest orthosphere of abc. See Figure 4 for an illustration.

Whenever a subsegment is encroached we split it by inserting its orthocenter. We need that the orthocenter lies

on the subsegment. This is indeed true because of the vertex gap property.

LEMMA 3.2. *If the vertex gap property holds, then the orthocenter of any subsegment ab lies on ab.*

Similarly, if a subfacet is split with its orthocenter we need that the orthocenter lies inside the subfacet.

LEMMA 3.3. *If no subsegment is encroached and the boundary and vertex gap properties hold, then the orthocenter of any subfacet abc lies on the facet containing abc.*

Proof. Let $\widehat{x}$ be the smallest orthosphere of a subfacet in a facet F. Suppose that x lies outside F. Let a be a vertex of this subfacet such that the line segment ax crosses some boundary segment of F. Let s denote the boundary segment crossed.

By the boundary gap property, $\widehat{a}$ does not intersect s. So $\widehat{x}$ must intersect s in order that $\widehat{x}$ intersects $\widehat{a}$. Let $\widehat{y}$ be the diametric sphere of $\widehat{x} \cap s$ with center y. The bisector plane of $\widehat{x}$ and $\widehat{y}$ contains s and is perpendicular to F. Since a and x lie on opposite sides of this bisector plane, $\pi(\widehat{y}, \widehat{a}) < \pi(\widehat{x}, \widehat{a}) = 0$.

Consider the subsegment uv on s that contains y. Observe that $\widehat{y}$ is not closer than orthogonal to $\widehat{u}$ or $\widehat{v}$, otherwise $\widehat{x}$ would also be closer than orthogonal to $\widehat{u}$ or $\widehat{v}$ which is a contradiction. Increase the radius of $\widehat{y}$ until $\widehat{y}$ is orthogonal to $\widehat{u}$ or $\widehat{v}$, say $\widehat{u}$. Let $\widehat{z}$ be the smallest orthosphere of uv. Assume that uv is horizontal as shown in Figure 3. We have three observations. First, $\pi(\widehat{y}, \widehat{a}) < 0$ is preserved. Second, y lies between u and z. Third, the bisector plane H_1 of $\widehat{y}$ and $\widehat{z}$ lies to the left of u, otherwise u and z lie on opposite sides of H_1. This implies that $\pi(\widehat{u}, \widehat{z}) > \pi(\widehat{u}, \widehat{y}) = 0$, a contradiction. If a and y lie on opposite sides of H_2, then

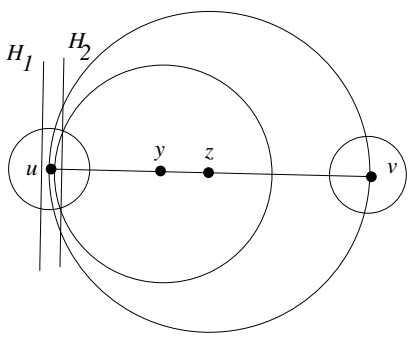

Figure 3: Cross sections of the spheres $\widehat{z}, \widehat{y}$ and $\widehat{u}$ and $\widehat{v}$ in Lemma 3.3 with the plane of F.

$\pi(\widehat{u}, \widehat{a}) < \pi(\widehat{y}, \widehat{a}) < 0$ which contradicts the vertex gap property. On the other hand, if a and y lie on the same side of H_2, then a, y and z lie on the same side of H_1. Thus, $\pi(\widehat{z}, \widehat{a}) \leq \pi(\widehat{y}, \widehat{a}) < 0$. This implies that a encroaches upon uv, a contradiction. □

3.5 Projection. Instead of finding any encroached subfacet, we will focus on one that contains the projection of its encroaching vertex. The next lemma shows the existence of such a subfacet. The lemma is a weighted version of the projection lemma of Shewchuk [17].

LEMMA 3.4. *Suppose that the boundary and vertex gap properties hold. If there is no encroached subsegment and p encroaches upon some subfacet on a facet F, then there exists a subfacet h on F which is encroached upon by p and h contains the orthogonal projection of p on F.*

Proof. Let T denote the 2D weighted Delaunay triangulation of the weighted vertices on F. Note that any subfacet on F is a triangle in T. We view the exterior of T (in the 2D plane containing T) as an infinite triangle whose orthosphere is centered at infinity and has an infinite radius. Let t_1 be a subfacet of F that is encroached upon by p. Suppose that p projects to a triangle t_2. When we walk along the directed segment from the projection of p (inside t_2) to an interior point of t_1, the power distance of p from the smallest orthospheres of edges and triangles of T encountered increases. If t_2 is outside F, then the walk will exit a triangle t outside F and enter a triangle t' inside F at some point, i.e., $t \cap t'$ is a subsegment on the boundary of F. Since p encroaches upon t_1, we have $\pi(\widehat{p}, \widehat{x}) \leq \pi(\widehat{p}, \widehat{y}) < 0$, where $\widehat{x}$ is the smallest orthosphere of $t \cap t'$ and $\widehat{y}$ is the smallest orthospheres of t_1. It follows that p encroaches upon $t \cap t'$ which is a contradiction. □

3.6 Algorithm. The input to our algorithm is a bounded PLC, that is, the underlying space of its outermost boundary is a 2-manifold. This is not a serious restriction since any PLC can be enclosed within a large enough bounded box whose elements are included in the extended PLC. Once the meshing is done conforming to the extended PLC, one can peel away the simplices that do not lie on the actual input PLC. This standard technique has been used in [16], also see [9]. We also assume that no input angle is less than $\frac{\pi}{2}$. This means that two adjacent facets, two adjacent edges, or a facet and an adjacent edge cannot make an angle less than $\frac{\pi}{2}$.

Our algorithm QUALMESH() takes a bounded PLC and outputs a mesh with guaranteed quality tetrahedra.

QUALMESH(P)

1. Compute the Delaunay triangulation of the input vertices of P;
2. Repeatedly apply a rule from the following list until no rule is applicable. Rule i is applied only if it is applicable and no Rule j with $j < i$ is applicable.

RULE 1. If there is an encroached subsegment, then we insert its orthocenter.

RULE 2. If there is an encroached subfacet, then there exists an encroached subfacet h that contains the projection of its encroaching vertex on the facet containing h. We insert the orthocenter of h provided that it does not encroach upon any subsegment. Otherwise, we reject the orthocenter and apply rule 1 to split all subsegments encroached upon by it.

RULE 3. If there is a skinny tetrahedron, then we insert its orthocenter provided that the orthocenter does not encroach upon any subsegment or subfacet. Otherwise, we reject the orthocenter and apply rules 1 and 2 repeatedly to split all subsegments and subfacets encroached upon by the orthocenter.

RULE 4. If there is a sliver, then we claim that it has a vertex v with zero weight. Pump v with weight in between $[0, \omega_0^2 f(v)^2]$ to remove the sliver (ω_0 is determined later in section 4).

When no rule is applicable, the Delaunay tetrahedra lying inside the outer boundary of P form the desired mesh.

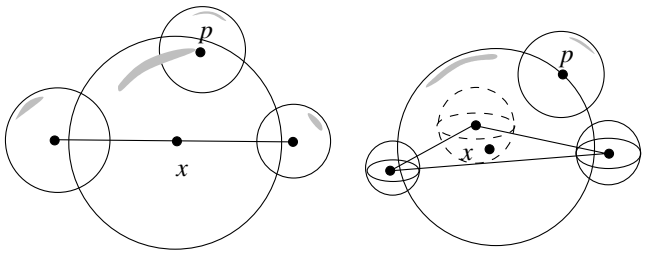

Figure 4: A subsegment is encroached by $\widehat{p}$ on the left and a subfacet is encroached by $\widehat{p}$ on the right. Both are split by their orthocenter x.

4 Analysis of Distances.

The foremost concern for QUALMESH() is its termination. This guarantee follows from the fact that inter-point distances cannot decrease below a threshold. Since only finitely many such vertices can be accommodated in a bounded domain, the algorithm must terminate. In this section we prove this lower bound on distances.

4.1 Adjacent elements and encroachment. A key property in ordinary Delaunay refinement (without pumping) is that vertices on one element of P cannot encroach upon subsegments and subfacets on an adjacent element of P, provided that no input angle is less than $\pi/2$. This property is needed for the algorithm to terminate. We show that this property also holds for weighted Delaunay refinement.

LEMMA 4.1. *If the vertex gap property holds, then for any subsegment ab on an edge e of P, ab cannot be encroached by any vertex on an edge or facet adjacent to e.*

Proof. Let $\widehat{x}$ be the smallest orthosphere of ab. By Lemma 3.2, x lies on ab. Let u be a vertex on an edge or facet adjacent to e. Let v be the common vertex of e and the edge or facet containing u. By the vertex gap property, $\widehat{u}$ does not contain v. Clearly, $\widehat{x}$ does not contain v. $\pi(\widehat{x}, \widehat{u}) = \|u - x\|^2 - X^2 - U^2 > \|u - x\|^2 - \|v - x\|^2 - \|u - v\|^2$. Since all input angles are at least $\pi/2$, $\|u - x\|^2 \geq \|v - x\|^2 + \|u - v\|^2$. So $\pi(\widehat{x}, \widehat{u}) \geq 0$ and u does not encroach upon ab. □

LEMMA 4.2. *Let abc be a subfacet on a facet F of P. If there is no encroached subsegment, then abc cannot be encroached by any vertex on an edge or facet adjacent to F.*

Proof. Assume that F is horizontal. Let T denote the triangulation formed by the subfacets on F. By duality, the vertical plane containing each edge of T is a bisector plane in the power diagram of the smallest orthospheres of edges and triangles in T. (These may not be all the bisector planes in the power diagram though.) Let u be a vertex on an edge or facet adjacent to F. Since no input angle is less than $\pi/2$, the vertical projection of u falls outside or on the boundary of F. It follows that from all smallest orthospheres of edges and triangles in T, the minimum power distance to $\widehat{u}$ is achieved by the smallest orthosphere of some subsegment on the boundary of F. Thus, u cannot encroach upon some subfacet on F unless u encroaches upon some subsegment on the boundary of F. □

4.2 Insertion radius and parent-child relation. We define a notion of *insertion radius* for each vertex inserted or rejected by the algorithm. A lower bound on the insertion radii implies a lower bound for inter-vertex distances. For any vertex x in the input PLC, the insertion radius r_x is the distance from the nearest input vertex. For any vertex x that the algorithm inserts or rejects, the insertion radius r_x is $\min\{\sqrt{\pi(\widehat{u}, x)}\}$ among all vertices u in the current set of vertices $\mathcal{V}$.

The analysis uses a parent-child relation among all vertices that are input vertices or inserted/rejected by the algorithm. This is similar to the parent-child relation defined by Shewchuk [17] for the 3D Delaunay refinement, but we need some modifications due to weighted Delaunay triangulation.

Parents of input vertices are undefined. If the algorithm inserts or rejects a vertex x using Rule i, $1 \leq i \leq 3$, then x has type i. The parent of x is defined as follows.

1. x has type 3. The parent of x is the most recently

inserted endpoint of the shortest edge of the skinny tetrahedron.

2. x has type 1 or 2. There is a vertex that encroaches upon the subsegment or subfacet split by x. The parent of x is the encroaching vertex p such that $\pi(\widehat{p}, x)$ is minimized. This is also equivalent to minimizing $\pi(\widehat{p}, x)$ over all vertices p in the current set of vertices $\mathcal{V}$.

By the above definition, $r_x = \sqrt{\pi(\widehat{p}, x)}$, where p is the parent of x.

4.3 Relation between insertion radii. For each vertex x inserted or rejected by the algorithm, we will relate r_x to the insertion radius of the parent of x in Lemma 4.6. This is the main result of this subsection. First, we need three technical results: Lemma 4.3, Lemma 4.4 and Lemma 4.5.

LEMMA 4.3. *Suppose that the boundary gap property holds. If x has type 1 or 2 and the parent p of x was not rejected, then $r_x \geq \sqrt{1-\omega_2^2}\|x-p\|$.*

Proof. Let $\widehat{x}$ be the smallest orthosphere of the subsegment or subfacet split by x. Since p did not encroach on anything when it was inserted, p encroaches upon $\widehat{x}$ only after pumping. Since p was not rejected, $r_x^2 = \pi(\widehat{p}, x) = \|x-p\|^2 - P^2$. By the boundary gap property, $P \leq \omega_2\|x-p\|$. So $r_x \geq \sqrt{1-\omega_2^2}\|x-p\|$. □

LEMMA 4.4. *Let ax be a line segment on a plane H. Let C_a and C_x be two orthogonal circles on H centered at a and x. For any point p in the strip on H bounded by lines perpendicular to ax and through the endpoints of ax, if p lies outside C_a and C_x, then* $\sin\angle pax \geq \text{radius}(C_x)/\sqrt{2\text{radius}(C_x)^2 + \text{radius}(C_a)^2}$.

Proof. When p is directly above x, $\angle pax$ is minimized. Then $\sin\angle pax = \text{radius}(C_x)/\sqrt{\|a-x\|^2 + \text{radius}(C_x)^2} = \text{radius}(C_x)/\sqrt{2\text{radius}(C_x)^2 + \text{radius}(C_a)^2}$. □

LEMMA 4.5. *Suppose that the vertex gap property holds. Let p be a vertex inserted by the algorithm. Let F be a segment or facet of P not containing p. The distance from p to F is least $c \cdot r_p$, where $c = \sqrt{(1-4\omega_1^2)/2}$.*

Proof. Since p was inserted by the algorithm, p is not an input vertex. Go back to the time when p was inserted. At that time, p was unweighted and p lied outside any orthosphere centered on F (otherwise, p would be rejected).

Consider the case where F is a segment. If the nearest point on F to p is an endpoint u of F, then the minimum distance from p to F is $\|p-u\|$. Since $\pi(\widehat{u}, p) = \|p-u\|^2 - U^2$, the distance from p to F is $\|p-u\| \geq \sqrt{\pi(\widehat{u}, p)} \geq r_p$. If the nearest point on F to p is a point in a subsegment ab, then let $\widehat{x}$ be the smallest orthosphere of ab. Assume that p is closer to a than b. By the vertex gap property and Lemma 3.1, $A \leq 2\omega_1 X/\sqrt{1-4\omega_1^2}$. Apply Lemma 4.4 to the cross-section of $\widehat{a}$ and $\widehat{x}$ in the plane spanned by a, b, and p. We obtain $\sin\angle pax \geq \sqrt{(1-4\omega_1^2)/(2-4\omega_1^2)} \geq \sqrt{(1-4\omega_1^2)/2}$. The distance from p to f is $\|a-p\|\sin\angle pax \geq r_p \sin\angle pax$.

Consider the case where F is a facet. If the nearest point of F to p lies on a boundary segment of F, then the analysis of the previous paragraph gives the result. Otherwise, the nearest point to p lies inside some subfacet abc on F. Let q denote the nearest point to p. Note that q is the orthogonal projection of p onto abc. Let $\widehat{x}$ be the orthosphere of abc. The power diagram of $\widehat{a}$, $\widehat{b}$ and $\widehat{c}$ divides abc into three regions, each containing a vertex of abc. Assume that q lies inside the region containing a. Let H be the plane perpendicular to abc that contains a and p. Let $C_x = H \cap \widehat{x}$. Applying Lemma 4.4 to the cross-section of $\widehat{a}$ and $\widehat{x}$ on H, we obtain $\sin\angle pax \geq \text{radius}(C_x)/\sqrt{2\text{radius}(C_x)^2 + A^2}$. This bound attains its minimum when radius(C_x) is minimized. Since q lies inside abc, radius(C_x) is minimized when H contains b or c. That is, radius(C_x) is the orthoradius of ab or ac. In this configuration, the vertex gap property and Lemma 3.1 imply that $A \leq 2\omega_1 \text{radius}(C_x)/\sqrt{1-4\omega_1^2}$. This gives $\sin\angle pax \geq \sqrt{(1-4\omega_1^2)/(2-4\omega_1^2)} \geq \sqrt{(1-4\omega_1^2)/2}$. The distance from p to f is $\|a-p\|\sin\angle pax \geq r_p \sin\angle pax$. □

We can now relate r_x to the insertion radius of the parent of x.

LEMMA 4.6. *Suppose that the boundary gap and vertex gap properties hold. Let x be any vertex inserted or rejected. Let p be the parent of x, if it exists.*

1. *If x is an input vertex or p is an input vertex, then $r_x \geq c_1 \cdot f(x)$, where $c_1 = \sqrt{1-4\omega_1^2}$.*
2. *Otherwise, $r_x \geq \sqrt{1-\omega_2^2} \cdot f(x)$ or $r_x \geq c_2 \cdot r_p$, where $c_2 = (1-4\omega_1^2)/\sqrt{2}$ if x has type 1 or 2 and $c_2 = \rho_0$ if x has type 3.*

Proof. If x is an input vertex, then $r_x \geq f(x)$ by definition of local feature size. Suppose that the algorithm inserted x or rejected x by rule 1, 2, or 3. Let $\widehat{x}$ denote the corresponding orthosphere centered at x.

Consider the case where p is a vertex of P. If x has type 1 or 2, then x lies on a segment or facet of P that does not contain p. (The 1D or 2D weighted Delaunay triangulation of vertices on a segment or facet gives the subsegments and subfacets. So a vertex on a segment (facet) cannot encroach upon subsegments (subfacets) on the same segment (facet).) Starting with Lemma 4.3, we get $r_x \geq \sqrt{1-\omega_2^2}\|x-p\| \geq \sqrt{1-\omega_2^2} f(x)$. Suppose that x has

type 3. Let τ be the skinny tetrahedron split by x. By Lemma 3.1, $X \geq \sqrt{1-4\omega_1^2}R$, where R is the circumradius of τ. One endpoint of the shortest edge of τ is p. Let q denote the other endpoint. Since p is an input vertex, q is also an input vertex. This is because q must have existed before p by the definition of parent-child relation which can only happen if q is also an input vertex. Therefore, $P \leq \omega_1\|p-q\| \leq 2\omega_1 R$. Since $\widehat{x}$ and $\widehat{p}$ are orthogonal, $\|p-x\|^2 = X^2 + P^2 \leq X^2 + 4\omega_1^2R^2 \leq X^2/(1-4\omega_1^2)$. Similarly, $\|q-x\|^2 \leq X^2/(1-4\omega_1^2)$. We conclude that $r_x = X \geq \sqrt{1-4\omega_1^2}\max\{\|p-x\|, \|q-x\|\} \geq \sqrt{1-4\omega_1^2}f(x)$.

Consider the case where p is not a vertex of P. If x has type 3, then r_p is at most the shortest edge length of the tetrahedron split by x. So $r_x = X \geq \rho_0 r_p$. Suppose that x has type 1 or 2. There are three subcases depending on the types of p and x.

Case 1: p has type 1. Let e be the segment that contains the subsegment split by p. If x has type 1 and lies on a segment e' of P, then by Lemma 4.1, e and e' cannot be adjacent for p to be the parent of x. Similarly, if x has type 2 and lies on a facet F of P, then e cannot be adjacent to F by Lemma 4.2. We conclude that $f(x) \leq \|x-p\|$. Starting with Lemma 4.3, we get $r_x \geq \sqrt{1-\omega_2^2}\|x-p\| \geq \sqrt{1-\omega_2^2}f(x)$.

Case 2: p has type 2 and x has type 2. The facets that contain p and x cannot be adjacent by Lemma 4.2. So $f(x) \leq \|x-p\|$. It follows from Lemma 4.3 that $r_x \geq \sqrt{1-\omega_2^2}f(x)$.

Case 3: p has type 2 and x has type 1, or p has type 3. There are two subcases depending on whether p is weighted or not.

Case 3.1: p is unweighted. Let g be the subsegment or subfacet split by x. Let $\widehat{x}$ denote the smallest orthosphere of g. Since p is unweighted, p lies inside $\widehat{x}$ in order that p encroaches upon g. Let a be the vertex of g that has minimum power distance from p. If x has type 1, then let H be the plane containing g and p, otherwise let H be the plane perpendicular to g that contains a and p. Let y be the center of $\widehat{x} \cap H$. Let R denote radius$(\widehat{x} \cap H)$. Observe that g contains the projection of p onto the segment or facet containing g (if x has type 1, this is clearly true; otherwise, rule 2 enforces this condition). This observation implies that the angle $\angle pya$ on H is at most $\pi/2$. We have $\pi(p,\widehat{a}) = \|p-a\|^2 - A^2 \leq \|x-p\|^2 + \|x-a\|^2 - A^2$. Since the circle $\widehat{x} \cap H$ is orthogonal to the circle $\widehat{a} \cap H$, $\|x-a\|^2 - A^2 = R^2$. Since p lies inside $\widehat{x}$, $\|x-p\| \leq R$. It follows that $\pi(p,\widehat{a}) \leq 2R^2$. Since $R \leq X = r_x$ and $r_p^2 \leq \pi(p,\widehat{a})$, we conclude that $r_x \geq r_p/\sqrt{2} > (1-4\omega_1^2)r_p/\sqrt{2}$.

Case 3.2: p is weighted. Since p has type 2 and p is weighted, p was inserted by the algorithm. By Lemma 4.3, $r_x \geq \sqrt{1-\omega_2^2}\|x-p\|$. By Lemma 4.5, $\|x-p\| \geq r_p\sqrt{(1-4\omega_1^2)/2}$. It follows that $r_x \geq r_p\sqrt{(1-\omega_2^2)(1-4\omega_1^2)/2} \geq r_p(1-4\omega_1^2)/\sqrt{2}$.

□

4.4 Lower bound for insertion radius. In this section, we prove a lower bound of the insertion radius of a vertex x in terms of $f(x)$. This will allow us to prove a lower bound between any two vertices at any time in $\mathcal{V}$. In turn, this will allows us to prove in Section 5 several guarantees provided by weighted Delaunay refinement.

We first prove a relation involving local feature sizes and insertion radius.

LEMMA 4.7. *Suppose that the boundary and vertex gap properties hold. Let x be a vertex with parent p. If $r_x \geq c \cdot r_p$, then $f(x)/r_x \leq f(p)/(c \cdot r_p) + 1/\sqrt{\lambda}$, where $\lambda = 1-4\omega_1^2$.*

Proof. Let $\widehat{x}$ be the orthosphere centered at x before x was inserted. If x has type 3, then $r_x^2 = X^2 = \|x-p\|^2 - P^2$. By the vertex gap property, P is at most $2\omega_1$ times the circumradius of the tetrahedron split by x. So $P \leq 2\omega_1 X/\sqrt{1-4\omega_1^2}$ by Lemma 3.1. It follows that $r_x \geq \sqrt{1-4\omega_1^2}\|x-p\|$. Suppose that x has type 1 or type 2. If p was not rejected by the algorithm, then by Lemma 4.3, $r_x \geq \sqrt{1-\omega_2^2}\|x-p\|$ which is at least $\sqrt{1-4\omega_1^2}\|x-p\|$ as $\omega_2 \leq 2\omega_1$. If p was rejected by the algorithm, then p was unweighted and it lies inside $\widehat{x}$ giving $r_x = X > \|x-p\|$. We conclude that $r_x \geq \sqrt{\lambda}\|x-p\|$. Starting with the Lipschitz condition, we get

$$\begin{aligned} f(x) &\leq f(p) + \|x-p\| \\ &\leq \frac{f(p)}{c \cdot r_p} \cdot r_x + \frac{r_x}{\sqrt{\lambda}} \end{aligned}$$

which implies that $f(x)/r_x \leq f(p)/(c \cdot r_p) + 1/\sqrt{\lambda}$. □

We are ready to lower bound the insertion radius and inter-vertex distance. It is done in Lemma 4.8 using induction. To this end, we need the following constants.

$$\omega_1 = \frac{1}{2\sqrt{2}}$$

$$\omega_2 = \frac{1}{\sqrt{2}}$$

$$C_1 = \frac{\rho_0(2+\sqrt{2}\lambda+\lambda^2)}{\sqrt{\lambda}(\rho_0\lambda^2-2)}$$

$$C_2 = \frac{\sqrt{\lambda}(\sqrt{2}\rho_0+\lambda\rho_0+\sqrt{2})}{\rho_0\lambda^2-2}$$

$$C_3 = \frac{\sqrt{\lambda}(\lambda\rho_0+\lambda+\sqrt{2})}{\rho_0\lambda^2-2}$$

$$\omega_0 = \frac{1}{2\sqrt{2}(1+C_1)}$$

FACT 4.1. *If $\rho_0 > 8$, then we have $C_1 > C_2 > C_3 > 1/\sqrt{\lambda} > 1$ and $2\omega_1 = \omega_2 < 1$.*

So far we have proved results on inter-vertex distances, assuming that the vertex gap and boundary gap properties hold. Proof of these two properties in turn require these results. We break this circularity with induction in the next lemma which is the culmination of all previous results.

LEMMA 4.8. *Let x be a vertex of P or a vertex inserted, rejected, or pumped by the algorithm. We have the following invariants for $\rho_0 > 8$.*

1. *If x is a vertex of P or the parent of x is a vertex of P, then $r_x \geq \sqrt{\lambda} f(x) \geq f(x)/C_3$. Otherwise, if x has type i, for $1 \leq i \leq 3$, then $r_x \geq f(x)/C_i$.*
2. *If x is a vertex of P or x is inserted, then for any other vertex y that exists currently, $\|x-y\| \geq \max\{f(x), f(y)\}/(1+C_1)$.*
3. *The vertex and boundary gap properties hold after inserting, rejecting, or pumping x.*

Proof. We prove by induction. Invariants (1) and (3) hold before the algorithm starts (the basis case). Clearly, invariant (1) is not affected by pumping a vertex. So it suffices to prove invariant (1) where x is inserted or rejected by the algorithm. Let p be the parent of x. If p is a vertex of P, then $r_x \geq \sqrt{\lambda} f(x)$ by Lemma 4.6. Otherwise, assume inductively that invariant (1) holds for p. We conduct a case analysis. If x has type 3, then $r_x \geq \rho_0 \cdot r_p$ by Lemma 4.6. By induction we get $f(p) \leq C_1 r_p$ regardless of the type of p. By Lemma 4.7, we get

$$\frac{f(x)}{r_x} \leq \frac{C_1}{\rho_0} + \frac{1}{\sqrt{\lambda}} = C_3.$$

If x has type 2, then the proof of Lemma 4.6 reveals that $r_x \geq \sqrt{\lambda} f(x)$ if p has type 1 or 2 (we are done in this case) and $r_x \geq \lambda r_p/\sqrt{2}$ if p has type 3. If p has type 3, $f(p) \leq C_3 r_p$ by induction assumption. Then by Lemma 4.7, we get

$$\frac{f(x)}{r_x} \leq \frac{\sqrt{2}C_3}{\lambda} + \frac{1}{\sqrt{\lambda}} = C_2.$$

If x has type 1, then the proof of Lemma 4.6 reveals that $r_x \geq \sqrt{\lambda} f(x)$ if p has type 1 (we are done in this case) and $r_x \geq \lambda r_p/\sqrt{2}$ if p has type 2 or 3. If p has type 2 or 3, $f(p) \leq C_2 r_p$ by induction assumption. Then by Lemma 4.7, we get

$$\frac{f(x)}{r_x} \leq \frac{\sqrt{2}C_2}{\lambda} + \frac{1}{\sqrt{\lambda}} = C_1.$$

This proves that invariant (1) holds in general.

Consider invariant (2). Recall that x is a vertex of P or x is inserted right now. For any vertex y that exists currently, $\|x-y\| \geq r_x \geq f(x)/C_1 \geq f(x)/(1+C_1)$. Since $f(x) \geq f(y) - \|x-y\|$, we also have $\|x-y\| \geq f(y)/C_1 - \|x-y\|/C_1$ implying $\|x-y\| \geq f(y)/(1+C_1)$.

Consider invariant (3). It follows from invariant (2) that for any vertices u and v, $\|u-v\| \geq f(u)/(1+C_1) = 2\sqrt{2}\omega_0 f(u) = \omega_0 f(u)/\omega_1 \geq U/\omega_1$. This proves that the vertex gap property still holds. By Lemma 4.5, the distance from u to any segment or facet of P is at least $r_u\sqrt{(1-4\omega_1^2)/2} \geq f(u) \cdot 1/(1+C_1) \cdot \sqrt{(1-4\omega_1^2)/2} = \omega_0 f(u)/\omega_2 \geq U/\omega_2$. This proves that the boundary gap property still holds. □

Remark. Invariant (2) in Lemma 4.8 implies that for any two vertices x and y that are vertices of P or inserted by the algorithm, $\|x-y\| \geq \max\{f(x), f(y)\}/(1+C_1)$. This follows directly from invariant (2) if y is a vertex of P or y exists when x is inserted. If y is inserted later, then x exists when y is inserted. So we can reverse the roles of x and y to obtain the same bound by invariant (2).

5 Termination and Mesh Quality.

We establish four guarantees for QUALMESH. The first two guarantees, termination and gradedness, follow directly from Lemma 4.8. The absence of sliver and the size optimality are guaranteed using Lemma 4.8 and some other previous results [7, 13, 16].

THEOREM 5.1. (TERMINATION.) QUALMESH *terminates.*

Proof. It follows from Lemma 4.8 that any two vertices u and v at any stage of the algorithm must satisfy $||u-v|| \geq \max(f(u), f(v))/(1+C_1) \geq f_{min}/(1+C_1)$ where $f_{min} = \min_{x \in P} f(x)$ is a positive constant by definition. Since one can pack at most $24(1+C_1)^3 \text{vol}(D)/(4\pi f_{min}^3)$ disjoint balls inside a bounded domain D, the algorithm must terminate. □

THEOREM 5.2. (GRADEDNESS.) *The nearest neighbor distance of a vertex v in the output mesh is at least $\frac{f(v)}{1+C_1}$.*

Proof. This is only the restatement of (2) of Lemma 4.8. □

5.1 No sliver. Slivers incident to an unweighted vertex p are eliminated by pumping p. In the weight interval where p is pumped, tetrahedra incident to p change at discrete instances. A result in [7] implies that only constant number of tetrahedra can be incident to p over the entire interval of pumping. Choosing a proper parameter to define a tetrahedron as sliver, the intervals over which slivers remain incident to p can be made small enough so that there is a subinterval over which no sliver is incident to p. Algorithmically, we can use flips to generate new tetrahedra as pumping progresses and stop when no sliver is incident to p. We recast this idea of sliver exudation [7] in the context of weighted Delaunay triangulation.

Let Del $\mathcal{V}$ denote the weighted triangulation of $\mathcal{V}$. Del $\mathcal{V}$ has *orthoradius-edge property* $[\rho_0]$ if the ratio of orthoradius to shortest edge length of every tetrahedron in Del $\mathcal{V}$ is at most ρ_0. $\mathcal{V}$ has *weight property* $[\omega]$ for some $\omega \in (0, 1/2)$ if for each weighted vertex $\widehat{u} \in \mathcal{V}$, $U \leq \omega N(u)$, where $N(u)$ is the nearest neighbor Euclidean distance of u. Let $\mathcal{V}^*$ be $\mathcal{V}$ after pumping one or more unweighted vertices of $\mathcal{V}$. We state three results that are instrumental to proving the absence of slivers. Their proofs follow from the results in [7].

The first result states that if $\mathcal{V}^*$ has weight property $[\omega]$, then the orthoradius-edge ratio of each tetrahedron in Del $\mathcal{V}^*$ is still bounded by a constant.

LEMMA 5.1. *Assume* Del $\mathcal{V}$ *has orthoradius-edge property* $[\rho_0]$ *and weight property* $[\omega]$. *If* $\mathcal{V}^*$ *has weight property* $[\omega]$, *then* Del $\mathcal{V}^*$ *has orthoradius-edge ratio property* $[\rho]$ *for some constant* ρ.

The second and third results are concerned with the union of weighted Delaunay triangulations. Consider all possible weight assignments to unweighted vertices of $\mathcal{V}$ such that $\mathcal{V}^*$ has weight property $[\omega]$. Let K be the union of all edges in all the possible weighted Delaunay triangulation obtained.

The second result states that the incident edges of a vertex have similar lengths.

LEMMA 5.2. *If* Del $\mathcal{V}$ *has orthoradius-edge property* $[\rho_0]$, *then the lengths of any two adjacent edges in* K *is within a constant factor* ν_0.

The third result states that the degree of each vertex in K is bounded by a constant.

LEMMA 5.3. *If* Del $\mathcal{V}$ *has orthoradius-edge property* $[\rho_0]$, *then the degree of every vertex in* K *is bounded by some constant* δ_0.

The above results are generalizations of those proved in [7] for sliver exudation. In [7], all the points in $\mathcal{V}$ are unweighted and it is assumed that the radius-edge ratio of every tetrahedron in Del $\mathcal{V}$ is at most ρ_0. We generalize it so that $\mathcal{V}$ may contain weighted points, $\mathcal{V}$ has weight property $[\omega]$, and Del $\mathcal{V}$ has the orthoradius-edge property $[\rho_0]$. The results are the same except for the above changes.

We are ready to outline the proof that the output mesh contains no sliver.

THEOREM 5.3. (NO SLIVER.) *There is no sliver in the output of* QUALMESH.

Proof. (Sketch) First, we prove that if there is a sliver τ then τ must be incident to an unweighted vertex. Initially the assertion is true before any pumping. Consider an action by the algorithm which changes the triangulation. If it is pumping a vertex v, then only tetrahedra incident to v change and we claim that the new tetrahedra cannot be slivers due to pumping. If it is insertion of a new vertex v, then new tetrahedra and hence new slivers will be incident to v which is unweighted establishing the claim.

Next, we argue that pumping an unweighted vertex p can remove slivers incident to p. Assume that the shortest edge incident to p in K has length 1. When we pump p, we vary P^2 within the range $[0, \omega_0^2 f(p)^2]$. By invariant (2) of Lemma 4.8 and our assumption that the length of the shortest edge incident to p is 1, there exists a constant ω such that $\omega_0^2 f(p)^2 \leq \omega^2$. So the pumping preserves the weight property $[\omega]$. Let $pqrs$ be a sliver that we try to eliminate by pumping p. Note that q, r and s may be weighted. Figure 5(a) shows q, r and s lying on a horizontal plane, p above the plane, and the orthosphere of $pqrs$ centered at z before pumping. We assume that the orthocenter z lies below the plane containing q, r and s. When p is being

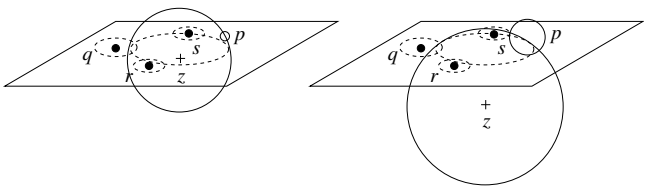

(a) (b)

Figure 5: Pumping action.

pumped, the orthocenter moves vertically downward and the orthoradius increases at the same time. Figure 5(b) shows the situation when p is being pumped. Let D denote the Euclidean distance of p from the plane containing qrs. Let $H(P)$ denote the Euclidean distance of the orthocenter z from the plane containing qrs as a function of the weight P of p. It is proved in [7] that $H(P) = H(0) - P^2/2D$.

Let Y be the radius of the smallest orthosphere of qrs. Since $pqrs$ is a sliver, Y is proportional to each side length of $pqrs$. By Lemma 5.1, for the sliver $pqrs$ to remain weighted

Delaunay when p is pumped, the orthosphere centered at z cannot become arbitrarily large. In particular, $|H(P)|$ should be at most $c \cdot Y$ for some constant c. Since $H(P) = H(0) - P^2/2D$, we conclude that $2D \cdot (H(0) - c \cdot Y) \leq P^2 \leq 2D \cdot (H(0) + c \cdot Y)$. So P^2 can only take on value within an interval of length $4cD \cdot Y$ for the sliver $pqrs$ to remain weighted Delaunay. We denote this interval by W_{qrs}. Since $pqrs$ is a sliver, the ratio of vol$(pqrs)$ to the shortest edge length of $pqrs$ is less than the threshold σ_0. By the orthoradius-edge property, this ratio is proportional to D/Y. So $D \leq c_1 \cdot \sigma_0 Y$ for some constant c_1. This implies that $|W_{qrs}| \leq c_2 \cdot \sigma_0 Y^2$ for some constant c_2.

Let $\mathcal{F}$ denote the family of triangles that form weighted Delaunay tetrahedron with p while p is pumped. By Lemma 5.3, $|\mathcal{F}| \leq \delta_0^3$. Let qrs be a triangle in $\mathcal{F}$ and let Y be the orthoradius of qrs. By the orthoradius-edge property $[\rho]$ in Lemma 5.1, $Y \leq \rho \cdot \|p - q\|$. Recall that the shortest edge incident to p has length 1. By Lemma 5.2, $\|p - q\| \leq \nu_0$. So $Y \leq \rho \cdot \nu_0$. It follows that for each $qrs \in \mathcal{F}$, $|W_{qrs}| \leq c_2 \rho^2 \nu_0^2 \cdot \sigma_0$. Thus, the sum of $|W_{qrs}|$ over triangles $qrs \in \mathcal{F}$ is at most $c_2 \rho^2 \nu_0^2 \delta_0^3 \cdot \sigma_0$.

Lastly, it can be shown that $f(p)$ is at least constant times the length of the shortest edge incident to p in K which is 1. Thus, the allowable range of P^2 is a superset of $[0, c_3 \omega_0^2]$ for some constant c_3. It follows that if we pick $\sigma_0 < c_3 \omega_0^2 / (c_2 \rho_1^2 \nu_0^2 \delta_0^3)$, then we can pick a value of P^2 such that slivers incident to p will be eliminated.

□

5.2 Size optimality. The size of our Delaunay mesh is within a constant factor of the size of any mesh that has bounded aspect ratio, provided that the outermost boundary of P coincides with the outermost boundary of the convex hull of P and a conforming tetrahedralization of the entire convex hull of P is needed. Our proof is a combination of ideas in Ruppert's proof for the 2D case [16] and ideas in Mitchell and Vavasis's proof for their octree algorithm in higher dimensions [15].

Let T be a conforming tetrahedralization of the convex hull of P with bounded aspect ratio. Let τ be a tetrahedron in T. Denote the minimum height of τ from a vertex by $h(\tau)$. Let v_0, v_1, v_2, and v_3 be the vertices of τ. Each v_i is to be viewed as a column vector consisting of the coordinates of the vertex. Define M_τ to be the 3×3 matrix $(v_1 - v_0, v_2 - v_0, v_3 - v_0)$. Although M_τ depends on the numbering of the vertices of τ, the numbering does not affect the properties of M_τ that will be used. As for vectors, we use the 2-norm as our matrix norm. The following technical result will be useful.

LEMMA 5.4. (LEMMA 2 IN [15]) *For any tetrahedron τ in T, $\|(M_\tau^{-1})^t\| = \|M_\tau^{-1}\| \leq k_1/h(\tau)$ for some constant k_1.*

Proof. It is proved in [15] that $\|M_\tau^{-1}\| \leq k_1/h(\tau)$ for some constant k_1. Since the transpose operation does not affect the value of a 2-norm, $\|(M_\tau^{-1})^t\| = \|M_\tau^{-1}\|$. □

Between two points on two disjoint elements of T (vertices, edges, triangles, or tetrahedra), we show in the following lemma that one can always find a tetrahedron that is relatively small compared with the distance between the two points.

LEMMA 5.5. (ANALOGOUS TO THEOREM 2 IN [15]) *Let p and q be two points in the interior of T. Let σ_p and σ_q be the elements of T of lowest dimension containing p and q respectively. If σ_p and σ_q do not share any vertex, then there is a tetrahedron τ in T intersecting pq such that*

1. *$h(\tau) \leq k_2 \|p - q\|$ for some constant k_2.*
2. *τ shares a vertex with σ_q. (τ can also be forced to share a vertex with σ_p instead, but τ cannot be guaranteed to share vertices with both σ_p and σ_q.)*

Proof. Define a simplicial map $\psi : T \to \mathbb{R}$ by setting $\psi(v) = 1$ for each vertex v of σ_q and $\psi(w) = 0$ for all other vertices w. It follows that $\psi(x) = 1$ for all point $x \in \sigma_q$ and $\psi(x) = 0$ for all point $x \in \sigma_p$. Since ψ is continuous, there exists a point u on pq such that the directional derivative of ψ at u has magnitude at least $1/\|p - q\|$. By convexity, there is a tetrahedron τ in T containing u. By linearity of ψ on τ, $\bigtriangledown\psi$ is constant on τ. Therefore,

$$\| \bigtriangledown \psi\| \geq \frac{1}{\|p - q\|} \tag{5.1}$$

on τ. This implies that τ shares a vertex with σ_q, otherwise ψ and $\bigtriangledown\psi$ would be identically zero on τ which is a contradiction. This proves (2). We express $\bigtriangledown\psi$ on τ using M_τ as follows. Let v_0, v_1, v_2, v_3 be the vertices of τ. Define $r_i = \psi(v_i) - \psi(v_0)$ for $i = 1, 2, 3$. Observe that $0 \leq |r_i| \leq 1$. For any point z in τ, one can check that $\psi(z) - \psi(v_0) = (r_1, r_2, r_3) M_\tau^{-1} (z - v_0)$ (see the appendix for a proof). This implies that $\bigtriangledown\psi = (M_\tau^{-1})^t (r_1, r_2, r_3)^t$ on τ and so

$$\begin{aligned} \| \bigtriangledown \psi\| &\leq \|(M_\tau^{-1})^t\| \cdot \|(r_1, r_2, r_3)^t\| \\ &\leq \sqrt{3} \cdot \|(M_\tau^{-1})^t\|, \quad \text{since } 0 \leq |r_i| \leq 1 \\ &\leq \sqrt{3} k_1 / h(\tau), \quad \text{by Lemma 5.4.} \end{aligned}$$

Combining the above with inequality 5.1, we obtain $h(\tau) \leq \sqrt{3} k_1 \|p - q\|$. This proves (1). □

Since T has bounded aspect ratio, it enjoys properties similar to that stated in Lemma 5.2. In particular, two tetrahedra sharing a vertex have similar minimum heights.

LEMMA 5.6. ([14, 15]) *If two tetrahedra τ_1 and τ_2 in T share a vertex, then $h(\tau_1) \leq k_3 h(\tau_2)$ for some constant k_3.*

Next, we prove that the minimum heights of tetrahedra in T change fairly smoothly.

LEMMA 5.7. (SIMPLIFIED VERSION OF LEMMA 11 IN [15]) *Let p and q be two points and let τ_p and τ_q be the two tetrahedra in T that contain them respectively. We have $h(\tau_q) \leq k_4 \max(h(\tau_p), \|p-q\|)$ for some constant k_4.*

Proof. Let σ be the subface of τ_q of the lowest dimension that contains q. If σ shares a vertex with τ_p, then τ_p and τ_q share a vertex and by Lemma 5.6, $h(\tau_q) \leq k_3 h(\tau_p)$. Consider the case where σ does not share a vertex with τ_p. By Lemma 5.5, there is a tetrahedron τ in T intersecting pq such that

$$h(\tau) \leq k_2 \|p-q\|, \tag{5.2}$$

and τ shares a vertex with σ. It follows that τ shares a vertex with τ_q. By Lemma 5.6, we have

$$h(\tau_q) \leq k_3 h(\tau). \tag{5.3}$$

Combining inequalities 5.2 and 5.3, we conclude that $h(\tau_q) \leq k_2 k_3 \|p-q\|$. Thus, it suffices to choose $k_4 = k_2 k_3$. □

The minimum heights of tetrahedra in T are also related to the local feature sizes.

LEMMA 5.8. (ANALOGOUS TO LEMMA 5 IN [16]) *Let x be a point and let τ be a tetrahedron in T containing x. Then $h(\tau) \leq k_5 f(x)$ for some constant k_5.*

Proof. Let B be the ball centered at x of radius $f(x)$. B contains two points p and q on two non-incident elements of P. By Lemma 5.5, there is a tetrahedron τ' in T intersecting pq such that $h(\tau') \leq k_2\|p-q\|$. Let u be a point in the intersection of τ' and pq. By applying Lemma 5.7 to u and x, we have

$$\begin{aligned} h(\tau) &\leq k_4 \max(h(\tau'), \|u-x\|) \\ &\leq k_4 \max(k_2\|p-q\|, \|u-x\|) \\ &\leq k_4 \max(k_2 \cdot 2f(x), f(x)) \\ &\leq 2k_2 k_4 f(x) \end{aligned}$$

□

We are now ready to prove the main theorem of this section.

THEOREM 5.4. (CLOSE TO OPTIMAL SIZE.) *The number of tetrahedra produced by* QUALMESH *is within a constant factor of the size of any mesh of bounded aspect ratio for the same domain under the convexity assumption.*

Proof. Let P denote the domain to be meshed. Let n be the number of vertices in the Delaunay mesh output by QUALMESH. First, we show that n is at most some constant times $\int_P \frac{dx}{f(x)^3}$. Since $N(v) \geq f(v)/(1+C_1)$, we can center a ball B_v of radius $d_v = f(v)/(2+2C_1)$ at each vertex v so that all balls are disjoint. Therefore,

$$\begin{aligned} \int_P \frac{dx}{f(x)^3} &\geq \sum_v \int_{B_v} \frac{dx}{f(x)^3} \\ &\geq \sum_v \frac{4\pi d_v^3}{3(f(v)+d_v)^3} \\ &\geq \sum_v \frac{4\pi}{3(3+2C_1)^3}. \end{aligned}$$

Therefore, we have

$$n = \sum_v 1 \leq \frac{3(3+2C_1)^3}{4\pi} \int_P \frac{dx}{f(x)^3}.$$

By Lemma 5.3 the number of tetrahedra in the Delaunay mesh is within a constant factor of n. More formally, let m be the number of tetrahedra, we have

$$m \leq k_6 \int_P \frac{dx}{f(x)^3},$$

for some constant k_6.

Next, consider a mesh T of P of bounded aspect ratio. For each point x in P, define $\ell(x)$ to be the minimum height of the tetrahedron in T containing x. By Lemma 5.8, $\ell(x) \leq k_5 f(x)$ and so

$$\begin{aligned} \int_P \frac{dx}{f(x)^3} &\leq k_5^3 \int_D \frac{dx}{\ell(x)^3} \\ &= k_5^3 \sum_{\tau \in T} \int_\tau \frac{dx}{h(\tau)^3} \\ &= k_5^3 \sum_{\tau \in T} \frac{\text{vol}(\tau)}{h(\tau)^3} \\ &\leq k_5^3 k_7 \sum_{\tau \in T} 1, \end{aligned}$$

since $\text{vol}(\tau) \leq k_7 h(\tau)^3$ for some constant k_7. Combining the above with the upper on m, we conclude that m is within a constant factor of the size of T. □

6 Conclusions.

A series of developments starting with Chew [4] and Ruppert [16], and continuing with Shewchuk [17] and Cheng et al. [7] brought the difficult problem of quality 3D Delaunay meshing of bounded domains close to the solution. Li and

Teng [12] recently developed a randomized point-placement strategy to generate a provably good 3D Delaunay mesh of bounded domains. This paper introduces a new paradigm, weighted Delaunay refinement, which gives the first deterministic algorithm for the problem.

Of course, as with previous algorithms the constants derived for the theory are miserably unsatisfactory for all practical purposes. For example, the constant $\rho_0 > 8$ is large for any practical purpose and the constant σ_0 is extremely small. Experiments show that these constants need not be that bad in practice when sliver exudation and Delaunay refinement are used separately [11]. Will the same remain true when we combine the two into our weighted Delaunay refinement algorithm? We plan future experiments to answer this question.

References

[1] N. Amenta, M. Bern, and D. Eppstein. Optimal point placement for mesh smoothing. *Proc. 8th Annu. ACM-SIAM Sympos. Discrete Alg.*, (1997), 528-537.

[2] M. Bern and D. Eppstein. Mesh generation and optimal triangulation. *Computing in Euclidean Geometry*, 2nd Ed., World Scientific, 1995, 47–123.

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

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

[5] L. P. Chew. Guaranteed-quality mesh generation for curved surfaces. *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, (1993), 274–280.

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

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

[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] H. Edelsbrunner, X.-Y. Li, G. L. Miller, A. Stathopoulos, D. Talmor, S.-H. Teng, A. Üngör and N. Walkington. Smoothing cleans up slivers. *Proc. 32nd Annu. ACM Sympos. Theory Comput.*, (2000), 273–277.

[11] H. Edelsbrunner and D. Guoy. An experimental study of sliver exudation. *Proc. 10th Intl. Meshing Roundtable*, (2001), 307–316.

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

[13] 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.

[14] G. L. Miller, S.-H. Teng, W. Thurston and S. Vavasis. Geometric separators for finite element meshes. *SIAM J. Sci. Comput.*, 19 (1998), 364–386.

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

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

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

[18] G. Strang and G. J. Fix. *An Analysis of the Finite Element Method.* Prentice Hall, 1973.

Appendix.

Let τ be a tetrahedron with vertices, v_0, v_1, v_2, v_3. Define $r_i = \psi(v_i) - \psi(v_0)$ for $i = 1, 2, 3$. We show that any point z in τ, $\psi(z) - \psi(v_0) = (r_1, r_2, r_3) M_\tau^{-1}(z - v_0)$. The point z can be written as a convex combination of the vertices of τ: $z = \sum_{i=0}^{3} \lambda_i v_i$. This implies that

$$z - v_0 = \sum_{i=1}^{3} \lambda_i (v_i - v_0).$$

We view M_τ^{-1} as three row vectors $\alpha_1, \alpha_2, \alpha_3$ ordered from top to bottom. Recall that $M_\tau = (v_1 - v_0, v_2 - v_0, v_3 - v_0)$. Since $M_\tau^{-1} M_\tau = I$, $\alpha_k \cdot (v_k - v_0) = 1$ and $\alpha_k \cdot (v_j - v_0) = 0$ for all $j \neq k$. It follows that

$$M_\tau^{-1}(z - v_0) = (\lambda_1, \lambda_2, \lambda_3)^t.$$

Hence, $(r_1, r_2, r_3) M_\tau^{-1}(z - v_0) = \sum_{i=1}^{3} \lambda_i (\psi(v_i) - \psi(v_0)) = (\sum_{i=0}^{3} \lambda_i \psi(v_i)) - \psi(v_0) = \psi(z) - \psi(v_0)$. This completes the proof.